

Robert Aumann's Game and Economic Theory*

Sergiu Hart[†]

March 7, 2006

Abstract

An overview of the landmark contributions of Robert J. Aumann, winner of the 2005 Nobel Memorial Prize in Economic Sciences.

Keywords: Robert Aumann; game theory; repeated games; correlated equilibrium; common knowledge; continuum of agents

JEL Classification: C7, D4, D8, B31

I. Introduction

Robert J. Aumann has been a central figure in developing game theory and establishing its key role in modern economics. Aumann has shaped the field through his fundamental and pioneering work. His contributions during the past half century are profound—conceptually, and often also mathematically—and at the same time crystal clear, elegant, and illuminating.

Among Aumann's main scientific contributions, three areas stand out:

*Partially supported by a grant of the Israel Science Foundation. The author thanks Elchanan Ben-Porath, Avinash Dixit, Andreu Mas-Colell, Eyal Winter, and Shmuel Zamir for their comments. A presentation is available at <http://www.ma.huji.ac.il/hart/abs/aumann-p.html>.

[†]Center for the Study of Rationality, Department of Economics, and Institute of Mathematics, The Hebrew University of Jerusalem. *E-mail:* hart@huji.ac.il. *Web page:* <http://www.ma.huji.ac.il/hart>.

1. *Repeated Games*: the study of long-term interactions and the accounting for cooperative and other patterns of behavior in terms of the classical “selfish” utility-maximizing paradigm.
2. *Knowledge, Rationality, and Equilibrium*: the analysis of knowledge in multi-agent environments and the laying of the foundations of rationality and equilibrium.
3. *Perfect Competition*: the modeling of perfect competition by a continuum of agents and the relations between price equilibria and cooperative outcomes.

The next three sections discuss each of these topics in turn (with special emphasis on the area of repeated games, the main topic for which Aumann was awarded the 2005 Nobel Memorial Prize in Economic Sciences). We conclude with a short overview of some of his other contributions.

II. Repeated Games

Most relationships between decision-makers last a long time. Competition between firms in markets, insurance contracts, credit relationships, and negotiations, are often long-term affairs. The same is true of employer–employee, lawyer–client, and firm–subcontractor relationships, of conflicts and agreements between political parties and nations, of evolutionary processes in biology.

In such long-term interactions, the different stages are naturally interdependent. Decision-makers react to past experience, and take into account the future impact of their choices. Many interesting and important patterns of behavior—like rewarding and punishing, cooperation and threats, transmitting information and concealing it—can only be seen in multi-stage situations.

The general framework for studying strategic interaction is game theory. The “players” are decision-makers—be they individuals, collectives, computer programs, or genes—whose actions mutually affect one another. Game

theory studies interactions (“games”) from a rational decision-theoretic point of view, and develops models and methodologies that are universal and widely applicable—particularly in economics.

Foremost among the multi-stage models are *repeated games*, where the same game is played at each stage. Such models allow us to untangle the complexities of behavior from the complexities of the interaction itself, which here is simply stationary over time.

The Classical Folk Theorem

The simplest setup is as follows. Given an n -person game G , let G^* be the *supergame* of G : the same n players are repeatedly playing the game G , at periods $t = 1, 2, \dots$. (It is customary to call G the “one-shot game,” and G^* the “repeated game”; to avoid confusion, the choices of the players in G are referred to as “actions,” and the choices in G^* as “strategies.”) At the end of each period, every player is informed of the actions chosen by all players in that period; thus, before choosing his action at time t , each player knows what everyone did in all the previous periods $1, 2, \dots, t - 1$.

The payoffs in G^* are defined as an appropriate average of the payoffs received in each period. (There are technical difficulties here; however, any reasonable definition that is consistent with the “long-term” idea turns out to lead to essentially the same result; see Aumann 1959, 1989 [Chapter 8], and the survey of Sorin 1992).

The question is, What are the Nash¹ (1951) *equilibria* of G^* , i.e., those strategy combinations for the n players such that no player can increase his payoff by unilaterally changing his strategy? What are the resulting outcomes? The answer is given in the following result, which emerged in the late fifties. Its precise authorship is unknown, and it is part of the “folklore” of game theory; it is thus known as the “Folk Theorem.”²

¹John F. Nash, Jr. was awarded the 1994 Nobel Prize in Economics for this work.

²See the sections *The Import of the Folk Theorem* and *A Historical Note* below.

The Folk Theorem. *The set of Nash equilibrium outcomes of the repeated game G^* is precisely the set of feasible and individually rational outcomes of the one-shot game G .*

An “outcome”³ is represented as an n -dimensional vector $a = (a_1, a_2, \dots, a_n)$, whose i th coordinate a_i is the payoff (or utility) of player i . Such an outcome is *feasible in G* if it can be obtained from some combination of actions in the game G , or more generally from a probability distribution over such action combinations.⁴ For example, if G is the two-person “Battle of the Sexes” game⁵

	L	R
T	3, 1	0, 0
B	0, 0	1, 3

then the set of feasible outcomes is the triangle with vertices $(3, 1)$, $(1, 3)$, and $(0, 0)$ (see Figure 1). Note that some feasible outcomes require a *joint* (or *correlated*) randomization: for instance, $(2, 2)$ corresponds to the two action combinations TL and BR being played with equal probabilities of $1/2$ each (thus $(2, 2)$ is *not* a Nash equilibrium outcome in G ; in fact, it *cannot* be obtained from any pair of *independently* randomized actions of the two players).⁶

A payoff of a player i is “individually rational” in G if player i cannot be prevented by the other players from getting it. Formally, let r_i be the minimax value of the zero-sum game where i wants to maximize, and the

³Or “payoff vector.”

⁴Formally, an outcome is *feasible* if it lies in the convex hull of the vectors $(u_1(s), u_2(s), \dots, u_n(s))$, where s ranges over all pure action combinations (s_1, s_2, \dots, s_n) , and u_i denotes the payoff function of player i .

⁵Here and in the other examples, the actions are called “T(op),” “B(ottom),” “L(eft),” and “R(ight).”

⁶It would thus be more appropriate (but unwieldy) to call this “coordinated feasibility,” or “cooperative feasibility,” or “joint feasibility.”




Figure 1: The Folk Theorem for the Battle of the Sexes game

other players want to minimize, the payoff of player⁷ i . An outcome a is *individually rational in G* if $a_i \geq r_i$ for all players $i = 1, 2, \dots, n$. Returning to the “Battle of the Sexes” example, we have⁸ $r_1 = r_2 = \frac{3}{4}$, and the set of feasible and individually rational outcomes in G is the darkly shaded area in Figure 1. The Folk Theorem says that this is precisely the set of outcomes of all Nash equilibria of the infinitely repeated game G^* .

An informal proof of the Folk Theorem is as follows (complete formal proofs require many technical details; see for instance Aumann 1959, 1989 [Chapter 8]).

One direction consists of showing that the payoff vectors of the Nash equilibria of G^* are feasible and individually rational in G . First, any payoff

⁷Take player 1; then $r_1 = \min_{x_2, \dots, x_n} \max_{x_1} u_1(x_1, x_2, \dots, x_n)$, where each x_j ranges over the randomized actions of player j , and u_1 is the payoff function of player 1.

