

## MD INTERVIEW

# AN INTERVIEW WITH ROBERT AUMANN

*Interviewed by Sergiu Hart  
The Hebrew University of Jerusalem*

January 2005

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 also 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.

Address correspondence to: Sergiu Hart, 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>.



**FIGURE 1.** Bob Aumann, circa 2000.

**Sergiu HART:** Good morning, Professor Aumann. Well, I am not going to call you Professor Aumann. But what should I call you—Yisrael, Bob, Johnny?

**Robert AUMANN:** You usually call me Yisrael, so why don't you continue to call me Yisrael. But there really is a problem with my given names. I have at least three given names—Robert, John, and Yisrael. Robert and John are my given names from birth and Yisrael is the name that I got at the circumcision. Many people call me Bob, which is of course short for Robert. There was once a trivia quiz at a students' party at the Hebrew University, and one of the questions was, which faculty member has four given names and uses them all? Another story connected to my names is that my wife went to get approval of having our children included in her passport. She gave me the forms to sign on two different occasions. On one I signed Yisrael and on one I signed Robert. The clerk, when she gave him the forms, refused to accept them, saying, "Who is this man? Are there different fathers over here? We can't accept this."

**H:** I remember a time, when you taught at Tel Aviv University, you were filling out a form when suddenly you stopped and phoned your wife. "Esther," you asked, "what's my name in Tel Aviv?"

Let's start with your scientific biography, namely, what were the milestones on your scientific route?

**A:** 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 postdoc 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. I used to get up in the morning at 6:15, go to the university in uptown New York from Brooklyn—an hour and a quarter on the subway—then study calculus for an hour, then go back to the yeshiva on the lower east side for most of the morning, then go back up to City College at 139th Street and study there until 10 P.M., then go home and do some homework or whatever, and then I would get up again at 6:15. I did this for one semester, and then it became too much for me and I made the hard decision to quit the yeshiva and study mathematics.

**H:** How did you make the decision?

**A:** I really can't remember. I know the decision was mine. My parents put a lot of responsibility on us children. I was all of seventeen at the time, but there was no overt pressure from my parents. Probably math just attracted me more, although I was very attracted by Talmudic studies.

At City College, there was a very active group of mathematics students. The most prominent of the mathematicians on the staff was Emil Post, a famous logician. He was in the scientific school of Turing and Church—mathematical logic, computability—which was very much the “in” thing at the time. This was the late forties. Post was a very interesting character. I took just one course from him and that was functions of real variables—measure, integration, etc. The entire course consisted of his assigning exercises and then calling on the students to present the solutions on the blackboard. It's called the Moore method—no lectures, only exercises. It was a very good course. There were also other excellent teachers there, and 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. A schoolchild can understand what a prime number is, but understanding the distribution of prime numbers requires the theory of functions of a complex variable; it is closely related to the Riemann hypothesis, whose very formulation requires at least two or three years of university mathematics, and which remains unproved to this day. Another interesting aspect of number theory was that it was absolutely useless—pure mathematics at its purest.

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 prime numbers or 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 [1]; but the proof was essentially in place in the fall of 1954. Shortly thereafter, my research interests turned from knot theory to the areas that have occupied me to this day.

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:** This is indeed fascinating. Incidentally, has the “big, famous” problem ever been solved?

**A:** Yes. About a year after my thesis was published, a mathematician by the name of Papakyriakopoulos solved the general problem of asphericity. He had been working on it for eighteen years. He was at Princeton, but didn’t have a job there; they gave him some kind of stipend. He sat in the library and worked away on this for eighteen years! During that whole time he published almost nothing—a few related papers, a year or two before solving the big problem. Then he solved this big problem, with an amazingly deep and beautiful proof. And then, he disappeared from sight, and was never heard from again. He did nothing else. It’s like these cactuses that flower once in eighteen years. Naturally that swamped my result; fortunately mine came before his. It swamped it, except for one thing. Papakyriakopoulos’s result does not imply that alternating knots will not come apart. What he proved is that a knot that does not come apart is aspheric. What I proved is that *all* alternating knots are aspheric. It’s easy to see that a knot that comes apart is not aspheric, so it follows that an alternating knot will not

come apart. So that aspect of my thesis—which is the easily formulated part—did survive.

A little later, but independently, Dick Crowell also proved that alternating knots do not come apart, using a totally different method, not related to asphericity.

**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 postdoc, 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 Forrester 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.

**H:** You started reading game theory at that point?

**A:** I just did the minimum necessary of reading in order to be able to attack the problem.

**H:** Who were the game theorists at Princeton at the time? Did you have any contact with them?

**A:** I had quite a bit of contact with the Princeton mathematics department. Mainly at that time I was interested in contact with the knot theorists, who included John Milnor and of course R. H. Fox, who was the high priest of knot theory. But there was also contact with the game theorists, who included Milnor—who was both a knot theorist and a game theorist—Phil Wolfe, and Harold Kuhn. Shapley was already at RAND; I did not connect with him until later.

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.

Let me tell you an interesting anecdote. When I felt it was time to go on sabbatical, I started looking for a job, and made various applications. One was to Princeton—to Morgenstern. One was to IBM Yorktown Heights, which was also quite a prestigious group. I think Ralph Gomory was already the director of the math department there. Anyway, I got offers from both. The offer from IBM was for \$14,000 per year. \$14,000 doesn't sound like much, but in 1960 it was a nice bit of money; the equivalent today is about \$100,000, which is a nice salary for a young guy just starting out. Morgenstern offered \$7,000, exactly half. The offer from Morgenstern came to my office and the offer from IBM came home; my wife Esther didn't open it. I naturally told her about it and she said, "I know why they sent it home. They wanted *me* to open it."

I decided to go to Morgenstern. Esther asked me, "Are you sure you are not doing this just for *ipcha mistabra?*," which is this Talmudic expression for doing just the opposite of what is expected. I said, "Well, maybe, but I do think it's better to go to Princeton." Of course I don't regret it for a moment. It is at Princeton that I first saw the Milnor–Shapley paper, which led to the "Markets with a Continuum of Traders" [16], and really played a major role in my career; and I have no regrets over the career.

**H:** Or you could have been a main contributor to computer science.

**A:** Maybe, one can't tell. No regrets. It was great, and meeting Morgenstern and working with him was a tremendous experience, a tremendous privilege.

**H:** Did you meet von Neumann?

**A:** I met him, but in a sense, he didn't meet me. We were introduced at a game theory conference in 1955, two years before he died. I said, "Hello, Professor von Neumann," and he was very cordial, but I don't think he remembered me afterwards unless he was even more extraordinary than everybody says. I was a young person and he was a great star.

But Morgenstern I got to know very, very well. 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.



**FIGURE 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 [14]—

**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 and was quite prominent in the history of Israel. After the Yom Kippur War in 1973, he came to Israel and to Egypt to try to broker an arrangement between the two countries. He shuttled back and forth between Cairo and Jerusalem. When in Jerusalem, he stayed at the King David Hotel, which is acknowledged to be the best hotel here. Many people were appalled at what he was doing, and thought that he was exercising a lot of favoritism towards Egypt. One of these people was my cousin Steve Strauss, who was the masseur at the King David. Kissinger often went to get a massage from Steve. Steve told us that whenever Kissinger would, in the course of his shuttle diplomacy, do something particularly outrageous, he would slap him really hard on the massage table. I thought that Steve was kidding, but this episode appears also in Kissinger's memoirs; so there is another connection between game theory and the Aumann family.

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.

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 and we are again called upon to play a game—to be involved in an interactive situation—then the second time around 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 [4]. 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 [42]. The name has stuck. Incidentally, the term "folk theorem" is nowadays also used in other areas for classic results: the folk theorem of evolution, of computing, and so on.

**A:** The original Folk Theorem is quite similar to my '59 paper, but a good deal simpler, less deep. As you said, that became quite prominent in the later literature. 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 of the things that are most important are things that a mathematician would consider trivial.

**H:** I remember once in class that you got stuck in the middle of a proof. You went out, and then came back, thinking deeply. Then you went out again. Finally you came back some twenty minutes later and said, “Oh, it’s trivial.”

**A:** Yes, I got stuck and started thinking; the students were quiet at first, but got noisier and noisier, and I couldn’t think. I went out and paced the corridors and then hit on the answer. I came back and said, this is trivial; the students burst into laughter. So “trivial” is a bad term.

Take something like the Cantor diagonal method. Nowadays it would be considered trivial, and sometimes it really is trivial. But it is extremely important; for example, Gödel’s famous incompleteness theorem is based on it.

**H:** “Trivial to explain” and “trivial to obtain” are different. Some of the confusion lies there. Something may be very simple to explain once you get it. On the other hand, thinking about it and getting to it may be very deep.

**A:** Yes, and hitting on the right formulation may be very important. The diagonal method illustrates that 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 [4]. 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 repeated games, dynamic games, 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 ’64 they started working with the United States Arms Control and Disarmament Agency. Mike Maschler worked with them on the first project, which had to do with inspection; obviously there is a game between an inspector and an inspectee, who may want to hide what he is doing. Mike made an important contribution to that. There were other people working on that also, including Frank Anscombe. This started in ’64, and the second project, which was larger, started in ’65. It 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, 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 [v] 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:** You mentioned something else and I want to pick up on that: the Milnor–Shapley paper on oceanic games. That led you to another major work, “Markets with a Continuum of Traders” [16]: modeling perfect competition by a continuum.

**A:** As I already told you, 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 *International Economic Review* 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. This was published in *Econometrica* in 1964 [16].

**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 [23] on existence of competitive equilibria in continuum markets; in '75 came the paper on values of such markets, also in *Econometrica* [32]. Then there came later papers using a continuum, by me with or without coauthors [28, 37, 38, 39, 41, 44, 52], 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 nonatomic games [i]; this generated a huge literature. Many of your students worked on that. What's a nonatomic game, by the way? There is a story about a talk on "Values of nonatomic games," where a secretary thought a word



**FIGURE 3.** Werner Hildenbrand with Bob Aumann, Oberwolfach, 1982.

was missing in the title, so it became “Values of nonatomic war games.” So, what are nonatomic games?

**A:** It has nothing to do with war and disarmament. On the contrary, in war you usually have two sides. Nonatomic means the exact opposite, where you have a continuum of sides, a very large number of players.

**H:** None of which are atoms.

**A:** Exactly, in the sense that I was explaining before. It is like Milnor and Shapley’s oceanic games, except that in the oceanic games there were atoms—“large” players—and in nonatomic 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 nonatomic 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. Even the 2000 U.S. presidential election is a nonatomic game—no single voter, even in Florida, could have affected the outcome. (The people who did affect the outcome were the Supreme Court judges.) In a nonatomic 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 [32]. This result had several precursors, the first of which was a '64 RAND Memorandum of Shapley.

**H:** Values of nonatomic 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 (*Journal of Mathematical Economics*, '74 [29]). 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 [29, Examples 2.5 and 2.6]. 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 [29, Section 9j]. 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 [34]. 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 paper [53]. 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 talk now about the Hebrew University. You came to the Hebrew University in '56 and have been there ever since.

