ACADEMIC PRESS

Available online at www.sciencedirect.com
SCIENCE DIRECT®

Games and Economic Behavior 42 (2003) 327–332

GAMES and Economic Behavior

www.elsevier.com/locate/geb

Reviews and Comments

With the intent of stimulating discussion, this section is reserved for book reviews, comments and letters; your input is welcome. By nature, this material may be subjective, reflecting the opinions of the authors; your responses are therefore encouraged.

The Collected Papers of Robert J. Aumann. Robert J. Aumann, MIT Press, Cambridge, MA, 2000. Vol. 1, 720 pp; Vol. 2, 672 pp.

The Collected Papers of Robert J. Aumann is a two-volume set composed of 73 chapters, most of which are individual papers. This remarkable collection includes numerous seminal contributions to both game theory and economics. For example:

- *Cooperative game theory*: "Acceptable points in general cooperative n-person games," Chapter 20; "Von Neumann–Morgenstern solutions to cooperative games without side payments," Chapter 38; "An axiomatization of the non-transferable utility value," Chapter 60.
- *Non-cooperative game theory*: "Mixed and behavior strategies in infinite extensive games," Chapter 28; "Subjectivity and correlation in randomized strategies," Chapter 31; "Agreeing to disagree," Chapter 32; "Correlated equilibrium as an expression of Bayesian rationality," Chapter 33; "Backward induction and common knowledge of rationality," Chapter 36.
- *Decision theory*: "A definition of subjective probability," Chapter 17.
- *Economics*: "Markets with a continuum of traders," Chapter 46; "Existence of competitive equilibria in markets with a continuum of traders," Chapter 47.

It is, of course, impossible to do justice to the entire two-volume collection of papers in a single review article. It is my modest hope to convey some of the beauty, scope and significance of the work and to make evident the benefits awaiting those who read the collected papers, as a whole, themselves.

Section 1 is titled "General" and it collects together Aumann's philosophical views on the role of game theory (Chapter 1), a wonderful historical account of the developments of game theory from 1910 to 1986 (Chapter 2), as well as several other pieces.

The philosophical paper is titled "What Is Game Theory Trying to Accomplish?" Most game theorists have probably been asked this question in one form or another, but I suspect that far fewer have actually thought deeply about the matter. Aumann is an exception, and he was bold (confident) enough to put his thoughts in writing.

Aumann's answer to the question he poses in the title is, at first blush, a little surprising. Neither predictive power nor practical applications figure prominently in his view. Rather,

0899-8256/2003 Published by Elsevier Science (USA).
doi:10.1016/S0899-8256(02)00540-7

Aumann suggests that game theory should be judged primarily on its ability to enhance our understanding of the world around us. Indeed, he views this as the central motivation of any science. In his words (Vol. I, p. 6),

> On the most basic level, what we are trying to do in science is to understand our world. Predictions are an excellent means of testing our comprehension, and once we have the comprehension applications are inevitable; but the basic aim of scientific activity remains the comprehension itself.

Aumann goes on to discuss what he calls the three components of comprehension: relationships, unification, and simplicity. These ideas are beautifully presented and well worth reading.

While Aumann's view that comprehension—more so than practical applications or predictive power—is at the heart of scientific inquiry is surely defensible, even commendable, a skeptic might be hard-pressed not to conclude that such a view is born out of necessity. "After all," s/he would say, "game theory has almost no predictive power and few practical applications."

But Aumann does not concede this much. For example, while he acknowledges that rationality is but one of a number of factors affecting human behavior, he also holds the view that (Vol. I, p. 12) "... the cumulative effect of numbers and time and learning ... pushes people 'in general' in the direction of rational decision-making." Moreover, Aumann observes that a vast amount of economics owes much of its clarity and usefulness to the development of game theory (e.g., models of search, entry and exit, product quality, auctions, insurance, principal agent relationships, etc.). Thus, Aumann argues, the unquestionably large body of economic modeling with game theory at its core is an important and relevant measure of game theory's success.