⁸The row player can guarantee himself a payoff of $\frac{3}{4}$ by playing T with probability $\frac{1}{4}$ and B with probability $\frac{3}{4}$: his payoff is then $\frac{1}{4} \cdot 0 + \frac{3}{4} \cdot 1 = \frac{3}{4}$ when the column player plays L, and $\frac{1}{4} \cdot 3 + \frac{3}{4} \cdot 0 = \frac{3}{4}$ when the column player plays R (and so it is also $\frac{3}{4}$ when the column player plays a randomized action).

vector in G^* is an average of payoff vectors of G , and thus feasible in G . And second, any Nash equilibrium of G^* (and of G) must yield individually rational payoffs; otherwise, if some player, say player i , were to get less than r_i , then, given the other players' strategies, i could find a reply that gives him a payoff of at least r_i in each period and thus also in the long run—which means that player i has a profitable deviation, in contradiction to the Nash equilibrium requirement.

The other direction introduces two basic ingredients of repeated games: a “plan,” and “punishments” (or “threats”). Let a be a feasible and individually rational payoff vector of G ; from the definitions of these two notions we immediately obtain:

- (i) There exists a sequence of pure action combinations in G whose payoff in G^* is precisely a (in the above example, to get $(2, 2)$, alternate between TL in odd periods, $t = 1, 3, 5, \dots$, and BR in even periods, $t = 2, 4, 6, \dots$)—this is the “*plan*.”
- (ii) For each player i there exist (mixed) actions of the other players in G such that no matter what i plays his payoff in G cannot exceed r_i —these constitute the “*punishment of i* .”

Now define for each player j a strategy in G^* as follows: play according to the plan (i.e., play the sequence of actions in (i) above), so long as everyone did so in the past; if someone deviated from the plan, say player i was the first to do so, at some period⁹ d , then at every period thereafter, $t = d + 1, d + 2, \dots$, punish player i (i.e., play the mixed action as defined in (ii) above which guarantees that, whatever player i does, his payoff will be at most r_i).

The combination of these strategies results in all players following the plan, and so the payoff is indeed a by (i). Moreover, it constitutes a Nash equilibrium of G^* since any deviation by a player i will make the punishments against him go into effect,¹⁰ which will reduce his long-run payoff in G^* to at most r_i (which is less than or equal to the payoff a_i that he is getting

⁹If more than one player deviated at the same time, choose one of them.

¹⁰Since the plan consists of a sequence of pure actions, any deviation is immediately detected.

under the plan)—so no increase in payoff is possible. Thus the threat of punishment ensures that each player fulfils his part of the joint plan.

We will refer to the special Nash equilibria constructed above as “canonical equilibria.”

The Import of the Folk Theorem

By their nature, repeated games are complex objects—for players to play, and for theorists to analyze. There are a huge number of possible strategies, even when the game is repeated just a few times¹¹; in addition, many of these strategies are extremely complex. This makes the equilibrium analysis appear, at first, unmanageable.

The Folk Theorem shows that in fact this is not so. The resulting geometric characterization of the equilibrium outcomes is extremely simple (again, see Figure 1). What is more important however is the resulting *behavioral characterization*: every equilibrium of G^* is outcome-equivalent to a canonical equilibrium, which consists of a coordinated plan, supported by the threat of appropriate punishments.¹²

The most valuable insight from the analysis is the connection that is established between the so-called “*noncooperative*” and “*cooperative*” theories. The Folk Theorem relates noncooperative behavior in the repeated game (i.e., equilibrium in G^*) to cooperative behavior in the one-stage game (i.e., joint feasibility in G). This is best described in a survey of Aumann:

“The theory of repeated games of complete information is concerned with the evolution of fundamental patterns of interaction between people (or for that matter, animals; the problems it attacks are similar to those of social biology). Its aim is to account for phenomena such as cooperation, altruism, revenge, threats (self-destructive or otherwise), etc.—phenomena which may at

¹¹The number of strategies is doubly exponential in the number of repetitions.

¹²The “Revelation Principle” in mechanism design is of a similar type: everything that can be implemented can also be implemented by a simple “direct mechanism.”

first seem irrational—in terms of the usual “selfish” utility-maximizing paradigm of game theory and neoclassical economics. [...]

The significance of the Folk Theorem is that it relates cooperative behavior in the game G to non-cooperative behavior in its supergame G^* . This is the fundamental message of the theory of repeated games of complete information; that cooperation may be explained by the fact that the “games people play”—i.e., the multiperson decision situations in which they are involved—are not one-time affairs, but are repeated over and over. In game-theoretic terms, an outcome is cooperative if it requires an outside enforcement mechanism to make it “stick.” Equilibrium points are self-enforcing; once an equilibrium point is agreed upon, it is not worthwhile for any player to deviate from it. Thus it does not require any outside enforcement mechanism, and so represents non-cooperative behavior. On the other hand, the general feasible outcome does require an enforcement mechanism, and so represents the cooperative approach. In a sense, the repetition itself, with its possibilities for retaliation, becomes the enforcement mechanism.” (Aumann 1981 [Section 1])

Thus, the Folk Theorem shows, first, that one can succinctly analyze complex repeated interactions; second, that simple, natural, and familiar behaviors emerge; and third, how noncooperative strategic behavior brings about cooperation. This emergence of cooperation from a noncooperative setup makes repeated games a fascinating and important topic.

The result of the Folk Theorem has turned out to be extremely robust. The extensions and generalizations are concerned with varying the equilibrium concept, the long-run evaluation of the payoffs, imposing restrictions on strategies, bounded rationality, modifying the structure of the game, introducing asymmetric information, imperfect monitoring, and so on. It should be emphasized that many of the results in this vast literature are not just simple extensions; they almost always embody new ideas and important insights, while overcoming the many conceptual, and at times also technical,

complexities of the models.

A Historical Note

The Folk Theorem was essentially known to most people working in the area in the late fifties. However, it had not been published. Perhaps it was considered too simple a result; perhaps, too complicated to write rigorously.¹³

The discounted repeated Prisoner's Dilemma with discount factor close enough to 1 is mentioned in Luce and Raiffa (1957 [page 102]) as having equilibria yielding efficient outcomes; Shubik (1959b [Section 10.4]) presents a more detailed analysis.

Aumann (1959, 1960, 1961) was the first to provide an extensive analysis of infinitely repeated games. After setting up the model in an explicit and rigorous manner, he showed that the two approaches—noncooperative for the repeated game G^* , and cooperative for the one-shot game G —lead to the same solution.¹⁴ The 1959 paper of Aumann is the fundamental paper on repeated games; while it goes beyond the classical folk theorem, it addresses and resolves the basic issues that were necessary for the development of this area.