**A:** I'll tell you something. Mathematical game theory is a branch of applied mathematics. When I was a student, applied mathematics was looked down



upon by many pure mathematicians. They stuck up their noses and looked down upon it.

**H:** At that time most applications were to physics.

**A:** Even that—hydrodynamics and that kind of thing—was looked down upon. That is not the case anymore, and hasn't been for quite a while; but in the late fifties when I came to the Hebrew University that was still the vogue in the world of mathematics. At the Hebrew University I did not experience any kind of inferiority in that respect, nor in other respects either. Game theory was accepted as something worthwhile and important. In fact, Aryeh Dvoretzky, who was instrumental in bringing me here, and Abraham Fränkel (of Zermelo–Fränkel set theory), who was chair of the mathematics department, certainly appreciated this subject. It was one of the reasons I was brought here. Dvoretzky himself had done some work in game theory.

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

**A:** I don't know whether it was the brainchild of Yoram Ben-Porath or Menahem Yaari or both together. Anyway, Ben-Porath, who was the rector of the university, asked Yaari, Itamar Pitowsky, Motty Perry, and me to make a proposal for establishing a center for rationality. It wasn't even clear what the center was to be called. Something having to do with game theory, with economics, with philosophy. We met many times. Eventually 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, the Center for Operations Research and Econometrics in Louvain, Belgium, 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. We should have a member from political science, but we don't; that's a hole in the program. We should have one from medicine too, because medicine is a field in which rational utility-maximizing behavior is very important, and not at all easy. But at this time we don't have one. 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 for the readers of this journal? They 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, just like war is a grim game, and the stakes are high; but it is a game. That's another kind of interaction between computers and game theory.

Still another kind comes from computers that solve games, play games, and design games—like auctions—particularly on the Web. These are applications of computers to games, whereas before, we were discussing applications of games to computers.

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*. In economics people are assumed to act in order to maximize their utility; at least, until Tversky and Kahneman came along and said that people do not necessarily act in their self-interest. That is one way in which psychology is represented in the Center for Rationality: the study of irrationality. But the subject is still rationality. We'll discuss Kahneman and Tversky and the new school of "behavioral economics" later. Actually, using the term "behavioral economics" is already biasing the issue. The question is whether behavior really is that way or not.

We have mentioned computer science, psychology, economics, politics. 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. For example, the State of Israel is struggling with that. Also, I just came back from Paris, where Michel Balinsky told me about the problems of elections in American politics. There is apparently a tremendous amount of gerrymandering in American politics, and it's becoming a really big problem. So it is not only in Israel that we are struggling with the problem of how to conduct elections.

Another aspect is forming a government coalition: if it is too small—a minimal winning coalition—it will be unstable; if too large, the prime minister will have too little influence. What is the right balance?

**L:** 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. But when you come to think of it, it is a very complex statement. What does roundness mean? 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 or something that you cannot understand by rational means, an extra nonrationally driven player?

**A:** My response is that each player has to see to his own actions. In discussing the laws, the rules by which we live, the Talmud sometimes says that a certain action is not punishable by mortal courts but is punished by Heaven, and then discusses such punishments in detail. Occasionally in such a discussion somebody will say, well, we can only determine what the reaction of human courts will be to this or that action. We cannot dictate to Heaven how to react, and therefore it's useless for us to discuss it. That cuts off the discussion. As a religious person I must ask myself how *I* will act. I cannot discuss the rationality or irrationality of G-d.

**H:** The point is not the rationality or irrationality of that player, of G-d, but how that player affects what other players do and in what ways rational players can take this into account. Let me make it very simplistic. As you said, you don't know what Heaven will do, so how can I make rational decisions if I don't know that?

**A:** We don't know what Heaven will do, but we do have rules of conduct. We have the Pentateuch, the Torah, the Talmud.

**H:** I am talking more on the philosophical level, rather than on a practical level. The point is that that player is not reducible to standard mortal arguments

or understanding. Because if he were he would not be a special entity, which G-d is. However, he is part of the world. Not only is he part of the world, he is an important part of the religious world. He is not just a side player. He is the main player. Not only is he the main player, he is a player who by definition cannot be reduced to rational analysis.

**A:** I wouldn't say that He is irrational. By the way, it is interesting that this should come up just today, because there is a passage in the Torah reading of yesterday that relates to this. "This commandment that I command you today is not far away from you. It is not in Heaven so that one would have to say, 'Who will go up to Heaven and will take it from there and tell us about it?'" (*Deuteronomy* 30, 11–12). These verses were interpreted in the Talmud as saying that in the last analysis, commands in the Torah, the religious commandments, the whole of Scripture must be interpreted by human beings, by the sages and wise men in each generation. So the Torah must be given practical meaning by human beings.

The Talmud relates a story of a disagreement between one of the sages, Rabbi Eliezer ben Horkanos, and all the other sages. Rabbi Eliezer had one opinion and all the others had a different opinion. Rabbi Eliezer said, if I am right then let the water in the aqueduct flow upwards. Sure enough, there was a miracle, and the water started flowing uphill. So the other sages said, we are sorry, the law is not determined by the way the water flows in an aqueduct. It is determined by majority opinion. He asked for several other miracles and they all happened—Heaven was on his side. Nevertheless, his opinion was rejected. Each time the majority rejected it and said this is irrelevant. In the end he said, if I am right let a voice come from Heaven and say so. And sure enough, a voice came from Heaven and said, why do you argue with Rabbi Eliezer, whatever he says is always right. This was again rejected by the majority, who quoted the verse I just cited, "It is not in Heaven." The Torah was given to *us* by Heaven, and now it is *our* prerogative to interpret it. The story goes on to say that Elijah (the prophet who never died and keeps going back and forth between Heaven and earth) was asked by one of the sages who met him, were you in Heaven when that happened? He said, yes, I was there. The sage said, how did G-d react to his opinion being rejected by the earthly sages? So Elijah said, G-d smiled and said, "My children have vanquished me."

This is an example of what is 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 as human beings.

**H:** This is very interesting. 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 work with Michael Maschler [46], 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.

**H:** That’s a good example. In fact my wife has said many times, after yet another guest suddenly dropped in on us on Saturday, or we had to go and do something or other: “I wish we would become religious and have a really quiet Saturday once in a while.” So I can definitely understand the advantages of having a nice, quiet day of rest.

**A:** The day before the Sabbath, Friday, is a very hectic day for the person in charge of the house, who has to prepare for the Sabbath. On Friday in Israel, like on Saturday in most of the western world, many offices are closed. It is a semi-day-of-rest. But for religious people, especially for the houseperson, it is very hectic. We have a seminar series at the Center for Rationality called “Rationality on Friday”; my wife used to say that she could understand rationality on any other day, but *not* on Friday.

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. I don’t have to worry about these questions because I have a religion that tells me what to do. So I went to my rabbi—a Holocaust survivor, a very renowned, pious person. I figured he won’t even know what software is—I’ll have to explain it to him. Maybe there is a Talmudic rule about this kind of intellectual property not really being property. Whatever he’ll say, I’ll do. I went to him. He said, ask my son-in-law. So I said, no, I am asking *you*. He said, okay, come back in a few days. I’ll make a long story short. I went back again and again. He didn’t want to give me an answer. Finally I insisted and he said, “Okay, if you really want to know, it’s absolutely forbidden to do this, absolutely forbidden.” So I ordered the software.

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.

Unfortunately it has been my lot to spend more time in hospitals than I would have wanted. I have witnessed some very beautiful things. People coming to hospital wards and saying, look, we have private ambulances. We can take people from this hospital to wherever you want to go, from Metulla to Eilat (the northern and southern extremities of Israel), for nothing. We'll take anybody, religious, irreligious, Jews, Arabs, anybody. These were people who obviously were religious. They were going around with a beard and sidelocks. You have people who come around on Friday afternoon to make *kiddush* for the sick, and people who come around at any time of the week playing the violin and things like that.

The religious community, by the way, is very close. This 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.

**H:** Returning to the rules and their interpretation: do you mean that you would not go to the rabbi to ask him, say, whether to enter into a certain partnership, or how to vote in an election?

**A:** Well, I would not, and many others like me would not. But others—for example, "Khassidim"—might well consult their rabbi on such matters. In Khasidic circles, the rabbi is often much more than a scholar and legal and spiritual authority. He is a fountain of advice on all kinds of important decisions—medical, business, family, whatever. And often he gives very good advice! How come? Is he smarter than others? Yes, he often is. But that's not the important reason. The important reason is that everybody comes to him, so he gets a whole lot of inside information. We have a very interesting strategic equilibrium there—it's optimal for everyone to go to him, given that everybody goes to him! Of course, for that it is important that he be honest and straightforward, and that's already dictated

by the moral rules. But it's also part of the equilibrium, because the whole thing would fall apart if he weren't.

There is, incidentally, a phenomenon like this also among the "Mitnagdim," like me. There is a person in Israel called Rabbi Firer, who is absolutely the top source of medical information in the whole country, possibly in the whole world. And he is *not* a physician. Anybody who has an unusual or serious medical problem can go to him, or phone him. You make a phone appointment for, say, 1:17 A.M., you describe your problem, and he tells you where to go for treatment. Often the whole thing takes no more than a minute. Sometimes, in complicated cases, it takes more; he will not only direct you to a treatment center in Arizona, he'll arrange transportation when necessary, make the introductions, etc., etc. The whole point is that he is *not* a physician, so he has no special interests, no axe to grind. How it works is that he, like the Khassidic rabbi, gets information from everybody, patients and doctors alike, and he is also unusually brilliant. And he is deeply religious, which, again, is what keeps him honest. I have made use of him more often than I would have liked.

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.

Still another side is the "descriptive." Religion will not explain how evolution formed various species, or why competition works.

But I must immediately correct myself: the Talmud *does* in fact discuss both evolution and competition. Evolution is discussed in the tractate *Shabbat* on page 31a. The sage Hillel was asked why the eyes of certain African tribesmen are smaller than usual, and why the feet of other African tribesmen are broader than usual. Hillel's answers were adaptive: the eyes are smaller because these tribesmen live in a windy, sandy region, and the smallness of the eyes enables them better to keep the sand out; and the feet are broader because that tribe lives in a swampy region, and the broad feet enable easier navigation of the swamps.

Competition is also discussed in the Talmud. In the tractate *Baba Bathra* 89a, the Talmud says that the authorities must appoint inspectors to check the accuracy of the weights and measures used by marketplace vendors, but *not* to oversee prices. The twelfth-century commentator Samuel ben Meier (Rashbam) explains the reason: if a vendor overcharges, another vendor who needs the money will undercut him, all the customers will go to him, and the original vendor will have to match the lower price. The invisible hand—six hundred years before Adam Smith!

Other game-theoretic and economic principles are also discussed in the Talmud. The nucleolus makes an implicit appearance in the tractate *Kethuboth* 93a [46]; risk aversion shows up in *Makkoth* 3a [80]; moral hazard, in *Kethuboth* 15a, and the list can be made much longer.

But of course, all these discussions are only the barest of hints. We still need the game theory to understand these matters. The Talmud speaks about adaptation,



but one can hardly say that it anticipated the theory of evolution. The Talmud discusses competition, but we can hardly say that it anticipated the formulation of the equivalence theorem, to say nothing of its proof.