While the reader will have to consult the original paper to obtain a full airing of Aumann's carefully articulated view, let me conclude by recommending this paper for any game theory reading list. Students (and professors!) will enjoy it, and whether or not they ultimately agree with it, their own view of what game theory is trying to accomplish will be the better for it.

The second chapter in this section is titled "Game Theory," and it appears in the *New Palgrave Dictionary of Economics* under the entry with the same name. Here the reader will find a concise historical account of some of the best ideas in game theory from 1910 to 1986. As one would expect, there are accounts of the theorems of Zermelo, Kuhn, and von Neumann. Less expected, however, is a description of the connection between game theory and the foundations of mathematics! Apparently, assuming that every infinite zero-sum game of perfect information has a value contradicts the axiom of choice. Moreover, replacing the axiom of choice with this "value" assumption provides an alternative axiomatization for set theory. How wonderful! It is tidbits such as this, sprinkled here and there, that make this entire collection a delight to read.

Von Neumann and Morgenstern's *Theory of Games and Economic Behavior* is also nicely summarized here. There is an especially cogent treatment of the cooperative theory, where solutions, interpretations, and applications are provided. Later we hear of Nash's (1951) contribution as well as those of Harsanyi (1967, 1973), Selten (1965, 1975), Maynard Smith (1982) and a host of others. Throughout, Aumann's own insights connect each work with the next, providing context and motivation. This historical retrospective

might well have been titled "What Game Theory Has Accomplished." It is a terrific complement to its philosophical cousin.

Section 2 is devoted to a single paper, Aumann's doctoral thesis, titled "Asphericity of Alternating Knots." I do not often wish that I had taken a course in algebraic topology, but this is one such instance. I might then be capable of fully appreciating this apparently deep result.

Section 3 is "Decision Theory: Utility and Subjective Probability." Among the contents here are Aumann's papers on utility theory without the completeness axiom (Chapter 14), the well-known paper with Anscombe on subjective probability (Chapter 17), some early work on measurable selection theorems (Chapter 15), and a correspondence between Aumann and Savage on a conceptual aspect of Savage's utility theory (Chapter 18). The conceptual issue is related to state dependent utilities. Aumann's letter provides several interesting examples, all of which pose a problem of sorts for Savage's axioms. Savage's response reinforces passages from his *Foundations of Statistics* (F. of S.), where the consideration of inconsistent consequences, such as, "being dry in the rain," is proposed as a method for eliciting subjective probabilities. Savage's letter leaves little doubt that his position on the matter had not changed substantially since the publication of F. of S. He writes (Vol. I, p. 307):

> To some—perhaps to you—it will seem grotesque if I say that I should not mind being hung so long as it be done without damage to my health or reputation, but I think it desirable to adopt such language so that the danger of being hung can be contemplated in the framework of F. of S.

Both Aumann's letter and Savage's reply are delightful to read in their entirety. It is not only revealing to find that Aumann sought further clarification on the matter, but a privilege to get a sense of Savage's own doubts about a deep issue that he clearly struggled with.

Sections 4 and 5 are about repeated and extensive games, respectively. What appears to be Aumann's first use of correlated strategies can be found in the two papers on acceptable points in cooperative games (Chapters 20 and 21). There, correlated strategies are employed within an n-person game to express cooperation among the players. The main result establishes a connection between an equilibrium of the static game, expressing a kind of reactive stability to deviating coalitions (this later came to be called the β-core of the static game; see Chapter 39), and the set of *strong equilibria* (i.e., those in which coalitions cannot profitably deviate) in correlated strategies of the repeated game.

It is quite interesting to learn of Aumann's views at that time on the use of correlated strategies. He writes (Vol. I, p. 322):

> In general, to make use of a correlated strategy vector, the players have to agree to consult the *same* random device that chooses the pure strategies they will play. Correlated strategy vectors are thus not usable in non-cooperative games.