Credibility and Perfect Equilibria

The “canonical equilibria” of G^* (recall the end of the section *The Classical Folk Theorem* above) entail punishments that are constructed so as to decrease the payoff of a deviating player. However, these punishment actions may hurt some of the punishers as well. In such a case, after a deviation of a player i , a punishing player j may find it to his own advantage *not* to punish i , since j 's own payoff would decrease as a consequence. The threat of punishment appears not to be rational, and thus not credible.¹⁵

¹³Luce and Raiffa (1957 [top of page 102]) point out that there are difficulties here in translating intuition into precise arguments.

¹⁴See the section *Coalitions and Strong Equilibria* below.

¹⁵Aumann (1959 [beginning of Section 10]) points out that there are difficulties with equilibria that entail unrelenting punishment, which moreover may also hurt the punishers, and that there are other, more reasonable equilibrium points, “which give deviating players a chance to return to the fold.”

But wouldn't the equilibrium requirements take care of this problem? After all, in equilibrium any player, in particular player j , cannot increase his payoff by modifying his strategy. But the payoff is determined by the plan alone; deviations do not occur along the equilibrium play. Therefore changes in the strategy of player j following a deviation by player i do not affect player j 's payoff. This (perhaps slightly technical) argument shows that the equilibrium requirements have no implication for the behavior of the players in the punishment phase. In equilibrium, the threat of punishment deters player i from deviating; but once i has in fact deviated, a punishing player j may benefit by not carrying out the punishment.

What this suggests is a need to supplement the notion of equilibrium with additional requirements, namely, that after *every* history of play—whether or not it occurs in equilibrium—no player should be able to increase his payoff in the continuation of the game by unilaterally changing his strategy there. These conditions are embodied in the concept of (subgame) *perfect equilibrium* (Selten¹⁶ 1965, 1975), which requires the combination of strategies to constitute an equilibrium not only in the whole game but also in every continuation game (“subgame”).¹⁷

A strengthening of the Folk Theorem, due to Aumann and Shapley (1976) and to Rubinstein (1976, 1979), in independent work, shows that in the repeated game G^* , the equilibrium outcomes and the perfect equilibrium outcomes in fact coincide.

The Perfect Folk Theorem (Aumann–Shapley 1976, Rubinstein 1976). *The set of perfect Nash equilibrium outcomes of the repeated game G^* is precisely the set of feasible and individually rational outcomes of the one-shot game G .*

The idea is to refine the construction of equilibrium, so that the result-

¹⁶Reinhard Selten was awarded the 1994 Nobel Prize in Economics for this work.

¹⁷In the special case where the punishments form a Nash equilibrium of the one-shot game, like the Bertrand equilibrium in a duopoly model, or the Battle of the Sexes game, the canonical equilibria of the classical Folk Theorem turn out to be in fact perfect (since punishing is now an equilibrium in the continuation game); see Friedman (1971).

ing “canonical perfect equilibria” entail punishments that last only finitely many periods (during which the one-period gain made by a deviation from the cooperative plan is essentially wiped out), after which everyone returns to the plan. The return to the plan makes it rational for each player to punish all deviations by other players—and the equilibrium is indeed perfect (there are some subtleties here). In short: a plan, punishments—and also “*forgiveness*”; punishing is not forever, but just long enough to make deviations unprofitable.

The importance of the Perfect Folk Theorem is that the threats now become *credible*—they will indeed be carried out if needed—which makes the cooperative plan much more reasonable and robust.

Coalitions and Strong Equilibria

Another strengthening of the equilibrium requirements considers deviations by *groups* of players (rather than just single players). A *strong equilibrium*, a concept introduced by Aumann (1959), is a combination of strategies from which no group of players can profitably deviate. Aumann (1959) shows that the outcomes of the strong equilibria of the repeated game G^* constitute a cooperative solution of the one-shot game G , namely, the appropriate “core”¹⁸ of G .

The Strong Folk Theorem (Aumann 1959). *The set of strong Nash equilibrium outcomes of the repeated game G^* is precisely the core of the one-shot game G .*

Like the classical Folk Theorem, this result demonstrates the connection between the noncooperative approach (the strong equilibria of G^*) and the cooperative approach (the core of G). In fact, this work led Aumann and others to the development of the theory of general cooperative games—“non-transferable utility” games¹⁹—which are most important in economic theory (see Aumann 1967).

¹⁸Specifically, the so-called “ β -core.”

¹⁹Prior to this, only “transferable utility” or “side payment” games were studied.

Asymmetric Information

A most important and fascinating extension of the Folk Theorem is to the case of *asymmetric information*, where different players possess different knowledge of the relevant parameters of the one-shot game that is repeatedly played; for instance, a player may not know the utility functions of the other players.

Now, information is valuable. But how should it be used advantageously—to gain a competitive edge, to attain mutual benefits through cooperation, or both?

To illustrate the issues involved, assume for concreteness that one player, call him the “informed player,” has private information which the other players do not possess. In a one-shot interaction, the informed player will clearly utilize his information so as to gain as much as possible; this will often require him to play different actions depending on his private information. However, if the situation is repeated, then the other players, by observing the action taken by the informed player, may infer what his information was. But then the informed player will no longer have an informational advantage in future periods. So, what good is private information to the informed player if he cannot use it to his benefit? The problem here is to find the right balance: to use the information as much as possible, while revealing it as little as possible.

This is one side of the coin. The other is that there are situations where the informed player would like to convey his information to the others, so that their resulting actions will benefit him. But that is not always easy to do; can the uninformed players trust him? (Should one trust a shop owner who claims that buying from him is a good “deal”? Wouldn’t he claim the same even if it weren’t so?) The problem is how to make the revealed information credible, so that everyone benefits. Of course, repetition—i.e., the long-term relationship—is essential here.

In the mid-sixties, following the Harsanyi²⁰ (1967–1968) model of incomplete information games, Aumann and Michael Maschler founded and devel-

²⁰John C. Harsanyi was awarded the 1994 Nobel Prize in Economics for this work.

oped the theory of *repeated games with incomplete information*. In a series of path-breaking reports written in 1966, 1967, and 1968,²¹ Aumann and Maschler set up the basic models, and showed how the complex issues in the use of information alluded to above can actually be resolved in an explicit and elegant way.

As suggested by the Folk Theorem, there are two issues that need to be addressed: individual rationality and feasibility. We will deal with each in turn.

Individual Rationality and the Optimal Use of Information

We start with individual rationality: determining how much a player can *guarantee* (no matter what the other players do). We will illustrate the issues involved with three simple examples. In each example there will be two players who repeatedly play the same one-shot game, and we will consider individual rationality from the point of view of the row player.²² The row player's payoffs are given by either the matrix M1 or the matrix M2; the row player is informed of which matrix is the "true matrix," but the column player is not—he only knows that the two matrices are equally likely. After each stage both players observe the actions taken, but not the payoffs.²³

The first example is

²¹These unpublished reports to the Mathematica Institute were widely circulated in the profession. They are now collected in one book, Aumann and Maschler (1995), together with very extensive notes on the subsequent developments.

²²The game may thus be thought of as a *zero-sum game*, with the row player as the maximizer of his payoffs, and the column player as the minimizer.