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?

Is your view a common view of religious people?

**A:** Maybe not. One doesn't discuss this very much in religious circles. When I was young, there were many attempts by religious people to “reconcile” science and religion. For example, each of the six days of creation can be viewed as representing a different geological era. There was—and perhaps still is—a view that science contradicts religion, that one has to reconcile them. It is apologetic, and I don't buy it.

**H:** Take, for example, the six days of creation; whether or not this is how it happened is practically irrelevant to one's decisions and way of conduct. It's on a different level.

**A:** It is a different view of the world, a different way of looking at the world. That's why I prefaced my answer to your question with the story about the roundness of the world being one way of viewing the world. An evolutionary geological perspective is one way of viewing the world. A different way is with the six days of creation. Truth is in our minds. If we are sufficiently broad-minded, then we can simultaneously entertain different ideas of truth, different models, different views of the world.

**H:** I think a scientist will have no problem with that. Would a religious person have problems with what you just said?

**A:** Different religious people have different viewpoints. Some of them might have problems with it. By the way, I'm not so sure that no scientist would have a problem with it. Some scientists are very doctrinaire.

**H:** I was just reminded of Newcomb's paradox, with its “omniscient being.” We both share the view that it doesn't make much sense. On the other hand, perhaps it does make sense in a religious world.

**A:** No, no. It's a little similar to this question of the omnipotence of G-d. If G-d is omnipotent, can he create an immovable object? Atheists will come up with a question like that, saying, here, I've disproved the whole idea of religion.

By the way, it's not a Jewish view that G-d is omnipotent. But that's not the point; the point is that the question is simply nonsense.

Altogether, the Jewish tradition is not very strong on theology, on what it is that G-d can or cannot do. But there is a very strong tradition of human free will in Judaism. There is definitely one thing that G-d *cannot* do, namely, influence a person's free will, his decision-making capacity. So there is a lack of omnipotence at least in that aspect of the Jewish tradition.

**H:** 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.

Let's move now to your personal biography.

**A:** I was born in 1930 in Frankfurt, Germany, to an orthodox Jewish family. My father was a wholesale textile merchant, rather well to do. 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 the emigration was cancelled and 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 and 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.

I am saying all this to illustrate 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.

Anyway, that was our story. We did get away in time, in 1938. We left Germany, and made our way to the United States; we got an immigration visa with some difficulty. 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. I went to Jewish parochial schools for my elementary education and also for high school. It is called a yeshiva high school, and combines Talmudic and other Jewish studies with secular studies. I have already mentioned my math teacher in high school, Joe Gansler. I also had excellent Talmud and Jewish studies teachers.

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 Lebanon in the 1982 Peace for Galilee operation. My other children are all happily married. Shlomo's widow also remarried and she is like a daughter to us. Shlomo had two children before he was killed (actually the second one was 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.



**FIGURE 4.** Bob Aumann with fiancée Esther Schlesinger, Israel, January 1955.

**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 postdoc 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



**FIGURE 5.** Bob Aumann with his immediate family, Jerusalem, October 10, 2005.

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, when the academic year in Israel starts. Could I work until October at Bell Labs?" They said, "Sure, we'll be glad to have you." That was very nice of them.

That was a really good four months there. John McCarthy, a computer scientist, was one of the people I got to know during that period. John Addison, a mathematician, logician, Turing machine person, was also there. One anecdote about Addison that summer is that he had written a paper about Turing machines, and wanted to issue it as a Bell Labs discussion paper. The patent office at Bell Labs gave him trouble. They wanted to know whether this so-called improvement on Turing machines could be patented. It took him a while to convince them that a Turing machine is not really a machine.

I am telling this long story to illustrate the difficulties with practical decision making. The process of practical decision making is 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, in contradistinction to what the theory says. But that's not what behavioral economics is concerned with. 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 “real” decisions rewarded by monetary payoffs. But even then the monetary payoff is usually very small. More importantly, the decisions that people face are not ones that they usually take, with which they are familiar. The whole setup is artificial. It is not a decision that really affects them and to which they are used.

Let me give one example of this—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, in these reports, how well the actual behavior of people fits economic theory.

**H:** Can you give an example of that?

**A:** One example I remember is where there was a very strong effect of raising the tax on alcohol by a very small amount, let’s say ten percent. Now we are talking about raising the price of a glass of beer by two to two and a half percent. It had a

very significant effect on the number of automobile accidents in the States. There is a tremendous amount of price elasticity in people's behavior. Another example is how increasing the police force affects crime.

**H:** Let's be more specific. Take the alcohol example. Why does it contradict the behavioral approach?

**A:** The conclusion of so-called behavioral economics is that people don't behave in a rational way, that they don't respond as expected to economic incentives. Empirical economics shows that people do respond very precisely to economic incentives.

**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.

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

**A:** Yes, it does bring us to rule rationality. The basic premise there is that people grope around. They learn what is a good way of behaving. When they make mistakes they adjust their behavior to avoid those mistakes. It is a learning process, not an explicit optimization procedure. This is actually an old idea. For example, Milton Friedman had this idea that people behave *as if* they were rational.

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.

**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 one hundred Deutsch Marks, which in the early eighties was equivalent to 150–200 Euros of today—a highly nonnegligible 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, once the proposed split appears on his

computer screen. 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 ninety-nine for the offerer and one for the responder.

**H:** The idea being that the responder would not leave even one Deutsch Mark 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 twenty Deutsch Marks; and the offerer usually anticipated this and therefore offered him more.

Walking away from twenty Deutsch Marks 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 Deutsch Mark when one hundred Deutsch Marks 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.

Of course, 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.

That is another example of rule rationality. 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. Also, one has to take into account that we cannot make precise computations and evaluate every decision. We need to develop rules that are relatively simple and applicable in many situations. Once we take into account this cost of computation, it is clear that a rule that is relatively simple, but gives a good answer in many situations, is better than a rule that requires you to go to extreme lengths to find the right answer for every decision.

**A:** That's the reason for it. You are giving the fundamental reason why people develop rules rather than optimize each act. It is simply too expensive.

**H:** 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. What you are saying is, 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:** This connects to another area in which you are involved quite a lot lately, namely, biology and evolutionary ecology. Do you want to say something about that?

**A:** The connection of evolution to game theory has been one of the most profound developments of the last thirty or forty years. It is one of the major developments since the big economic contributions of the sixties, which were mainly in cooperative game theory. It actually predates the explosion of noncooperative game theory of the eighties and nineties.

It turns out that there is a very, very strong connection between population equilibrium and Nash equilibrium—strategic equilibrium—in games. The same mathematical formulae appear in both contexts, but they have totally different interpretations. In the strategic, game-theoretic interpretation there are players and strategies, pure and mixed. In the two-player case, for every pair of strategies,

each player has a payoff, and there is a strategic equilibrium. In the evolutionary context, the players are replaced by populations, the strategies by genes, the probabilities in the mixed strategies by population proportions, and the payoffs by what is called fitness, which is a propensity to have offspring. You could have a population of flowers and a population of bees. There could be a gene for having a long nectar tube in the flowers and a gene for a long proboscis in the bees. Then, when those two meet, it is good for the flower and good for the bee. The bee is able to drink the nectar and so flits from flower to flower and pollinates them.

What does that mean, “good”? It means that both the flowers and the bees will have more offspring. The situation is in equilibrium if the proportions of genes of each kind in both populations are maintained. This turns out to be formally the same as strategic equilibrium in the corresponding game.

This development has had a tremendous influence on game theory, on biology, and on economic theory. It’s a way of thinking of games that transcends biology; it’s a way of thinking of what people do as traits of behavior, which survive, or get wiped out, just like in biology. It’s not a question of conscious choice. Whereas the usual, older interpretation of Nash equilibrium is one of conscious choice, conscious maximization. This ties in with what we were saying before, about rule rationality being a better interpretation of game-theoretic concepts than act rationality.

**H:** Perhaps it is time now to ask, 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.

Game theory started formally with the von Neumann and Morgenstern book in the 1940s. Probably the war had a lot to do with the fact that many people got interested. Just to see how it developed, in the first international game theory workshop in 1965 in Jerusalem there were seventeen people.

**A:** There were three conferences on game theory in Princeton in the fifties: ’53, ’55, and ’57. Those were attended by more than seventeen people. The seventeen people in 1965 were seventeen selected people.

**H:** The discipline 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.

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* [iv] we wrote that game theory may be viewed as a sort of umbrella or unified field 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.

**H:** Perhaps it is like mathematics applied to other sciences, which is a tool, a language for formalizing and analyzing.

**A:** That's an interesting analogy. Mathematics helps in certain aspects of many sciences—those given to formalization. Game theory is similar in that respect: it helps in many disciplines, specifically in their interactively rational parts. Figure 6 is a stylized representation.

**H:** The Game Theory Society was established in '99. You were the first, founding president, up to 2003. By now you should have a good overview of what game theory is, and of what the Game Theory Society is.

**A:** Game theory has become a big discipline, or rather a big *interdiscipline*. It is time to have a tool for gathering game theorists in all kinds of senses. Conferences, journals, the Web. When discussing my education, I mentioned that at City College there were a couple of tables reserved for the more dedicated math students. People would come between classes, sit down, have an ice cream soda, and talk about math. The Game Theory Society is the game theory table in the cafeteria that's called the world. It is a place where people can discuss game theory and exchange ideas, in various senses and various ways.

**H:** Do you have any thoughts on where game theory is going?

**A:** It is difficult to tell. It is very hard to know where things are going. In the presidential address at the Game Theory Society Congress in Bilbao in 2000 [81], I discussed some directions for research in the future.

Let me say something of a more general nature. People are pushing in different directions; we are going to find a spreading of the discipline among different people. Some people go in a very strongly mathematical direction, very deep mathematics. We will see a separation of the more mathematical branches from the more applied branches like economic applications. We'll see a lot of experimental and engineering application of game theory. People in game theory will understand each other less in the future.

**H:** Do you expect a Tower of Babel syndrome to develop?

**A:** It is not something that I would like, but it's a sign of maturity. Tower of Babel syndrome is a very good way of putting it.



**FIGURE 6.** The blooming of game theory.

**H:** What is definitely true is that from a small community where essentially everybody could understand everybody else, game theory has grown to a big “city,” where people are much more specialized. As in any developing discipline, it’s natural that everybody goes deeper into one of the aspects and understands less and less of the others. Nevertheless, at this point there is still interplay between the various aspects and approaches, so everybody benefits from everybody else. Take physics or mathematics. I wouldn’t understand what algebraic topology does nowadays. Somebody in combinatorics may understand something about probability, but wouldn’t understand some of the things we do in game theory. Nevertheless, mathematics is a single discipline.

Would you like to say anything about the different approaches in game theory? For example, mathematical vs. conceptual; axiomatic and cooperative vs. strategic and noncooperative. Why is it that there are so many approaches? Are they contradictory or are they just different? And how about people who think that some approaches in game theory are valid, and other approaches are not?

**A:** You are quite right that there is a group of people, working in noncooperative, strategic games, who think that cooperative (coalitional) game theory is less important, not relevant, not applicable.