This natural view would, as we know, be refined when Aumann (roughly fifteen years later) introduces the notion of *correlated equilibrium*, wherein the use correlated strategies and non-cooperative behavior *are* compatible. The progression of ideas is delightfully evident.

In his "Survey of Repeated Games" (Chapter 23), Aumann provides a very informative description of the theory of repeated games of incomplete information (Vol. I, pp. 423–

429). A special case of the zero-sum theory is explained beautifully with very little notation, yet in a manner that leaves out no essential details. This is quite an accomplishment.

Several appendixes in this survey contain open problems and further examples. One of the examples concerns the effect of finite complexity in the repeated prisoner's dilemma. It is well known that the "always defect" equilibrium of this repeated game survives a vast array of equilibrium refinements; it is notoriously difficult to eliminate, despite its gross inefficiency. Aumann presents an example (Vol. I, p. 435) showing that if the players can condition their current action on at most the opponent's previous action, then *only* the cooperative outcome survives iterative elimination of weakly dominated strategies. Further, the unique surviving strategy is none other than tit-for-tat.

This example eventually led to Aumann and Sorin's study (Chapter 24) of cooperation in games of common interest with bounded recall, a beautiful paper combining perturbation ideas á la Kreps–Milgrom–Roberts–Wilson, bounds on complexity, and common interests, to yield a non-cooperative theory of coordination and efficiency. The discussion at the end of this paper has always, in my mind, been a model of clarity. The authors lay bare the lovely insights that make their result work, and carefully point out the difficulties that would arise if one or another hypothesis were dropped. This kind of in-depth epilogue discussion is a marvelously helpful device that is characteristic of Aumann's papers.

The final section, Section 6, of this first volume contains Aumann's work on equilibrium and knowledge. There are several classic papers here. I especially appreciate the sequential presentation (Chapters 31–33) of "Subjectivity and Correlation in Randomized Strategies" (SC), "Agreeing to Disagree" (AD), and "Correlated Equilibrium as an Expression of Bayesian Rationality" (CE). This gives the reader a birds-eye view of this deep, important, and fruitful research program initiated by Aumann.

In SC one sees the seeds of ideas that will be further developed in the papers to follow. For example, because differences in subjective probabilities play an important role in SC, Aumann is led to ask whether such differences can persist even after the agents reveal their differing probability assessments to one another. Aumann observes that (Vol. I, p. 587) "... some revision of probabilities would certainly be called for ... but whether such a revision must always lead to equal probabilities ... is not clear."

Of course, the matter would be clarified magnificently in the second paper, AD, a paper that has spawned the vast, and continuing, literature on the epistemic foundations of game theory.

Another example of an idea that would blossom concerns one of the main results in SC (Proposition 5.1, p. 572). The result says that, in a two-person game, so long as the players agree about probability zero events, the set of equilibrium payoffs is the same whether or not the players have objective or subjective probabilities, i.e., whether or not they agree on the probabilities of all events. What is really interesting here is the proof. Aumann shows how, from any subjective equilibrium (in which there might be disagreement about the probabilities of some events), one can construct an objective one (i.e., an ordinary one in mixed strategies). The revealing construction is as follows: player i's mixed strategy in the objective equilibrium is precisely player j's subjective beliefs about the pure strategy that i employs in the subjective equilibrium.

Thus, expressed here is what will eventually become the dominant interpretation of mixed strategy equilibria; i's mixed strategy represents j's beliefs about the pure strategy

i employs. But, as the focus was elsewhere, no explicit mention of this interpretation of objectively mixed strategies appears in SC, which seems to have been written concurrently with Harsanyi's (1973) paper on purification. Evidently, this wonderful idea was in the air at the time.