²³Thus information is transmitted only through actions (if the column player were to observe the payoffs, he could determine immediately which matrix is the true matrix). This is a simplifying assumption that allows a clear analysis. Once these games are studied and well understood, one goes on to the general model with so-called "signalling matrices" (where each combination of actions generates a certain signal to each player; the signal could include the payoff, or be more or less general; this is discussed already in Aumann 1959 [Section 6, second paragraph]).

<i>Game 1</i>		
		L R
M1 (probability = $\frac{1}{2}$) :	T	4 0
	B	0 0
		L R
M2 (probability = $\frac{1}{2}$) :	T	0 0
	B	0 4

Again, the question we ask is, How much can the row player guarantee? If he were to play optimally in each matrix, i.e., T when the true matrix is M1, and B when it is M2 (this is the only way he may get positive payoffs), then his action would reveal which matrix is the true matrix. Were the column player to play R after seeing T, and L after seeing B, the row player would get a payoff of²⁴ 0. If, instead, the row player were to ignore his information, he would face the “average non-revealing game” with payoff matrix²⁵

		L R
$\frac{1}{2} M1 + \frac{1}{2} M2 :$	T	2 0
	B	0 2

Here the best he can do is to randomize equally between T and B, which would guarantee him a payoff of 1. This is better than 0, and it turns out that 1 is in fact the most the row player can guarantee in the long run here (the proof is by no means immediate). So, in Game 1, *the row player can*

²⁴Except, perhaps, in the first period (which is negligible in the long run).

²⁵In each cell, the payoff is the average of the corresponding payoffs in the two matrices.

guarantee the most by concealing his information and playing as if he did not know which matrix is the true matrix.

We modify the payoffs and get the second example:

<i>Game 2</i>			
		L	R
M1 (probability = $\frac{1}{2}$) :	T	4	4
	B	4	0
		L	R
M2 (probability = $\frac{1}{2}$) :	T	0	4
	B	4	4

In this game, if the row player ignores his information, then the average game is

		L	R
$\frac{1}{2}$ M1 + $\frac{1}{2}$ M2 :	T	2	4
	B	4	2

in which he can guarantee a payoff of 3 (by randomizing equally between T and B). If, however, the row player were to play T when the true matrix is M1 and B when it is M2, he would be guaranteed a payoff of 4, which clearly is the most he can get in this game. So, in Game 2, *the row player can guarantee the most by using—and thus fully revealing—his information.*

The third example is

<i>Game 3</i>				
		L	C	R
M1 (probability = $\frac{1}{2}$) :	T	4	2	0
	B	4	-2	0
		L	C	R
M2 (probability = $\frac{1}{2}$) :	T	0	-2	4
	B	0	2	4

Using his information fully—playing T when the true matrix is M1 and B when it is M2—guarantees only a payoff of 0 to the row player (the column player plays R after seeing T, and L after seeing B). Ignoring the information leads to the average game

		L	C	R
$\frac{1}{2} M1 + \frac{1}{2} M2 :$	T	2	0	2
	B	2	0	2

in which the row player can again guarantee only a payoff of 0 (the column player plays C). At this point it may appear that the row player cannot guarantee more than 0 in this game (since 0 is the most he is guaranteed, whether concealing or revealing the information). But that turns out to be false: by *partially using* his information—and thus *partially revealing* it—the row player can guarantee more. Indeed, consider the following strategy σ for

the row player:

<i>Strategy σ of the row player</i>	
If M1 :	$\left\{ \begin{array}{ll} \text{play T forever,} & \text{with probability } \frac{3}{4}, \\ \text{play B forever,} & \text{with probability } \frac{1}{4}. \end{array} \right.$
If M2 :	$\left\{ \begin{array}{ll} \text{play T forever,} & \text{with probability } \frac{1}{4}, \\ \text{play B forever,} & \text{with probability } \frac{3}{4}. \end{array} \right.$

To see what happens, consider the case where the randomization yielded “play T forever.” After the first period (in which T was played), the posterior probability that the true matrix is M1 is now, by Bayes’ rule,

$$\begin{aligned} \mathbf{P}(M1|T) &= \frac{\mathbf{P}(T|M1) \cdot \mathbf{P}(M1)}{\mathbf{P}(T|M1) \cdot \mathbf{P}(M1) + \mathbf{P}(T|M2) \cdot \mathbf{P}(M2)} \\ &= \frac{\frac{3}{4} \cdot \frac{1}{2}}{\frac{3}{4} \cdot \frac{1}{2} + \frac{1}{4} \cdot \frac{1}{2}} = \frac{3}{4}, \end{aligned}$$

and, for matrix M2, it is $\mathbf{P}(M2|T) = \frac{1}{4}$. The “average game” following T in the first period is thus

		L	C	R
$\frac{3}{4} M1 + \frac{1}{4} M2 :$	T	3	1	1
	B	3	-1	1

in which playing T guarantees 1 to the row player. A similar computation applies to the case of “play B forever,” where $\mathbf{P}(M1|T) = \frac{1}{4}$ and $\mathbf{P}(M2|T) = \frac{3}{4}$, and playing B guarantees also 1. Therefore the strategy σ guarantees to the row player a payoff of 1—strictly more than either full revelation or full

concealment—which turns out to be the most he can guarantee in this game. Note that the choice between the two options of “play T forever” and “play B forever” is made by the row player *only once*, in the first period, based on the information he has: the probabilities depend on which matrix is the true matrix.²⁶ However, the action of the row player reveals his information only *partially*: the prior probability of $\frac{1}{2}$ on M1 now becomes either the posterior probability $\frac{3}{4}$ (after T) or $\frac{1}{4}$ (after B).²⁷ So, in Game 3, *the partial use of information—and thus the ensuing partial revelation—is the only way that the row player guarantees the most.*

To summarize the three examples: in some games the informed player uses his information and reveals it fully; in others he does not use his information at all and thus conceals it; and in still others he uses and reveals it partially. When exactly does each case occur? The following general result gives a precise answer. Some notations: let M1 and M2 be two matrices of the same size, viewed as two-person zero-sum games;²⁸ one of them is chosen, with probabilities p and $1 - p$, respectively. Let $G(p) := p M1 + (1 - p) M2$ be the one-shot average game, and let $G^*(p)$ be the supergame when the row player is informed of the chosen matrix and the column player is not (this is called “information on one side”).

Theorem (Aumann–Maschler 1966). *The minimax value function of the repeated two-person zero-sum game with information on one side G^* equals the concavification of the minimax value function of the one-shot average game G :*

$$\text{val } G^* = \widehat{\text{val } G}.$$

That is, for every prior probability p , let $w(p) := \text{val } G(p)$ be the minimax

²⁶In all subsequent periods, he “forgets” his private information and just plays the same action as in the first period.

²⁷The posteriors contain more information than the prior (they are closer to full information, i.e., to 1 and 0, respectively).