Let me backtrack and describe what we mean by noncooperative or strategic game theory, vis-à-vis cooperative or coalitional game theory. Strategic game theory is concerned with strategic equilibrium—individual utility maximization given the actions of other people, Nash equilibrium and its variants, correlated equilibrium, that kind of thing. It asks how people should act, or do act. Coalitional game theory, on the other hand, concentrates on division of the payoff, and not so much on what people do in order to achieve those payoffs.

Practically speaking, strategic game theory deals with various equilibrium concepts and is based on a precise description of the game in question. Coalitional game theory deals with concepts like the core, Shapley value, von Neumann–Morgenstern solution, bargaining set, nucleolus. Strategic game theory is best suited to contexts and applications where the rules of the game are precisely described, like elections, auctions, internet transactions. Coalitional game theory is better suited to situations like coalition formation or the formation of a government in a parliamentary democracy or even the formation of coalitions in international relations; or, what happens in a market, where it is not clear who makes offers to whom and how transactions are consummated. Negotiations in general, bargaining, these are more suited for the coalitional, cooperative theory.

**H:** On the one hand, negotiations can be analyzed from a strategic viewpoint, if one knows exactly how they are conducted. On the other hand, they can be analyzed from a viewpoint of where they lead, which will be a cooperative solution. There is the “Nash program”—basing cooperative solutions on noncooperative implementations. For example, the alternating offers bargaining, which is a very natural strategic setup, and leads very neatly to the axiomatic solution of Nash—as shown by Rubinstein and Binmore.

**A:** These “bridges” between the strategic and the coalitional theory show that these approaches are not disparate. In order to make a bridge like that you have to define precisely the noncooperative situation with which you are dealing. One of the bridges that we discussed earlier in this interview is the Folk Theorem for repeated games. There, the noncooperative setup is the repeated game. When you have a bridge like that to the noncooperative theory, the strategic side must be precisely defined. The big advantage of the cooperative theory is that it does *not* need a precisely defined structure for the actual game. It is enough to say what each coalition can achieve; you need not say *how*. For example, in a market context you say that each coalition can exchange among its own members whatever it wants. You don’t have to say how they make their offers or counteroffers. In a political context, it is enough to say that any majority of parliament can form a government. You don’t have to say how they negotiate in order to form a government. That already defines the game, and then one can apply the ideas of the coalitional theory to make some kind of analysis, some kind of prediction.

You asked about the sociology of game theorists, rather than game theory. There is a significant group of people in strategic game theory who have an attitude towards coalitional game theory similar to that of pure mathematicians towards applied mathematics fifty years ago. They looked down their noses and

said, this is not really very interesting; we're not going to sully our hands with this stuff.

There is no justification for this in the game-theoretic sociology, just as there was no justification for it in the mathematics sociology. Each one of these branches of the discipline makes its contribution. In many ways, the coalitional theory has done better than the strategic theory in giving insight into economic and other environments. A prime example of this is the equivalence theorem, which gave a game-theoretic foundation for the law of supply and demand. There has been nothing of that generality or power in strategic game theory. Strategic game theory has made important contributions to the analysis of auctions, but it has not given that kind of insight into economics, or into any other discipline.

Another example of an important insight yielded by coalitional game theory is the theory of matching markets. This whole branch of game theory—and it is highly applied—grew out of the '62 paper of Gale and Shapley “College Admissions and the Stability of Marriage.” It is not quite as fundamental as the equivalence theorem, but it is a very important application, certainly of comparable importance to the work on auctions in strategic game theory, which is very important. There is no reason to denigrate the contributions of coalitional game theory, either on the applied or the theoretical level.

**H:** Indeed, Adam Brandenburger said that his students at Harvard Business School found cooperative game theory much more relevant to them than the noncooperative theory.

Let's switch to another topic. You have had an enormous impact on the profession by influencing many people. I am talking first of all about your students. By now you have had thirteen doctoral students. I think twelve of them are by now professors, in Israel and abroad, who are well recognized in the field and also in related fields.

**A:** Almost all the students eventually ended up in Israel, after a short break for a postdoc or something similar abroad.

**H:** That's not surprising since most of them—all except Wesley—started in Israel and are Israelis.

**A:** There is quite a brain drain from Israel. A large proportion of prominent Israeli scientists who are educated in Israel end up abroad—a much larger proportion than among my students.

These are my doctoral students up until now: 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. Of these, three are currently abroad—Kohlberg, Wesley, and 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 good. I don't mean good morally, but capable as scientists and specifically as mathematicians. All of my students came from mathematics. In most cases I knew them from my classes. In some cases not, and then I looked carefully at their grades and accepted only the very best. I usually worked quite closely with them, meeting once a week or so at least, hearing about progress, making suggestions, asking questions. When the final thesis was written I very often didn't read it carefully. Maybe this is news to Professor Hart, maybe it isn't. But by that time I knew the contents of the work because of the periodic meetings that we would have.

**H:** Besides, you don't believe anything unless you can prove it to yourself.

**A:** I read very little mathematics—only when I need to know. Then, when reading an article I say, well, how does one prove this? Usually I don't succeed, and then I look at the proof.

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 is a lot to say here, but it's your interview, so I am making it very short. There are many stories among your students, who are still very close to one another.



**FIGURE 7.** At the GAMES 1995 Conference in honor of Aumann's sixty-fifth birthday, Jerusalem, June 1995: Abraham Neyman, Bob Aumann, John Nash, Reinhard Selten, Ken Arrow, Sergiu Hart (from left to right).

Next, how about your collaborators? Shapley, Maschler, Kurz, and Drèze are probably your major collaborators. Looking at your publications I see many other coauthors—a total of twenty—but usually they are more focused on one specific topic.

**A:** I certainly owe a lot to all those people. Collaborating with other people is a lot of work. It makes things a lot more difficult, because each person has his own angle on things and there are often disagreements on conceptual aspects. It's not like pure mathematics, where there is a theorem and a proof. There may be disagreements about which theorem to include and which theorem not to include, but there is no room for substantive disagreement in a pure mathematics paper. Papers in game theory or in mathematical economics have large conceptual components, on which there often is quite substantial disagreement between the coauthors, which must be hammered out. I experienced this with all my coauthors.

You and I have written several joint papers, Sergiu. There wasn't too much disagreement about conceptual aspects there.

**H:** The first of our joint papers [50] was mostly mathematical, but over the last one [82] there was some . . . perhaps not disagreement, but clarification of the concepts. The other two papers [69, 70], together with Motty Perry, involved a lot of discussion. I can also speak from experience, having collaborated with other people, including some longstanding collaborations. Beyond mathematics, the arguments are about identifying the right concept. This is a question of judgment; one cannot prove that this is a good concept and that is not. One can only have a feeling or an intuition that *that* may lead to something interesting, that studying *this* may be interesting. Everybody brings his own intuitions and ideas.

**A:** But there are also sometimes real substantive disagreements. There was a paper with Maschler—"Some Thoughts on the Minimax Principle" [27]—where we had diametrically opposed opinions on an important point that could not be glossed over. In the end we wrote, "Some experts think A, others think 'Not A'." That's how we dealt with the disagreement. Often it doesn't come to that extreme, but there *are* substantial substantive disagreements with coauthors. Of course these do not affect the major message of the paper. But in the discussion, in the conceptualization, there are nuances over which there are disagreements. All these discussions make writing a joint paper a much more onerous affair than writing a paper alone. It becomes much more time-consuming.

**H:** But it is time well consumed; having to battle for your opinion and having to find better and better arguments to convince your coauthor is also good for your reader and is also good for really understanding and getting much deeper into issues.

That is one reason why an interdisciplinary center is so good. When you must explain your work to people who are outside your discipline, you cannot take anything for granted. All the things that are somehow commonly known and commonly accepted in your discipline suddenly become questionable. Then you realize that in fact they shouldn't be commonly accepted. That is a very



good exercise: explain what you are doing to a smart person who has a general understanding of the subject, but who is not from your discipline. It is one of the great advantages of our Rationality Center. A lot of work here has been generated from such discussions. Suddenly you realize that some of the basic premises of your work may in fact be incorrect, or may need to be justified. The same goes for collaborators. When you think by yourself, you gloss over things very quickly. When you have to start explaining it to somebody, then you have to go very slowly, step by step, and you cannot err so easily.

**A:** That's entirely correct, and I'd like to back it up with a story from the Talmud. A considerable part of the Talmud deals with pairs of sages, who consistently argued with each other; one took one side of a question and the other took the other side. One such pair was Rabbi Yochanan and Resh Lakish. They were good friends, but also constantly taking opposite sides of any given question. Then Resh Lakish died, and Rabbi Yochanan was inconsolable, grieved for many days. Finally he returned to the study hall and resumed his lectures. Then, for everything that Rabbi Yochanan said, one of the sages adduced thirty pieces of supporting evidence. Rabbi Yochanan broke down in tears and said, what good are you to me? You try to console me for the loss of Resh Lakish, but you do exactly the opposite. Resh Lakish would come up with thirty challenges to everything I said, thirty putative proofs that I am wrong. Then I would have to sharpen my wits and try to prove that he is wrong and thereby my position would be firmly established. Whereas you prove that I'm right. I know that I'm right; what good does it do that you prove that I am right. It doesn't advance knowledge at all.

This is exactly your point. When you have different points of view and there is a need to sharpen and solidify one's own view of things, then arguing with someone makes it much more acceptable, much better proved.

With many of my coauthors there were sharp disagreements and very close bargaining as to how to phrase this or that. I remember an argument with Lloyd Shapley at Stanford University one summer in the early seventies. I had broken my foot in a rock-climbing accident. Shapley came to visit me in my room at the Stanford Faculty Club, and I was hobbling around on crutches. This is unbelievable, but we argued for a full half hour about a comma. I don't remember whether I wanted it in and Lloyd wanted it out, or the other way around. Neither do I remember how it was resolved. It would not have been feasible to say "some experts would put a comma here, others would not." I always think that my coauthors are stubborn, but maybe I am the stubborn one.

I will say one thing about coauthorship. Mike Maschler is a wonderful person and a great scientist, but he is about the most stubborn person I know. One joint paper with Maschler is about the bargaining set for cooperative games [17]. The way this was born is that in my early days at the Hebrew University, in 1960, I gave a math colloquium at which I presented the von Neumann–Morgenstern stable set. In the question period, Mike said, I don't understand this concept, it sounds wrongheaded. I said, okay, let's discuss it after the lecture. And we did.

I tried to explain and to justify the stable set idea, which is beautiful and deep. But Mike wouldn't buy it. Exasperated, I finally said, well, can you do better? He said, give me a day or two. A day or two passes and he comes back with an idea. I shoot this idea down—show him why it's no good. This continues for about a year. He comes up with ideas for alternatives to stable sets, and I shoot them down; we had well-defined roles in the process. Finally, he came up with something that I was not able to shoot down with ease. We parted for the summer. During that summer he wrote up his idea and sent it to me with a byline of Robert Aumann and Michael Maschler. I said, I will have no part of this. I can't shoot it down immediately, but I don't like the idea. Maschler wouldn't take no for an answer. He kept at me stubbornly for weeks and months and finally I broke down and said, okay, I don't like it, but go ahead and publish it. This is the original "Bargaining Set for Cooperative Games" [17]. I still don't like that idea, but Maschler and Davis revised it and it eventually became, with their revision, a very important concept, out of which grew the Davis–Maschler kernel and Schmeidler's nucleolus. Because of where it led more than because of what it is, this became one of my most cited papers. Maschler's stubbornness proved justified. Maybe it should have waited for the Davis–Maschler revision in the first place, but anyway, in hindsight I'm not sorry that we published this. Michael has always been extremely stubborn. When he wants something, it gets done. As you say, Sergiu, coauthorship is much more exacting, much more painful than writing a paper alone, but it also leads to a better product.