The second volume contains Aumann's work on cooperative games, or, as he prefers to call them, *coalitional* games, as well as his work on markets with a continuum of traders. Prior to reading this volume, I was unaware (though not surprised to learn) that Aumann was one of the founding fathers (together with Peleg, Shapley, and Shubik) of the study of NTU games. Chapter 38 contains several of Aumann and Peleg's basic definitions and results from which the field "took off." Chapter 41 is Aumann's survey of the topic. There are also two entire sections, 10 and 11, devoted to the Shapley value. The first focuses on TU settings, while the second on NTU settings.

Section 11, the NTU section, contains an informative series of papers on the NTU Shapley value. Chapter 60 is Aumann's classic paper "An Axiomatization of the Non-Transferable Shapley Value," in which he provides the first characterization of the NTU value without reference to an underlying TU game. We are then treated to a lively debate on the merits/pitfalls of the NTU value, with Roth and Schafer, separately, and Scafuri and Yannelis, jointly, arguing "against" and Aumann arguing "in favor" of the NTU value. Several counterintuitive examples are carefully discussed and analyzed, and issues ranging from the general import of such examples to the role of solutions in game theory are touched upon. This is a discussion of the highest order.

The two papers "Markets with a Continuum of Traders," Chapter 46, and "Existence of Competitive Equilibria in Markets with a Continuum of Traders," Chapter 47, provide a remarkably elegant and general expression of the core equivalence theorem. The first of these is the core equivalence result per se, showing that when there are a continuum of consumers, the core and set of competitive equilibria coincide, while the second settles the question of non-emptiness of the core, or equivalently, the existence of competitive equilibrium, in the continuum agent setting.

Aumann's well-known results here rely heavily on new mathematical machinery (i.e., the integration of correspondences; see Chapters 68 and 69) that he developed expressly for this purpose. The relevant and important economic content of these results is that it is only in the idealized continuum agent setting in which the competitive assumption of price taking behavior actually obtains. Moreover, in this setting, one can substantially relax the assumptions under which competitive equilibria exist and under which core equivalence holds. In particular, one need not assume that consumer preferences are convex. Further, in production economies with a continuum of firms (Aumann chose to consider only pure exchange economies), it is similarly not necessary to assume that production sets are convex; fixed costs, in particular, can be accommodated. Hence, Aumann's continuum agent model not only provides the appropriate foundation for price-taking behavior, it also permits the theory to account for non-convexities in both preferences and technologies. This is an extremely important contribution.

I cannot help but mention a paper that has made a lasting impression upon me since I first encountered it as a graduate student. It appears here as Chapter 45 and is titled "A Bankruptcy Problem from the Talmud," joint with M. Maschler. It has every ingredient of which great scholarly works are composed. It confronts a long-standing and difficult

problem—the interpretation of a particularly vexed passage from the Talmud that had puzzled Talmudic scholars for over two millennia!—and presents a stunning solution obtained from a totally unexpected source, namely, from the nucleolus, a rather technical game theoretic solution concept. As if this were not enough, the authors, in addition, provide several entirely new perspectives from which to view the game theoretic solution, revealing not only new insights into the theoretical solution itself, but ultimately providing a deep understanding of the solution to the original problem. This, in my view, is an absolutely must read.

What is game theory trying to accomplish? Aumann's answer, in a nutshell, is *comprehension*. *The Collected Papers of Robert J. Aumann* provides deep insights, careful analysis, and thorough discussions—that is, comprehension—on a remarkable range of game theoretic and economic topics. These are papers to read again and again, with additional readings inevitably revealing gems not previously spotted. This two-volume set is not merely *highly recommended*, it is *essential material* for every game theorist and mathematical economist. The opportunity to learn from one of the great theorists of the twentieth century should not be missed.

Acknowledgment

I thank Motty Perry for providing helpful comments on an earlier draft.

Philip J. Reny
University of Chicago,
Chicago, IL, USA
E-mail address: p-reny@uchicago.edu