²⁸The result applies to any number of matrices; we present it here for two matrices for simplicity.

value of the one-shot average game $G(p)$; let \hat{w} denote the “concavification” of w , i.e., the minimal concave function that is everywhere greater than or equal to w (i.e., $\hat{w}(p) \geq w(p)$ for all $0 \leq p \leq 1$); the Aumann–Maschler Theorem says that $\text{val } G^*(p)$, the value of the repeated game $G^*(p)$, is precisely $\hat{w}(p)$, the evaluation at p of the concavification \hat{w} of w . Figure 2 illustrates this for our third example, Game 3 (note how the value of $G^*(p)$ depends on the values of $G(p')$ for p' -s that are different from p).




Figure 2: Value function of the one-shot average game (left), and of the repeated game (right, bold line) in Game 3

This is a striking result. It tells the precise amount of information that is optimal for the informed player to reveal. In some cases, he should use the information fully, by playing differently according to his information: a “completely revealing” or “separating” strategy. In other cases, he should conceal the information and play the same actions no matter what his information is: a “non-revealing” or “pooling” strategy. And there are still other cases where it is strictly better for him to use the information partially, by mixing it with an appropriate random “noise;” this yields posterior probabilities that are more informative than the prior probabilities, but not fully informative. The difficult problem turns out to have a precise, elegant—and

at times surprising—solution.²⁹

Feasibility and Strategic Communication

Unlike the complete information case, where the feasible outcomes were all those obtained from an action combination or from a (weighted) average of these, here one needs to take into account the informational issues. Not every physically possible outcome can arise in equilibrium; for instance, if an informed player is called upon to reveal his information, he needs to be motivated to do so.

Consider the following example, which has the same structure as the three games above:

<i>Game 4</i>			
		L	R
M1 (probability = $\frac{1}{3}$) :	T	3, 3	0, 0
	B	3, 3	0, 0
		L	R
M2 (probability = $\frac{2}{3}$) :	T	4, 0	3, 3
	B	4, 0	3, 3

Two players repeatedly play a one-shot game with payoffs given by either the bimatrix M1 or the bimatrix M2 (this is now a “non-zero-sum game”). M2 is twice as likely as M1, and only the row player is informed of whether M1 or M2 is the true bimatrix; finally, the actions, but not the payoffs, are observed after each stage.³⁰

²⁹The Aumann–Maschler result provides optimal strategies also for the *uninformed* column player; we do not discuss this here.

³⁰The payoffs correspond to a standard signalling setup, as in principal-agent interac-

The outcome $(3, 3)$ appears in both bimatrices; to obtain it, the informed row player must communicate to the uninformed column player whether the true bimatrix is M1 or M2, e.g., by playing T in the first period if it is M1, and B if it is M2; after T the column player will play L, and after B he will play R. However, this plan is *not* “incentive-compatible”: the row player will “cheat” and play T in the first period also when the true bimatrix is M2, since he prefers that the column player choose L rather than R also when it is M2 (he gets a payoff of 4 rather than 3)! Thus the communication of the row player cannot be trusted. It turns out that the row player has no way to transmit his information *credibly* to the column player, and the outcome $(3, 3)$ is *not* reachable. (The only feasible outcome is for the column player always to play R, which yields an expected payoff vector of only $(2, 2)$.) Thus the presence of asymmetric information hinders cooperation—even when it is to the mutual advantage of both players. The informed player wants to convey his information, but he cannot do so credibly.

The above example is due to Aumann, Maschler, and Stearns (1968), where this theory is first developed. That paper exhibits equilibria that consist of a number of periods of “information transmission” (by the informed player) interspersed with randomizations (by both players), which ultimately lead to certain combinations of actions (like (i) in the section *The Classical Folk Theorem* above); this whole “plan” (of communications and actions) is sustained in equilibrium by appropriate punishments when a deviation is observed (like (ii) in the same section, based on the individual rationality results above). The complete characterization of all equilibria, and their canonical representation in terms of communications, joint plans, and punishments, is provided by Hart (1985) (see also Aumann and Maschler 1995 [Postscript to Chapter 5] and, for a related setup, Aumann and Hart 2003).

The study of repeated games of incomplete information—which has flourished since the pioneering work of Aumann and Maschler—clarifies in a beautiful way the strategic use of information: how much to reveal, how much to

tions: the row player possesses the information and the column player determines the outcome (the two rows yield identical payoffs, so the informed player’s actions have no direct effect on the outcome).

conceal, how much of the revealed information to believe.

Summary – Repeated Games

The model of repeated games is an extremely simple, fundamental, and universal model of multi-stage interactions. It allows participants truly to interact and react to one another; their behavior may be simple—or highly intricate and complex. The analysis is deep and challenging, both conceptually and technically. In the end, the results are elegant and most insightful: simple and natural behavior emerges.

Aumann has been the leader in this area. The highlights of his contribution are (in chronological order):

- (A) *The initial study of repeated games*: showing how to analyze repeated games, and, most importantly, showing how repeated interaction yields cooperative outcomes (the classical Folk Theorem, and Aumann 1959).
- (B) *Asymmetric information*: introducing the essential ingredient of information, and showing how to use information optimally and rationally in long-run strategic interactions (Aumann and Maschler 1966, and Aumann, Maschler, and Stearns 1968).
- (C) *Credible threats and perfectness*: making the equilibria more natural and robust and thus of much wider applicability, which paved the way for their wide use, in particular in economics (Aumann and Shapley 1976, and Rubinstein 1976).

Each one of these three on its own is a landmark contribution; taken together, they complement and strengthen one another, providing a cohesive and significant “big picture”: the evolution of cooperative (and other) patterns of behavior in repeated interactions between rational utility-maximizing individuals.

For further reading, see Aumann’s highly influential surveys and lecture notes throughout the years (1967, 1981, 1985, 1987b, 1992); and, more recently, Mertens, Sorin, and Zamir (1994), the *Handbook of Game Theory*,

with *Economic Applications*, in particular the chapters of Sorin (1992), Zamir (1992), and Forges (1992), the extended postscripts of Aumann and Maschler (1995), and the chapters in Part C of Hart and Mas-Colell (1997).

III. Knowledge, Rationality and Equilibrium

In this section we discuss the two topics of correlated equilibrium and interactive knowledge, together with a beautiful connection between them—all introduced by Aumann.

Correlated Equilibrium

Consider a given and known game, and assume that, before playing it, each player receives a certain signal that does not affect the payoffs of the game. Can such signals affect the outcome?

Indeed, they can: players may use these signals to “correlate” their choices. Aumann (1974) defined the notion of *correlated equilibrium*: it is a Nash equilibrium of the game with the signals. When the signals are (stochastically) independent, this is just a Nash equilibrium of the given game. At the other extreme, when the signals are perfectly correlated—for instance, when everyone receives the same “public” signal (such as “sunspots”)—it amounts to an average of Nash equilibria (this is called a *publicly correlated equilibrium*). But in general, when the signals are private and partially correlated, it leads to new equilibria, which may lie outside the convex hull of the Nash equilibria of the game.