**H:** This very naturally leads us to what you view as your main contributions. And, what are your most cited papers, which may not be the same thing.

**A:** One's papers are almost like one's children and students—each one is different, one loves them all, and one does not compare them. Still, one does keep abreast of what they're doing; so I also keep an eye on the citations, which give a sense of what the papers are "doing."

One of the two most cited papers is the Equivalence Theorem—the "Markets with a Continuum of Traders" [16]—the principle that the core is the same as the competitive equilibrium in a market in which each individual player is negligible. The other one is "Agreeing to Disagree" [34], which initiated "interactive epistemology"—the formal theory of knowledge about others' knowledge. After that come the book with Shapley, *Values of Non-Atomic Games* [i], the two papers on correlated equilibrium [29, 53], the bargaining set paper with Maschler [17], the subjective probability paper with Anscombe [14], and "Integrals of Set-Valued Functions" [21], a strictly mathematical paper that impacted control theory and related areas as well as mathematical economics. The next batch includes the repeated games work—the '59 paper [4], the book with Maschler [v], the survey [42], and the paper with Sorin on "Cooperation and Bounded Recall" [57]; also, the Talmud paper with Maschler [46], the paper with Drèze on coalition structures [31], the work with Brandenburger on "Epistemic Conditions for Nash Equilibrium" [65], the "Power and Taxes" paper with Kurz [37], some of the papers on NTU-games [10, 24], and others.

That sort of sums it up. Correlated equilibrium had a big impact. The work on repeated games, the equivalence principle, the continuum of players, interactive epistemology—all had a big impact.

Citations do give a good general idea of impact. But one should also look at the larger picture. Sometimes there is a body of work that all in all has a big impact, more than the individual citations show. In addition to the above-mentioned topics, there is incomplete information, NTU-values and NTU-games in general—with their many applications—perfect and imperfect competition, utilities and subjective probabilities, the mathematics of set-valued functions and measurability, extensive games, and others. Of course, these are not disjoint; there are many interconnections and areas of overlap.

There is a joint paper with Jacques Drèze [51] on which we worked very, very hard, for very, very long. For seven years we worked on it. It contains some of the deepest work I have ever done. It is hardly cited. This is a paper I love. It is nice work, but it hasn't had much of an impact.

**H:** Sometimes working very hard has two bad side effects. One is that you have solved the problem and there is nothing more to say. Two, it is so hard that nobody can follow it; it's too hard for people to get into.

We were talking about various stations in your life. Besides City College, MIT, Princeton, and Hebrew University you have spent a significant amount of time over the years at other places: Yale, Stanford, CORE, and lately Stony Brook.

**A:** Perhaps the most significant of all those places is Stanford and, specifically, the IMSSS, the Institute for Mathematical Studies in the Social Sciences—Economics. This was run by Mordecai Kurz for twenty magnificent years between 1971 and 1990. The main activity of the IMSSS was the summer gatherings, which lasted for six to eight weeks. They brought together the best minds in economic theory. A lot of beautiful economic theory was created at the IMSSS. The meetings were relaxed, originally only on Tuesdays and Thursdays, with the whole morning devoted to one speaker; one or two speakers in the afternoon, not more. A little later, Wednesday mornings also became part of the official program. All the rest of the time was devoted to informal interaction between the participants. Kenneth Arrow was a fixture there. So was Frank Hahn. Of course, Mordecai. I came every year during that period.

It was an amazing place. Mordecai ran a very tight ship. One year he even posted guards at the doors of the seminar room to keep uninvited people out. But he himself realized that that was going a little far, so that lasted only that one summer.

Another anecdote from that period is this: the year after Arrow got the Nobel Prize, he was vacationing in Hawaii at the beginning of July, and did not turn up for the first session of the summer. Mordecai tracked him down, phoned him and said, Kenneth, what do you think you are doing? You are supposed to be here; get on the next plane and come down, or there will be trouble. The audacity of the request is sufficiently astounding, but even more so is that Arrow did it. He cancelled the rest of his vacation and came down and took his seat in the seminar.

The IMSSS was tremendously influential in the creation of economic theory over those two decades. And it was also very influential in my own career. Some of my best work was done during those two decades—much of it with very important input from the summer seminar at the IMSSS. Also, during those two decades I spent two full sabbaticals at Stanford, in '75–'76 and in '80–'81. This was a very important part of my life. My children used to say that California is their second home. Being there every summer for twenty years, and two winters as well, really enabled me to enjoy California to the fullest. Later on, in the nineties, we were again at Stanford for a few weeks in the summer. I told my wife there was a friend whom I hadn't seen in a year. She said, who, and I said, the Sierra Nevada, the mountains. We had been there a few weeks and we hadn't gone to the mountains yet. We went, and it was a beautiful day, as always. Many times during those years we would get up at 3 or 4 in the morning, drive to eastern California, to the beautiful Sierra mountains, spend the whole day there from 7 or 8 A.M. until 9 P.M., and then drive back and get to Palo Alto at 1 A.M.; exhausted, but deeply satisfied. We climbed, hiked, swam, skied.

The Sierra Nevada is really magnificent. I have traveled around the whole world, and never found a place like it, especially for its lakes. There are grander mountains, but the profusion and variety of mountain lakes in the Sierra is unbelievable. I just thought I would put that in, although it has nothing to do with game theory.

**H:** Getting back to the IMSSS summers: besides those who came every year, there were always a few dozen people, from the very young who were in the advanced stages of their doctoral studies, to very senior, established economists. People would present their work. There would be very exciting discussions. Another thing: every summer there were one or two one-day workshops, which were extremely well organized, usually by the very senior people like you; for example, you organized a workshop on repeated games in 1978 [42]. One would collect material, particularly material that was not available in print. One would prepare notes. They were duplicated and distributed to everybody there. They served for years afterwards as a basis for research in the area. I still have notes from those workshops; they were highly influential.

In all the presentations, you couldn't just come and talk. You had to prepare meticulously, and distribute the papers and the references. The work was serious and intensive, and it was very exciting, because all the time new things were happening. It was a great place.

**A:** You are certainly right—I forgot to mention all the other people who were there, and who varied from year to year. Sometimes people came for two or three or four consecutive years. Sometimes people came and then didn't come the next summer and then came again the following summer. But there was always a considerable group of people there who were contributing, aside from the three or four “fixtures.”

Another point is the intensity of the discussion. The discussion was very free-wheeling, very open, often very, very aggressive. I remember one morning I was

supposed to give a two-hour lecture. The lectures were from 10 to 11, then a half-hour break, and then 11:30 to 12:30. I rose to begin my presentation at 10 in the morning, and it wasn't more than one minute before somebody interjected with a question or remark. Somebody else answered, and pandemonium broke loose. This lasted a full hour, from 10 to 11. After a few minutes I sat down and let the people argue with each other, though this was supposed to be my presentation. Then came the break. By 11:30 people had exhausted themselves, and I gave my presentation between 11:30 and 12:30. This was typical, though perhaps a little unusual in its intensity.

**H:** That was typical, exactly. There was no such thing as a twenty-minute grace period. There was no grace whatsoever. On the other hand, the discussions were really to the point. People were trying to understand. It was really useful. It clarified things. If you take those twenty years, probably a significant part of the work in economic theory in those years can be directly connected to the Stanford summer seminar. It originated there. It was discussed there. It was developed there in many different directions. There was nothing happening in economic theory that didn't go through Stanford, or was at least presented there.

**A:** We should move on perhaps to CORE, the Center for Operations Research and Econometrics at the Catholic University of Louvain, an ancient university, about seven or eight hundred years old. CORE was established chiefly through the initiative of Jacques Drèze. I was there three or four times for periods of several months, and also for many shorter visits. This, too, is a remarkable research institution. Unlike the IMSSS, it is really most active during the academic year. It is a great center for work in economic theory and also in game theory. The person I worked with most closely throughout the years—and with whom I wrote several joint papers—is Jacques Drèze. Another person at CORE who has had a tremendous influence on game theory, by himself and with his students, is Jean-François Mertens. Mertens has done some of the deepest work in the discipline, some of it in collaboration with Israelis like my students Kohlberg, Neyman, and Zamir; he established a Belgian school of mathematical game theory that is marked by its beauty, depth, and sophistication.

Another institution with which I have been associated in the last ten or fifteen years is the Center for Game Theory at Stony Brook. The focus of this center is the summer program, which lasts two or three weeks, and consists of a large week-long international conference that covers all of game theory, and specialized workshops in various special areas—mostly quite applied, but sometimes also in special theoretical areas. The workshops are for smaller groups of people, and each one is three days, four days, two days, whatever. This program, which is extremely successful and has had a very important effect on game theory, has been run by Yair Tauman ever since its inception in '91. In the past I also spent several periods of several months each during the academic year teaching game theory or doing research in game theory there with a small group of top researchers and a small group of graduate students; that's another institution with which I've been associated.

I should also mention Yale, where I spent the '64–'65 academic year on sabbatical. This was after publication of the work with Frank Anscombe, "A Definition of Subjective Probability" [14]; Frank was the chairman of the statistics department at Yale. At that time I was also associated with the Cowles Foundation; Herb Scarf and Martin Shubik were there. A very unique experience was the personal friendship that I struck up with Jimmy Savage during that year. I don't know how many people know this, but he was almost totally blind. Almost—not quite. He could read with great difficulty, and tremendous enlargement. Looking at his work there is no hint of this. I again spent about six weeks at Yale in the late eighties at the Cowles Foundation, giving a series of lectures on interactive epistemology.

One more place that influenced me was Berkeley, where I spent the summer of '64 and the spring of '72. There the main contact was Gerard Debreu, who was a remarkable personality. Other people there were John Harsanyi and Roy Radner. In addition to his greatness as a scientist, Gerard was also well known as a gourmet. His wife Françoise was a terrific cook. Once in a while they would invite us to dinner; Françoise would go out of her way to prepare something kosher. Occasionally we would invite them. It was his practice at a meal to praise at most one dish. Sometimes he praised nothing; sometimes, one dish. That totally transformed a compliment from Gerard from something trivial to something sublime. Nowadays, I myself cook and give dinners; when a guest leaves saying everything was wonderful, it means nothing. Though I allow myself to be kidded, it really means nothing. But when a guest leaves and says, the soup was the most delicious soup I ever had, that says something. He doesn't talk about the meat and not about the fish and not about the salad and not about the dessert, just the soup. Or somebody else says, this was a wonderful trout mousse. One dish gets praised. Then you know it's meaningful.

I also spent a month at NYU, in February of 1997. It was interesting. But for me, the attractions of New York City overwhelmed the academic activity. Perhaps Esther and I took the city a little too seriously. This was a very beautiful time for us, but what surrounds NYU was more important to us than the academic activity.

**H:** Maybe it's a good point to ask you, in retrospect, who are the people who have most influenced your life?

**A:** First of all my family: parents, brother, wife, children, grandchildren. My great-grandchild has not yet had a specific important influence on me; he is all of one and a half. But that will come also. My students have influenced me greatly. You have influenced me. All my teachers. Beyond that, to pick out one person in the family, just one: my mother, who was an extraordinary person. 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 Schubert 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 [66]. She always encouraged, always pushed us along, gently, unobtrusively, always allowed us to make our own decisions. Of course parents always have an influence, but she was unusual.

I've already mentioned my math teacher in high school—“Joey” Gansler. On the Jewish side, the high school teacher who influenced me most was Rabbi Shmuel Warshavchik. He had spent the years of the Second World War with the Mir Yeshiva in China, having escaped from the Nazis; after the war he made his way to the United States. He had a tremendous influence on me. He attracted me to the beauty of Talmudic study and the beauty of religious observance. He was, of course, *khareidi*, a term that is difficult to translate. Many people call them ultra-orthodox, but that has a pejorative flavor that I dislike. Literally, *khareidi* means worried, scared, concerned. It refers to trying to live the proper life and being very concerned about doing things right, about one's obligations to G-d and man. Warshavchik's enthusiasm and intensity—the fire in his eyes—lit a fire in me also. He eventually came to Israel, and died a few years ago in Haifa.

The next person who had quite an extraordinary influence on me was a young philosophy instructor at City College called Harry Tarter. I took from him courses called Philo 12 and 13—logic, the propositional calculus, a little set theory.

**H:** So your work in interactive epistemology had a good basis.

**A:** It was grounded in Philo 12 and Philo 13, where I learned about Russell's paradox and so on. We struck up a personal relationship that went far beyond the lecture hall, and is probably not very usual between an undergraduate and a university teacher. Later, my wife and children and I visited him in the Adirondacks, where he had a rustic home on the shores of a lake. When in Israel, he was our guest for the Passover Seder. What was most striking about him is that he would always question. He would always take something that appears self-evident and say, why is that so? At the Seder he asked a lot of questions. His wife tried to shush him; she said, Harry, let them go on. But I said, no, these questions are welcome. He was a remarkable person.

Another person who influenced me greatly was Jack W. Smith, whom I met in my postdoc period at Princeton, when working on the Naval Electronics Project. Let me describe this project briefly. One day we got a frantic phone call from Washington. Jack Smith was on the line. He was responsible for reallocating used naval equipment from decommissioned ships to active duty ships. These were very expensive items: radar, sonar, radio transmitters and receivers—large, expensive equipment, sometimes worth half a million 1955 dollars for each item. It was a lot of money. All this equipment was assigned to Jack Smith, who had to assign it to these ships. He tried to work out some kind of systematic way of doing it. The naval officers would come stomping into his office and pull out their revolvers and threaten to shoot him or otherwise use verbal violence. He was distraught. He

called us up and said, I don't care how you do this, but give me some way of doing it, so I can say, "The computer did this."

Now this is a classical assignment problem, which is a kind of linear programming problem. The constraints are entirely clear. There is only one small problem, namely, what's the objective function? Joe Kruskal and I solved the problem one way or another [3], and our solution was implemented. It is perhaps one of the more important pieces of my work, although it doesn't have many citations (it does have some). At that opportunity we formed a friendship with Jack Smith, his wife Annie and his five children, which lasted for many, many years. He was a remarkable individual. He had contracted polio as a child, so he limped. But nevertheless the energy of this guy was really amazing. The energy, the intellectual curiosity, and the intellectual breadth were outstanding. A beautiful family, beautiful people. He made a real mark on me.

Let's go back to graduate days. Of course my advisor, George Whitehead, had an important influence on me. He was sort of dry—not in spirit, but in the meticulousness of his approach to mathematics. We had weekly meetings, in which I would explain my ideas. I would talk about covering spaces and wave my hands around. He would say, Aumann, that's a very nice idea, but it's not mathematics. In mathematics we may discuss three-dimensional objects, but our proofs must be one-dimensional. You must write it down one word after another, and it's got to be coherent. This has stayed with me for many years.

We've already discussed Morgenstern, who promoted my career tremendously, and to whom I owe a big debt of gratitude.

The people with whom you interact also influence you. Among the people who definitely had an influence on me was Herb Scarf. I got the idea for the paper on markets with a continuum of traders by listening to Scarf; we became very good friends. Arrow also influenced me. I have had a very close friendship with Ken Arrow for many, many years. He did not have all that much direct scientific influence on my work, but his personality is certainly overpowering, and the indirect influence is enormous. Certainly Harsanyi's ideas about incomplete information had an important influence. As far as reading is concerned, the book of Luce and Raiffa, *Games and Decisions*, had a big influence.

Another important influence is Shapley. The work on "Markets with a Continuum of Traders" was created in my mind by putting together the paper of Shapley and Milnor on Oceanic Games and Scarf's presentation at the '61 games conference. And then there was our joint book, and all my work on nontransferable utility values, on which Shapley had a tremendous influence.

**H:** Let's go now to a combination of things that are not really related to one another, a potpourri of topics. They form a part of your worldview. We'll start with judicial discretion and restraint, a much disputed issue here in Israel.

**A:** There are two views of how a court should operate, especially a supreme court. One calls for judicial restraint, the other for judicial activism. The view of judicial restraint is that courts are for applying the laws of the land, not making





**FIGURE 8.** 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.

them; the legislature is for making laws, the executive for administering them, and the courts for adjudicating disputes in accordance with them.

The view of judicial activism is that the courts actually have a much wider mandate. They may decide which activities are reasonable, and which not; what is “just,” and what is not. They apply their own judgment rather than written laws, saying this is or isn’t “reasonable,” or “acceptable,” or “fair.” First and foremost this applies to activities of government agencies; the court may say, this is an unreasonable activity for a government agency. But it also applies to things like enforcing contracts; a judicially active court will say, this contract, to which both sides agreed, is not “reasonable,” and therefore we will not enforce it. These are opposite approaches to the judicial function.

In Israel it is conceded all around that the courts, and specifically the Supreme Court, are extremely activist, much more so than on the Continent or even in the United States. In fact, the chief justice of the Israeli Supreme Court, Aharon Barak, and I were once both present at a lecture where the speaker claimed that the Supreme Court justifiably takes on legislative functions, that it is a legislative body as well as a judicial body. Afterwards, I expressed to Mr. Barak my amazement at this pronouncement. He said, what's wrong with it? The lecturer is perfectly right. We are like the Sages of the Talmud, who also took on legislative as well as judicial functions.

**H:** Do you agree with that statement about the Talmud?

**A:** Yes, it is absolutely correct.

There are two major problems with judicial activism. One is that the judiciary is the least democratically constituted body in the government. In Israel, it is to a large extent a self-perpetuating body. Three of the nine members of the committee that appoints judges are themselves Supreme Court judges. Others are members of the bar who are strongly influenced by judges. A minority, only four out of the nine, are elected people—members of the Knesset. Moreover, there are various ways in which this committee works to overcome the influence of the elected representatives. For example, the Supreme Court judges on the committee always vote as a bloc, which greatly increases their power, as we know from Shapley value analyses.

In short, the way that the judiciary is constituted is very far from democratic. Therefore, to have the judiciary act in a legislative role is in violation of the principles of democracy. The principles of democracy are well based in game-theoretic considerations; see, for example, my paper with Kurz called “Power and Taxes” [37], which discusses the relation between power and democracy. In order that no one group should usurp the political power in the country, and also the physical wealth of the country, it is important to spread power evenly and thinly. Whereas I do not cast any aspersions now on the basic honesty of the judges of the Israeli Supreme Court, nevertheless, an institution where so much power is concentrated in the hands of so few undemocratically selected people is a great danger. This is one item.

**H:** The court not being democratically elected is not the issue, so long as the mandate of the court is just to interpret the law. It becomes an issue when the judicial branch creates the law.

**A:** Precisely. What is dangerous is a largely self-appointed oligarchy of people who make the laws. It is the *combination* of judicial activism with an undemocratically appointed court that is dangerous.

The second problem with judicial activism is that of uncertainty. If a person considering a contract does not know whether it will be upheld in court, he will be unwilling to sign it. Activism creates uncertainty: maybe the contract will be upheld, maybe not. Most decision-makers are generally assumed to be risk-averse, and they will shy away from agreements in an activist atmosphere. So there will be many potential agreements that will be discarded, and the result will be distinctly suboptimal.

**H:** But incomplete contracts may have advantages. Not knowing in advance what the court will decide—isn't that a form of incompleteness of the contract?

**A:** Incomplete contracts may indeed sometimes be useful, but that is not the issue here. The issue is a contract on which the sides have explicitly agreed, but that may be thrown out by the court. *Ex ante*, that cannot possibly be beneficial to the parties to the contract. It might conceivably be beneficial to society, if indeed you don't want that contract to be carried out. A contract to steal a car *should* be unenforceable, because car theft should be discouraged. But we don't want to discourage legitimate economic activity, and judicial activism does exactly that.

**H:** The uncertainty about the court's decision may be viewed also as a chance device—which may lead to a Pareto improvement. Like mutual insurance.

**A:** Well, okay, that is theoretically correct. Still, it is farfetched. In general, uncertainty is a dampening factor.

In brief, for these two reasons—introducing uncertainty into the economy and into the polity, and its undemocratic nature—judicial activism is to be deplored.

**H:** Another topic you wanted to talk about is war.

**A:** Barry O'Neill, the game theory political scientist, gave a lecture here a few months ago. Something he said in the lecture—that war has been with us for thousands of years—set me thinking. It really is true that there is almost nothing as ever-present in the history of mankind as war. Since the dawn of history we have had constant wars. War and religion, those are the two things that are ever-present with us. A tremendous amount of energy is devoted on the part of a very large number of well-meaning people to the project of preventing war, settling conflicts peacefully, ending wars, and so on. Given the fact that war is so, so prevalent, both in time and in space, all over the world, perhaps much of the effort of preventing or stopping war is misdirected. Much of this effort is directed at solving specific conflicts. What can we do to reach a compromise between the Irish Catholics in the Republic of Ireland and the Protestants in North Ireland? What can we do to resolve the conflict between the Hindus in India and the Moslems in Pakistan? What can we do to resolve the conflict between the Jews and the Arabs in the Middle East? One always gets into the particulars of these conflicts and neglects the more basic problems that present themselves by the very fact that we have had wars continuously. War is only apparently based on specific conflicts. There appears to be something in the way human nature is constituted—or if not human nature, then the way we run our institutions—that allows war and in fact makes it inevitable. Just looking at history, given the constancy of war, we should perhaps shift gears and ask ourselves what it is that causes war. Rather than establishing peace institutes, peace initiatives, institutions for studying and promoting peace, we should have institutions for studying war. *Not* with an immediate view to preventing war. Such a view can come later, but first we should understand the phenomenon.