For example, take the two-person “Chicken” game

	LEAVE	STAY
LEAVE	4, 4	2, 5
STAY	5, 2	0, 0

There are two pure Nash equilibria, (STAY, LEAVE) and (LEAVE, STAY), and also a mixed Nash equilibrium where each player plays LEAVE with

probability $\frac{2}{3}$ and STAY with probability $\frac{1}{3}$; the payoffs are, respectively, $NE1 = (5, 2)$, $NE2 = (2, 5)$, and $NE3 = (3\frac{1}{3}, 3\frac{1}{3})$; see Figure 3.




Figure 3: Nash equilibria and correlated equilibria for the game of Chicken (see text; figure not drawn to scale)

Consider now a public fair-coin toss (i.e., the common signal is either “H” or “T,” with probabilities $\frac{1}{2} - \frac{1}{2}$); assume that after “H” the row player plays STAY and the column player plays LEAVE, whereas after “T” they play LEAVE and STAY, respectively. This is easily seen to constitute a Nash equilibrium of the extended game (with the signals), and so it is a publicly correlated equilibrium of the original Chicken game. The probability distribution of

outcomes is

	LEAVE	STAY
LEAVE	0	$\frac{1}{2}$
STAY	$\frac{1}{2}$	0

and the payoffs are $CE1 = (3\frac{1}{2}, 3\frac{1}{2})$ (see Figure 3; the set of payoffs of all publicly correlated equilibria, which is the convex hull of the Nash equilibrium payoffs, is the grey area there).

But there are other correlated equilibria, in particular, one that results in each action combination except (STAY, STAY) being played with an equal probability of $\frac{1}{3}$:

	LEAVE	STAY
LEAVE	$\frac{1}{3}$	$\frac{1}{3}$
STAY	$\frac{1}{3}$	0

Indeed, let the signal to each player be “L” or “S”; think of this as a “recommendation” to play LEAVE or STAY, respectively. When the row player gets the signal L, he assigns a (conditional) probability of $\frac{1}{2}$ to each one of the two pairs of signals (L, L) and (L, S); so, if the column player follows his recommendation, then the row player gets an expected payoff of $3 = \frac{1}{2} \cdot 4 + \frac{1}{2} \cdot 2$ from playing LEAVE, and only of $2\frac{1}{2} = \frac{1}{2} \cdot 5 + \frac{1}{2} \cdot 0$ from deviating to STAY. When the row player gets the signal S, he deduces that the pair of signals is necessarily (S, L), so if the column player indeed follows his recommendation and plays LEAVE then the row player is better off choosing STAY. Similarly for the column player. So altogether both players always follow their recommendations, and we have a correlated equilibrium. The payoffs of this correlated equilibrium are $CE2 = (3\frac{2}{3}, 3\frac{2}{3})$, which lies outside the convex hull of the Nash equilibrium payoffs (again, see Figure 3).

The notion of correlated equilibrium is most natural and arises in many setups. It is embodied in situations of communication and coordination,

when there are mediators or mechanisms, and so on. In fact, signals are all around us, whether public or private; it is unavoidable that they will find their way into the equilibrium notion.

Common Knowledge and Interactive Epistemology

In multi-agent interactive environments, the behavior of an agent depends on what he knows. Since the behavior of the other agents depends on what they know, it follows that what an agent knows about what the others know is also relevant; hence the same is true of what one knows about what the others know about what he knows—and so on. This leads to the notions of interactive knowledge and interactive beliefs, which are fundamental to the understanding of rationality and equilibrium.

An “event” E is *common knowledge* among a set of agents if everyone knows E , and everyone knows that everyone knows E , and everyone knows that everyone knows that everyone knows E , and so on. This concept was introduced by the philosopher David Lewis (1969) and, independently, by Aumann (1976). Aumann’s work went far beyond this, formalizing the concept and exploring its implications.

Formally, the information of an agent in a certain “state of the world” ω is described by the set of all states that he considers possible when ω is the true state, i.e., those states ω' that his information does not allow him to distinguish from ω . An *event* E is identified with the set of states where it occurs. The agent *knows* an event E at a state ω if all states that he considers possible at ω lie in E (i.e., E occurs at all these states). The event E is *common knowledge* at a state ω if, at ω , all agents know E , and all agents know (the event) that all know E , and so on (as Aumann showed, all this can be succinctly expressed using the algebra of partitions).

The fascinating result in Aumann’s paper is

The Agreement Theorem (Aumann 1976). *Consider two people who start with the same prior beliefs and may receive different information. If their posterior probabilities for some event A are common knowledge, then these posteriors must be identical.*

That is, in a state where each one knows the other one's posterior, and knows that the other one knows his own posterior, and knows that the other one knows that he knows the other one's posterior, and so on (and all this has been taken into account when computing the posteriors), the agents must have identical posteriors. Formally, let α and β be two numbers, and consider all states where the posterior probability for A of the first agent is α , and that of the second agent is β ; this set of states is an event, call it E . The Agreement Theorem says that if E is common knowledge at some state, then necessarily $\alpha = \beta$.

This result was the starting point of a whole new area, called *interactive epistemology*; it generated a huge literature, in economics (for instance, the “no trade” results), computer science, logic, philosophy, and beyond. Aumann went on to provide solid foundations (the so-called “syntactic approach”) of this whole area; see Aumann (1999a, 1999b).³¹ It is deep work that deals with basic questions, such as, Why can one assume without loss of generality that the information structures are commonly known?

Rationality and Equilibrium

In 1987 Aumann established an intriguing connection between two notions: correlated equilibrium and common knowledge of rationality.

Theorem (Aumann 1987a). *Consider a game whose players start with the same prior beliefs and may receive different information. If it is common knowledge that all players are Bayesian rational, then they are playing a correlated equilibrium of the game.*

A player is “Bayesian rational” if his action is optimal given his information. Let R denote the set of all states where all players are Bayesian rational (formally, R is an event), and assume that R is common knowledge at some state. The Aumann (1987a) Theorem says that the prior distribution yields a correlated equilibrium of the given game.

³¹These papers were circulated as lecture notes in the eighties.

It is remarkable that correlated equilibria are obtained from common knowledge of rationality and a common prior only; no assumptions of equilibrium behavior are needed.³²

IV. Perfect Competition

Perfect competition has always been a central theme in economics. It refers to a situation where there are many participants and the influence of each individual on the economy as a whole is negligible. But how should perfect competition be modelled?

Following various attempts,³³ Aumann (1964) introduced the general model of a *continuum* of agents.

“The influence of an individual participant on the economy cannot be mathematically negligible, as long as there are only finitely many participants. Thus a mathematical model appropriate to the intuitive notion of perfect competition must contain infinitely many participants. We submit that the most natural model for this purpose contains a continuum of participants, similar to the continuum of points on a line or the continuum of particles in a fluid. [...]

The continuum can be considered an approximation to the “true” situation in which there is a large but finite number of particles (or traders, or strategies, or possible prices). The purpose of adopting 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).” (Aumann 1964 [Section 1])

³²Regarding Nash equilibrium, the precise epistemic requirements are studied in Aumann and Brandenburger (1995) and Aumann (1995).

³³For instance, Shubik (1959a) and Debreu and Scarf (1963).