It's like fighting cancer. One way is to ask, given a certain kind of cancer, what can we do to cure it? Chemotherapy? Radiation? Surgery? Let's do statistical studies that indicate which is more effective. That's one way of dealing with

cancer, and it's an important way. Another way is simply to ask, what *is* cancer? How does it work? Never mind curing it. First let's understand it. How does it get started, how does it spread? How fast? What are the basic properties of cells that go awry when a person gets cancer? Just study it. Once one understands it one can perhaps hope to overcome it. But before you understand it, your hope to overcome it is limited.

**H:** So, the standard approach to war and peace is to view it as a black box. We do not know how it operates, so we try ad-hoc solutions. You are saying that this is not a good approach. One should instead try to go inside the black box: to understand the roots of conflict—not just deal with symptoms.

**A:** Yes. Violent conflict may be very difficult to overcome. A relevant game-theoretic idea is that, in general, neither side really knows the disagreement level, the “reservation price.” It's like the Harsanyi–Selten bargaining model with incomplete information, where neither side knows the reservation price of the other. The optimum strategy in such a situation may be to go all the way and threaten. If the buyer thinks that the seller's reservation price is low, he will make a low offer, even if he is in fact willing to pay much more. Similarly for the seller. So conflict may result even when the reservation prices of the two sides are compatible. When this conflict is a strike, then it is bad enough, but when it's a war, then it is much worse. This kind of model suggests that conflict may be inevitable, or that you need different institutions in order to avoid it. If in fact it is inevitable in that sense, we should understand that. One big mistake is to say that war is irrational.

**H:** It's like saying that strikes are irrational.

**A:** Yes, and that racial discrimination is irrational (cf. Arrow). 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:** Exactly as in strikes, the only way to transmit to the other side how important this thing is to you may be to go to war.

**A:** Yes. In fact Bob Wilson discussed this in his Morgenstern lecture here in '94—just after a protracted strike of the professors in Israel.

**H:** Here in Israel, we unfortunately have constant wars and conflicts. One of the “round tables” of the Rationality Center—where people throw ideas at each other very informally—was on international conflicts. You presented there some nice game-theoretic insights.

**A:** One of them was the blackmailer's paradox. Ann and Bob must divide a hundred dollars. It is not an ultimatum game; they can discuss it freely. Ann says to Bob, look, I want ninety of those one hundred. Take it or leave it; I will not walk out of this room with less than ninety dollars. Bob says, come on, that's crazy. We have a hundred dollars. Let's split fifty-fifty. Ann says, no. Ann—“the blackmailer”—is perhaps acting irrationally. But Bob, if he is rational, will accept the ten dollars, and that's the end.

**H:** The question is whether she can commit herself to the ninety. Because if not, then of course Bob will say, you know what, fifty-fifty. Now *you* take it or leave it. For this to work, Ann must commit herself credibly.

**A:** In other words, it's not enough for her just to say it. She has to make it credible; and then Bob will rationally accept the ten. The difficulty with this is that perhaps Bob, too, can credibly commit to accepting no less than ninety. So we have a paradox: once Ann credibly commits herself to accepting no less than ninety, Bob is rationally motivated to take the ten. But then Ann is rationally motivated to make such a commitment. But Bob could also make such a commitment; and if both make the commitment, it is not rational, because then nobody gets anything.

This is the blackmailer's paradox. It is recognized in game theory, therefore, that it is perhaps not so rational for the guy on the receiving end of the threat to accept it.

What is the application of this to the situation we have here in Israel? Let me tell you this true story. A high-ranking officer once came to my office at the Center for Rationality and discussed with me the situation with Syria and the Golan Heights. This was a hot topic at the time. He explained to me that the Syrians consider land holy, and they will not give up one inch. When he told me that, I told him about the blackmailer's paradox. I said to him that the Syrians' use of the term "holy," land being holy, is a form of commitment. In fact, they must really convince themselves that it's holy, and they do. Just like in the blackmailer's paradox, we could say that it's holy; but we can't convince *ourselves* that it is. One of our troubles is that the term "holy" is nonexistent in our practical, day-to-day vocabulary. It exists only in religious circles. We accept holiness in other people and we are not willing to promote it on our own side. The result is that we are at a disadvantage because the other side can invoke holiness, but we have ruled it out from our arsenal of tools.

**H:** On the other hand, we do have such a tool: security considerations. That is the "holy" issue in Israel. We say that security considerations dictate that we must have control of the mountains that control the Sea of Galilee. There is no way that anything else will be acceptable. Throughout the years of Israel's existence security considerations have been a kind of holiness, a binding commitment to ourselves. The question is whether it is as strong as the holiness of the land on the other side.

**A:** It is less strong.

**H:** Maybe that explains why there is no peace with Syria.

**A:** You know, the negotiations that Rabin held with the Syrians in the early nineties blew up over a few meters. I really don't understand why they blew up, because Rabin was willing to give almost everything away. Hills, everything.

Without suggesting solutions, it is just a little bit of an insight into how game-theoretic analysis can help us to understand what is going on, in this country in particular, and in international conflicts in general.

**H:** Next, what about what you refer to as "connections"?

**A:** A lot of game theory has to do with relationships among different objects. I talked about this in my 1995 “birthday” lecture, and it is also in the Introduction to my *Collected Papers* [vi].

Science is often characterized as a quest for truth, where truth is something absolute, which exists outside of the observer. But I view science more as a quest for *understanding*, where the understanding is that of the observer, the scientist. Such understanding is best gained by studying relations—relations between different ideas, relations between different phenomena; relations between ideas and phenomena. Rather than asking “How does this phenomenon work?” we ask, “How does this phenomenon resemble others with which we are familiar?” Rather than asking “Does this idea make sense?” we ask, “How does this idea resemble other ideas?”

Indeed, the idea of relationship is fundamental to game theory. Disciplines like economics or political science use disparate models to analyze monopoly, oligopoly, perfect competition, public goods, elections, coalition formation, and so on. In contrast, game theory uses the *same* tools in all these applications. The nucleolus yields the competitive solution in large markets [16], the homogeneous weights in parliaments (cf. Peleg), and the Talmudic solution in bankruptcy games [46]. The fundamental notion of Nash equilibrium, which a priori reflects the behavior of consciously maximizing agents, is the *same* as an equilibrium of populations that reproduce blindly without regard to maximizing anything.

The great American naturalist and explorer John Muir said, “When you look closely at anything in the universe, you find it hitched to everything else.” Though Muir was talking about the natural universe, this applies also to scientific ideas—how we *understand* our universe.

**H:** How about the issue of assumptions vs. conclusions?

**A:** There is a lot of discussion in economic theory and in game theory about the reasonableness or correctness of assumptions and axioms. That is wrongheaded. I have never been so interested in assumptions. I am interested in conclusions. Assumptions don’t have to be correct; *conclusions* have to be correct. That is put very strongly, maybe more than I really feel, but I want to be provocative. When Newton introduced the idea of gravity, he was laughed at, because there was no rope with which the sun was pulling the earth; gravity is a laughable idea, a crazy assumption, it still sounds crazy today. When I was a child I was told about it. It did not make any sense then, and it doesn’t now; but it does yield the right answer. In science one never looks at assumptions; one looks at conclusions. It does not interest me whether this or that axiom of utility theory, of the Shapley value, of Nash bargaining is or is not compelling. What interests me is whether the *conclusions* are compelling, whether they yield interesting insights, whether one can build useful theory from them, whether they are testable. Nowhere else in science does one directly test assumptions; a theory stands or falls by the validity of the conclusions, not of the assumptions.

**H:** Would you like to say something about the ethical neutrality of game theory?

**A:** Ethical neutrality means that game theorists don't necessarily advocate carrying out the normative prescriptions of game theory. Game theory is about selfishness. Just like I suggested studying war, game theory studies selfishness. Obviously, studying war is not the same as advocating war; similarly, studying selfishness is not the same as advocating selfishness. Bacteriologists do not advocate disease, they study 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.

### SCIENTIFIC PUBLICATIONS OF ROBERT AUMANN

#### BOOKS

- [i] *Values of Non-Atomic Games*, with Lloyd S. Shapley. Princeton, NJ: Princeton University Press, 1974, xi + 333.
- [ii] *Game Theory* (in Hebrew), with Yair Tauman and Shmuel Zamir. Tel Aviv: Everyman's University, 1981. Vol. 1: 211, Vol. 2: 203.
- [iii] *Lectures on Game Theory*, Boulder, CO: Underground Classics in Economics, Westview Press, 1989, ix + 120.
- [iv] *Handbook of Game Theory with Economic Applications*, coedited with Sergiu Hart. Amsterdam: Elsevier. Vol. 1, 1992, xxvi + 733, Vol. 2, 1994, xxviii + 787, Vol. 3, 2002, xxx + 858.
- [v] *Repeated Games with Incomplete Information*, with Michael Maschler. Cambridge, MA: MIT Press, 1995, xvii + 342.
- [vi] *Collected Papers*, Vol. 1, xi + 786, Vol. 2, xiii + 792, Cambridge, MA: MIT Press, 2000.

#### ARTICLES

- [1] Asphericity of alternating knots. *Annals of Mathematics* 64 (1956), 374–392.
- [2] The coefficients in an allocation problem, with Joseph B. Kruskal. *Naval Research Logistics Quarterly* 5 (1958), 111–123.
- [3] Assigning quantitative values to qualitative factors in the naval electronics problem, with Joseph B. Kruskal. *Naval Research Logistics Quarterly* 6 (1959), 1–16.
- [4] Acceptable points in general cooperative  $n$ -person games. In A.W. Tucker and R.D. Luce (eds.), *Contributions to the Theory of Games IV*, Annals of Mathematics Study 40, pp. 287–324. Princeton, NJ: Princeton University Press, 1959.
- [5] Von Neumann–Morgenstern solutions to cooperative games without side payments, with Bezalel Peleg. *Bulletin of the American Mathematical Society* 66 (1960), 173–179.
- [6] Acceptable points in games of perfect information. *Pacific Journal of Mathematics* 10 (1960), 381–417.
- [7] A characterization of game structures of perfect information. *Bulletin of the Research Council of Israel* 9F (1960), 43–44.
- [8] Spaces of measurable transformations. *Bulletin of the American Mathematical Society* 66 (1960), 301–304.

- [9] Linearity of unrestrictedly transferable utilities. *Naval Research Logistics Quarterly* 7 (1960), 281–284.
- [10] The core of a cooperative game without side payments. *Transactions of the American Mathematical Society* 98 (1961), 539–552.
- [11] Borel structures for function spaces. *Illinois Journal of Mathematics* 5 (1961), 614–630.
- [12] Almost strictly competitive games. *Journal of the Society for Industrial and Applied Mathematics* 9 (1961), 544–550.
- [13] (a) Utility theory without the completeness axiom. *Econometrica* 30 (1962), 445–462.  
(b) Utility theory without the completeness axiom: a correction. *Econometrica* 32 (1964), 210–212.
- [14] A definition of subjective probability, with Frank J. Anscombe. *Annals of Mathematical Statistics* 34 (1963), 199–205.
- [15] On choosing a function at random. In F.W. Wright (ed.), *Ergodic Theory*, pp. 1–20. New Orleans, LA: Academic Press, 1963.
- [16] Markets with a continuum of traders. *Econometrica* 32 (1964), 39–50.
- [17] The bargaining set for cooperative games, with Michael Maschler. In M. Dresher, L.S. Shapley, and A.W. Tucker (eds.), *Advances in Game Theory*, Annals of Mathematics Study 52, pp. 443–476. Princeton, NJ: Princeton University Press, 1964.
- [18] Mixed and behavior strategies in infinite extensive games. In M. Dresher, L.S. Shapley, and A.W. Tucker (eds.), *Advances in Game Theory*, Annals of Mathematics Study 52, pp. 627–650. Princeton, NJ: Princeton University Press, 1964.
- [19] Subjective programming. In M.W. Shelly and G.I. Bryan (eds.), *Human Judgments and Optimality*, pp. 217–242. New York, NY: John Wiley and Sons, 1964.
- [20] A variational problem arising in economics, with Micha Perles. *Journal of Mathematical Analysis and Applications* 11 (1965), 488–503.
- [21] Integrals of set-valued functions. *Journal of Mathematical Analysis and Applications* 12 (1965), 1–12.
- [22] A method of computing the kernel of  $n$ -person games, with Bezalel Peleg and Pinchas Rabinowitz. *Mathematics of Computation* 19 (1965), 531–551.
- [23] Existence of competitive equilibria in markets with a continuum of traders. *Econometrica* 34 (1966), 1–17.
- [24] A survey of cooperative games without side payments. In M. Shubik (ed.), *Essays in Mathematical Economics in Honor of Oskar Morgenstern*, pp. 3–27. Princeton, NJ: Princeton University Press, 1967.
- [25] Random measure preserving transformations. In L.M. LeCam and J. Neyman (eds.), *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, Vol. II, Part 2, pp. 321–326. Berkeley: University of California Press, 1967.
- [26] Measurable utility and the measurable choice theorem. In *La Decision*, pp. 15–26. Paris: Editions du Centre National de la Recherche Scientifique, 1969.
- [27] Some thoughts on the minimax principle, with Michael Maschler. *Management Science* 18 (1972), P-54–P-63.
- [28] Disadvantageous monopolies. *Journal of Economic Theory* 6 (1973), 1–11.
- [29] Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics* 1 (1974), 67–96.
- [30] A note on Gale's example, with Bezalel Peleg. *Journal of Mathematical Economics* 1 (1974), 209–211.
- [31] Cooperative games with coalition structures, with Jacques Drèze. *International Journal of Game Theory* 4 (1975), 217–237.
- [32] Values of markets with a continuum of traders. *Econometrica* 43 (1975), 611–646.
- [33] An elementary proof that integration preserves uppersemicontinuity. *Journal of Mathematical Economics* 3 (1976), 15–18.
- [34] Agreeing to disagree. *Annals of Statistics* 4 (1976), 1236–1239.



- [35] Orderable set functions and continuity III: orderability and absolute continuity, with Uri Rothblum. *SIAM Journal on Control and Optimization* 15 (1977), 156–162.
- [36] The St. Petersburg paradox: a discussion of some recent comments. *Journal of Economic Theory* 14 (1977), 443–445.
- [37] Power and taxes, with Mordecai Kurz. *Econometrica* 45 (1977), 1137–1161.
- [38] Power and taxes in a multi-commodity economy, with Mordecai Kurz. *Israel Journal of Mathematics* 27 (1977), 185–234.
- [39] Core and value for a public goods economy: an example, with Roy J. Gardner and Robert W. Rosenthal. *Journal of Economic Theory* 15 (1977), 363–365.
- [40] On the rate of convergence of the core. *International Economic Review* 19 (1979), 349–357.
- [41] Recent developments in the theory of the Shapley value. In O. Lehto (ed.), *Proceedings of the International Congress of Mathematicians, Helsinki, 1978*, pp. 995–1003. Helsinki: Academia Scientiarum Fennica, 1980.
- [42] Survey of repeated games. In V. Böhm (ed.), *Essays in Game Theory and Mathematical Economics in Honor of Oskar Morgenstern*, Vol. 4 of Gesellschaft, Recht, Wirtschaft, Wissenschaftsverlag, pp. 11–42. Mannheim: Bibliographisches Institut, 1981.
- [43] Approximate purification of mixed strategies, with Yitzhak Katznelson, Roy Radner, Robert W. Rosenthal, and Benjamin Weiss. *Mathematics of Operations Research* 8 (1983), 327–341.
- [44] Voting for public goods, with Mordecai Kurz and Abraham Neyman. *Review of Economic Studies* 50 (1983), 677–694.
- [45] An axiomatization of the non-transferable utility value. *Econometrica* 53 (1985), 599–612.
- [46] Game-theoretic analysis of a bankruptcy problem from the Talmud, with Michael Maschler. *Journal of Economic Theory* 36 (1985), 195–213.
- [47] What is game theory trying to accomplish? In K. Arrow and S. Honkapohja (eds.), *Frontiers of Economics*, pp. 28–76. Oxford: Basil Blackwell, 1985.
- [48] On the non-transferable utility value: a comment on the Roth–Shafer Examples. *Econometrica* 53 (1985), 667–677.
- [49] Rejoinder. *Econometrica* 54 (1986), 985–989.
- [50] Bi-convexity and bi-martingales, with Sergiu Hart. *Israel Journal of Mathematics* 54 (1986), 159–180.
- [51] Values of markets with satiation or fixed prices, with Jacques Drèze. *Econometrica* 54 (1986), 1271–1318.
- [52] Power and public goods, with Mordecai Kurz and Abraham Neyman. *Journal of Economic Theory* 42 (1987), 108–127.
- [53] Correlated equilibrium as an expression of Bayesian rationality. *Econometrica* 55 (1987), 1–18.
- [54] Value, symmetry, and equal treatment: a comment on Scafuri and Yannelis. *Econometrica* 55 (1987), 1461–1464.
- [55] Game theory. In J. Eatwell, M. Milgate, and P. Newman (eds.), *The New Palgrave, A Dictionary of Economics*, Vol. 2, pp. 460–482. London and Basingstoke: Macmillan, 1987.
- [56] Endogenous formation of links between players and of coalitions: an application of the Shapley value, with Roger Myerson. In A.E. Roth (ed.), *The Shapley Value: Essays in Honor of Lloyd S. Shapley*, pp. 175–191. Cambridge: Cambridge University Press, 1988.
- [57] Cooperation and bounded recall, with Sylvain Sorin. *Games and Economic Behavior* 1 (1989), 5–39.
- [58] CORE as a macrocosm of game-theoretic research, 1967–1987. In B. Cornet and H. Tulkens (eds.), *Contributions to Operations Research and Economics: The Twentieth Anniversary of CORE*, pp. 5–16. Cambridge and London: The MIT Press, 1989.
- [59] Nash equilibria are not self-enforcing. In J.J. Gabszewicz, J.-F. Richard, and L. Wolsey (eds.), *Economic Decision Making: Games, Econometrics and Optimisation (Essays in Honor of Jacques Drèze)*, pp. 201–206. Amsterdam: Elsevier Science Publishers, 1990.

- [60] Irrationality in game theory. In P. Dasgupta, D. Gale, O. Hart, and E. Maskin (eds.), *Economic Analysis of Markets and Games (Essays in Honor of Frank Hahn)*, pp. 214–227. Cambridge and London: MIT Press, 1992.
- [61] Long-term competition: a game-theoretic analysis, with Lloyd S. Shapley. In N. Megiddo (ed.), *Essays in Game Theory in Honor of Michael Maschler*, pp. 1–15. New York: Springer, 1994.
- [62] The Shapley value. In J.-F. Mertens and S. Sorin (eds.), *Game-Theoretic Methods in General Equilibrium Analysis*, pp. 61–66. Dordrecht: Kluwer Academic Publishers, 1994.
- [63] Economic applications of the Shapley value. In J.-F. Mertens and S. Sorin (eds.), *Game-Theoretic Methods in General Equilibrium Analysis*, pp. 121–133. Dordrecht: Kluwer Academic Publishers, 1994.
- [64] Backward induction and common knowledge of rationality. *Games and Economic Behavior* 8 (1995), 6–19.
- [65] Epistemic conditions for Nash equilibrium, with Adam Brandenburger. *Econometrica* 63 (1995), 1161–1180.
- [66] Reply to Binmore. *Games and Economic Behavior* 17 (1996), 138–146.
- [67] Reply to Margalit and Yaari. In K.J. Arrow, E. Colombatto, M. Perlman, and C. Schmidt (eds.), *The Rational Foundations of Economic Equilibrium*, pp. 106–107. London and Basingstoke: Macmillan, 1996.
- [68] Reply to Binmore and Samuelson. In K.J. Arrow, E. Colombatto, M. Perlman, and C. Schmidt (eds.), *The Rational Foundations of Economic Equilibrium*, pp. 130–131. London and Basingstoke: Macmillan, 1996.
- [69] The absent-minded driver, with Sergiu Hart and Motty Perry. *Games and Economic Behavior* 20 (1997), 102–116.
- [70] The forgetful passenger, with Sergiu Hart and Motty Perry. *Games and Economic Behavior* 20 (1997), 117–120.
- [71] Rationality and bounded rationality. *Games and Economic Behavior* 21 (1997), 2–14.
- [72] On the centipede game. *Games and Economic Behavior* 23 (1998), 97–105.
- [73] Common priors: a reply to Gul. *Econometrica* 66 (1998), 929–938.
- [74] Interactive epistemology I: knowledge. *International Journal of Game Theory* 28 (1999), 263–300.
- [75] Interactive epistemology II: probability. *International Journal of Game Theory* 28 (1999), 301–314.
- [76] Species survival and evolutionary stability in sustainable habitats, with Werner Güth. *Journal of Evolutionary Economics* 10 (2000), 437–447.
- [77] The rationale for measurability. In G. Debreu, W. Neufeind, and W. Trockel, *Economics Essays, A Festschrift for Werner Hildenbrand*, pp. 5–7. Berlin: Springer, 2001.
- [78] Harsanyi's sweater. *Games and Economic Behavior* 36 (2001), 7–8.
- [79] Incomplete information, with Aviad Heifetz. In R.J. Aumann and S. Hart (eds.), *Handbook of Game Theory with Economic Applications, Vol. 3*, pp. 1665–1686. Amsterdam: Elsevier, 2002.
- [80] Risk aversion in the Talmud. *Economic Theory* 21 (2003), 233–239.
- [81] Presidential address. *Games and Economic Behavior* 45 (2003), 2–14.
- [82] Long cheap talk, with Sergiu Hart. *Econometrica* 71 (2003), 1619–1660.

### OTHER

- [83] Letter to Leonard Savage, 8 January 1971. In J.H. Drèze, *Essays on Economic Decisions under Uncertainty*, pp. 76–78. Cambridge: Cambridge University Press, 1987.
- [84] On the state of the art in game theory, an interview. W. Albers, W. Güth, P. Hammerstein, B. Moldovanu, and E. van Damme (eds.), *Understanding Strategic Interaction (Essays in Honor of Reinhard Selten)*. Berlin: Springer, 1997, pp. 8–34. Reprinted in *Games and Economic Behavior* 24 (1998), 181–210.