The “continuum of agents” idea has turned out to be indispensable to the advancement of economic theory. As in the natural sciences, it has enabled precise and rigorous modelling and analysis, which otherwise would have been very hard or even impossible. Once the continuum model was studied and understood, one could go back and examine the “true” situations with finitely many agents, using appropriate approximations and limits.

The foremost result is the equivalence between *competitive* (or *Walrasian*) *equilibria* and *core* allocations in perfectly competitive markets: individual price-taking behavior and social stability yield the same outcomes (Aumann 1964; see the survey of Anderson 1994). This may be viewed as a modern and rigorous formalization of the “invisible hand” of Adam Smith and the limit “contract curve” of Edgeworth.

The core is a classical concept of cooperative game theory (recall also the Strong Folk Theorem above). Another central cooperative concept is that of *value*, originally introduced by Shapley (1953). The notion of value embodies standard economic ideas such as “expected outcome” and “marginal contribution.” Applying it to markets with a continuum of traders yielded a new equivalence result, this time between the price equilibria and the value allocations (Aumann and Shapley 1974, Aumann 1975; see the survey of Hart 2002).

All this led to an extensive study of the relationships between competitive equilibria and various game-theoretic concepts, which led to a very general “equivalence principle.” To quote Aumann:

“Perhaps the most remarkable single phenomenon in game and economic theory is the relationship between the price equilibria of a competitive market economy, and all but one³⁴ of the major solution concepts for the corresponding game. [...]

Intuitively, the equivalence principle says that the institution of market prices arises naturally from the basic forces at work in a market, (almost) no matter what we assume about the way in

³⁴The one exception is the *stable set*; see Hart (1974) (more recently, additional exceptions are certain NTU-values; see Hart 2002 [Section 5.4]).

which these forces work.” (Aumann 1987b [Section “1960–1970,” v])

The end result of the introduction of the continuum model in economics is that it has led to a much better understanding of perfect competition, its implications and limitations, from many disparate but converging viewpoints.

V. Other Contributions

In the previous three sections we presented three areas of Aumann’s landmark contributions. But Aumann’s pioneering work goes far beyond these. We can only list here some of the topics: cooperative games (non-transferable utility games, value, core, bargaining sets, nucleolus, consistency, bankruptcy and the Talmud, coalition formation, axiomatization), subjective probability and utility, power and taxes, games in extensive form, and mathematics (for example, set-valued functions).

But that is not all. Not only did Aumann lay down and deepen the “micro” foundations, he always espoused and emphasized the “macro” outlook. Aumann has been highly influential in providing the philosophical and conceptual groundwork for game theory (see, e.g., Aumann 1985 and 1987b). Of particular significance is his consistent promotion of the view of a *unified game theory*:

“Game Theory may be viewed as a sort of umbrella or “unified field” theory for the rational side of social science, where “social” is interpreted broadly, to include human individuals as well as other kinds of players (collectives such as corporations and nations, animals and plants, computers, etc.). 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.” (Aumann and Hart 1992 [Preface, pp. xi–xii])

VI. Aumann Inspires and Teaches

An overview of Aumann’s contributions would not be complete were it to cover just his own papers. Indeed, Aumann has greatly influenced and inspired generations of researchers, directly and indirectly: his students, collaborators, colleagues, and the many people who read his work, listened to his lectures, or conducted discussions with him. His papers are masterly written, combining deep analysis with informal explanations and intuitions; his surveys on game theory in general, and on various specific topics, have largely shaped the field. Aumann has been the driving force behind the development of many areas (such as those we have presented above). He has always applied the most rigorous standards, to his as well as others’ work; at the same time, he never lost sight of the emerging “big picture”: he aimed for it, and directed everyone toward it. (For a fascinating glimpse of Aumann’s scientific as well as personal world, see his interview in Hart 2005.)

Aumann has had, to date, thirteen doctoral students: Bezalel Peleg (Ph.D., 1964), David Schmeidler (1969), Shmuel Zamir (1971), Elon Kohlberg (1973), Benyamin Shitovitz (1974), Zvi Artstein (1974), Eugene Wesley (1975), Sergiu Hart (1976), Abraham Neyman (1977), Yair Tauman (1979), Dov Samet (1981), Ehud Lehrer (1987), and Yossi Feinberg (1997). All but one (who pursued a non-academic career) are now leading researchers in their own right. Aumann is very proud of his scientific “family”—as he calls it—and we are all proud to be part of it.

VII. A Short Biography of Aumann

Robert John Yisrael Aumann³⁵ was born in 1930 in Frankfurt, Germany. In 1938 the Aumann family left Germany and came to America. Robert

³⁵Aumann’s curriculum vitae and publications are available at <http://www.ma.huji.ac.il/~raumann>.

Aumann studied at City College in New York, and then at the Massachusetts Institute of Technology (MIT), where he got his doctoral degree in pure mathematics in 1955. After two years as a post-doc at Princeton University, he immigrated to Israel in 1956, and has been at the Hebrew University of Jerusalem since then. Though now officially retired, Aumann continues his scientific work, teaches, and supervises students.

In 1991 the multi-disciplinary Center for the Study of Rationality was established at the Hebrew University of Jerusalem; the two “founding fathers” were Aumann and Menahem Yaari. Today Aumann chairs its Academic Committee. Aumann was also instrumental in the foundation of the international Game Theory Society, and served as its first president from 1999 to 2003.

Aumann has received many awards and prizes, including the Harvey Prize, the Israel Prize, the Lanchester Prize, the Nemmers Prize, the EMET Prize, the von Neumann Prize, and the Nobel Prize.

References

- Anderson, R. M. (1994), The Core in Perfectly Competitive Economies, Chapter 14 in Aumann and Hart, Vol. 1 (1994), 413–457.
- Aumann, R. J. (1959), 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*, Princeton University Press, Princeton, NJ, 287–324.
- Aumann, R. J. (1960), Acceptable Points in Games of Perfect Information, *Pacific Journal of Mathematics 10*, 381–417.
- Aumann, R. J. (1961), The Core of a Cooperative Game without Side Payments, *Transactions of the American Mathematical Society 98*, 539–552.
- Aumann, R. J. (1964), Markets with a Continuum of Traders, *Econometrica 32*, 39–50.

- Aumann, R. J. (1967), A Survey of Cooperative Games without Side Payments, in M. Shubik (ed.), *Essays in Mathematical Economics in Honor of Oskar Morgenstern*, Princeton University Press, Princeton, NJ, 3–27.
- Aumann, R. J. (1974), Subjectivity and Correlation in Randomized Strategies, *Journal of Mathematical Economics* 1, 67–96.
- Aumann, R. J. (1975), Values of Markets with a Continuum of Traders, *Econometrica* 43, 611–646.
- Aumann, R. J. (1976), Agreeing to Disagree, *Annals of Statistics* 4, 1236–1239.
- Aumann, R. J. (1981), Survey of Repeated Games, in R. J. Aumann et al. (eds.), *Essays in Game Theory and Mathematical Economics in Honor of Oskar Morgenstern*, Vol. 4 of Gesellschaft, Recht, Wirtschaft, Wissenschaftsverlag, Bibliographisches Institut, Mannheim, 11–42.
- Aumann, R. J. (1985), What Is Game Theory Trying to Accomplish?, in K. J. Arrow and S. Honkapohja (eds.), *Frontiers of Economics*, Basil Blackwell, Oxford, 28–76.
- Aumann, R. J. (1987a), Correlated Equilibrium as an Expression of Bayesian Rationality, *Econometrica* 55, 1–18.
- Aumann, R. J. (1987b), Game Theory, in J. Eatwell, M. Milgate, and P. Newman (eds.), *The New Palgrave, A Dictionary of Economics*, Macmillan, London and Basingstoke, Vol. 2, 460–482.
- Aumann, R. J. (1989), *Lectures on Game Theory*, Underground Classics in Economics, Westview Press, Boulder, CO.
- Aumann, R. J. (1992), Rationality and Bounded Rationality, in Y. Moses (ed.), *Theoretical Aspects of Reasoning about Knowledge, Proceedings of the Fourth Conference (TARK 1992)*, Morgan Kaufmann Publishers, San Mateo, CA, 108–117; also *Games and Economic Behavior* 21 (1997), 2–14.

- Aumann, R. J. (1995), Backward Induction and Common Knowledge of Rationality, *Games and Economic Behavior* 8, 6–19.
- Aumann, R. J. (1999a), Interactive Epistemology I: Knowledge, *International Journal of Game Theory* 28, 263–300.
- Aumann, R. J. (1999b), Interactive Epistemology II: Probability, *International Journal of Game Theory* 28, 301–314.
- Aumann, R. J. and Brandenburger, A. (1995), Epistemic Conditions for Nash Equilibrium, *Econometrica* 63, 1161–1180.
- Aumann, R. J. and Hart, S. (eds.) (1992, 1994, 2002), *Handbook of Game Theory, with Economic Applications*, Elsevier/North-Holland, Amsterdam; Vol. 1, 1992; Vol. 2, 1994; Vol. 3, 2002.
- Aumann, R. J. and Hart, S. (2003), Long Cheap Talk, *Econometrica* 71, 1619–1660.
- Aumann, R. J. and Maschler, M. (1966), Game Theoretic Aspects of Gradual Disarmament, Mathematica, Inc. Report ST-80, V1–V55, Princeton. Reprinted as Chapter I in Aumann and Maschler (1995).
- Aumann, R. J. and Maschler, M. (1967), Repeated Games with Incomplete Information: A Survey of Recent Results, Mathematica, Inc. Report ST-116, 287–403, Princeton. Reprinted as Chapter II in Aumann and Maschler (1995).
- Aumann, R. J. and Maschler, M. (1968), Repeated Games of Incomplete Information: The Zero-Sum Extensive Case, Mathematica, Inc. Report ST-143, 37–116, Princeton. Reprinted as Chapter IV in Aumann and Maschler (1995).
- Aumann, R. J. and Maschler, M. (1995), *Repeated Games of Incomplete Information*, MIT Press, Cambridge, MA.

- Aumann, R. J., Maschler, M., and Stearns, R. E. (1968), Repeated Games of Incomplete Information: An Approach to the Non-Zero-Sum Case, *Mathematica*, Inc. Report ST-143, 117–216, Princeton. Reprinted as Chapter V in Aumann and Maschler (1995).
- Aumann, R. J. and Shapley, L. S. (1974), *Values of Non-Atomic Games*, Princeton University Press, Princeton, NJ.
- Aumann, R. J. and Shapley, L. S. (1976), Long-Term Competition – A Game-Theoretic Analysis (mimeo). Reprinted in N. Megiddo (ed.), *Essays in Game Theory*, Springer-Verlag (1994), New York, 1–15.
- Debreu, G. and Scarf, H. (1963), A Limit Theorem on the Core of an Economy, *International Economic Review* 4, 235–246.
- Forges, F. (1992), Repeated Games of Incomplete Information: Non-Zero-Sum, Chapter 6 in Aumann and Hart, Vol. 1 (1992), 155–177.
- Friedman, J. (1971), A Noncooperative Equilibrium for Supergames, *Review of Economic Studies* 38, 1–12.
- Harsanyi, J. C. (1967–1968), Games with Incomplete Information Played by Bayesian Players, Parts I, II, III, *Management Science* 14, 159–182, 320–334, 486–502.
- Hart, S. (1974), Formation of Cartels in Large Markets, *Journal of Economic Theory* 7, 453–466.
- Hart, S. (1985), Nonzero-Sum Two-Person Repeated Games with Incomplete Information, *Mathematics of Operations Research* 10, 117–153.
- Hart, S. (2002), Values of Perfectly Competitive Economies, Chapter 57 in Aumann and Hart, Vol. 3 (2002), 2169–2184.
- Hart, S. (2005), An Interview with Robert Aumann, *Macroeconomic Dynamics* 9, 683–740.

- Hart, S. and Mas-Colell, A. (eds.) (1997), *Cooperation: Game Theoretic Approaches*, Springer-Verlag, New York.
- Lewis, D. (1969), *Convention*, Harvard University Press, Cambridge, MA.
- Luce, R. D. and Raiffa, H. (1957), *Games and Decisions*, Wiley, New York.
- Mertens, J.-F., Sorin, S., and Zamir, S. (1994), Repeated Games, DP 9420, 9421, 9422, CORE, Université Catholique de Louvain (mimeo).
- Nash, J. F. (1951), Non-Cooperative Games, *Annals of Mathematics* 54, 286–295.
- Rubinstein, A. (1976), Equilibrium in Supergames (mimeo). Reprinted in N. Megiddo (ed.), *Essays in Game Theory*, Springer-Verlag (1994), New York, 17–28.
- Rubinstein, A. (1979), Equilibrium in Supergames with the Overtaking Criterion, *Journal of Economic Theory* 21, 1–9.
- Selten, R. (1965), Spieltheoretische Behandlung eines Oligopolmodells mit Nachfragerträgeit, *Zeitschrift für die gesamte Staatswissenschaft* 121, 301–324, 667–689.
- Selten, R. (1975), Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games, *International Journal of Game Theory* 15, 101–107.
- Shapley, L. S. (1953), A Value for n -Person Games, in H. W. Kuhn and A. W. Tucker (eds.), *Contributions to the Theory of Games II*, *Annals of Mathematics Study* 28, Princeton University Press, Princeton, NJ, 307–317.
- Shubik, M. (1959a), Edgeworth Market Games, in A. W. Tucker and R. D. Luce (eds.), *Contributions to the Theory of Games IV*, *Annals of Mathematics Study* 40, Princeton University Press, Princeton, NJ, 267–278.

Shubik, M. (1959b), *Strategy and Market Structure*, Wiley, New York.

Sorin, S. (1992), Repeated Games with Complete Information, Chapter 4 in Aumann and Hart, Vol. 1 (1992), 72–107.

Zamir, S. (1992), Repeated Games of Incomplete Information: Zero-Sum, Chapter 5 in Aumann and Hart, Vol. 1 (1992), 109–154.