

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.

Repeated Games with Incomplete Information. Robert J. Aumann and Michael B. Maschler, with the collaboration of Richard E. Stearns, MIT Press, Cambridge, MA, 1995. ISBN 0-262-01147-6, \$34.95.

ACDA ST/80, ACDA ST/116, ACDA ST/143 . . . these mysterious codes were well known to a small group of game theorists when I discovered them in 1978 while reading a paper by Mertens and Zamir (1971–1972). I was very frustrated not to be able to get three (out of seven) of the references but fortunately, I was lucky enough to meet J.-P. Ponssard, who kindly gave me a copy of these three reports to the Arms Control and Disarmament Agency, prepared by Mathematica, Princeton, New Jersey, in 1966, 1967, and 1968. I was so excited by the topic and the results that I started to write notes describing the basic results in the simplest case (Sorin, 1979). An English translation followed (1980) at a time when Mertens and Zamir had already written three chapters of a book on the subject. A few years later, I joined the project, and this work is still in progress (a preliminary version appeared in 1994).

But now Aumann and Maschler provide us with a tremendous publication: *Repeated Games with Incomplete Information*. This book is first of all a public proof of the existence of five of the chapters of these reports: Game Theoretic Aspects of Gradual Disarmament (Aumann and Maschler, 1966); Repeated Games with Incomplete Information: A Survey of Recent Results (Aumann and Maschler, 1967); A Formal Information Concept for Games with Incomplete Information (Stearns, 1967); Repeated Games of Incomplete Information: The Zero-Sum Extensive Case (Aumann and Maschler, 1968); and Repeated Games of Incomplete Information: An Approach to the Non-Zero-Sum Case (Aumann, Maschler, and Stearns, 1968). It is amazing that these reports, although never published and very difficult to obtain, have frequently been quoted and used and have laid the foundations for an impressive field of research.

But the book contains more: almost a third of it consists of comments on and further explanations and descriptions of new results. It is clearly destined to be the basic reference for all of the topics involving the strategic use of information in multimove games.

The preface explains the origin of the reports to the ACDA and recalls the role played by O. Morgenstern (to whom the book is dedicated) and H. Kuhn in bringing together R. Aumann, G. Debreu, J. Harsanyi, M. Maschler, J. Mayberry, H. Scarf, R. Selten, and R. Stearns to help understand the decision and game theoretic aspects of arms control. Then it shows how, on the basis of two theories, one concerning repeated games that emerged in the early 1960s, the other devoted to incomplete information games and newly created by Harsanyi (1967–1968), a new model was developed: repeated games with incomplete information where some game is repeated but the players have uncertainty about which game it will be.

Formally the basic model is a finite family of games with an initial probability that determines which one of these games will be played. Once a game is chosen, the players may obtain some information about it and the same game is played over and over, in stages. After each stage, the players receive information on their opponents' moves and/or on the game chosen.

Altogether this book is a beautiful introduction to the topic, not only because it presents the original initial proofs but also because it allows the reader to follow the building of the theory step by step and so to see how the main ideas and concepts emerge and then are refined or extended.

The book contains five chapters. Chapters I–IV treat zero-sum games, and the results in the non-zero-sum case given in Chapter V depend heavily on the previous analysis.

Chapter I deals with the simplest case where player I is fully informed about which game is being played, player II has no initial information, and the moves of both players are made public after each stage. The idea that the choice of a move for a given strategy of player I transmits information leads to the notion of nonrevealing and completely revealing strategies. Further it is shown that the use of mixed strategies, corresponding to partial revelation, may achieve better payoffs. The nonrevealing game is then introduced and the basic idea concerning the use of private information is presented: the so-called “splitting lemma” shows that the informed player can generate any martingale as a posterior probability. As a consequence, to analyze a specific situation with some initial probability p on the set of different games, one has to consider all initial probabilities and from this one infers that the amount player I can obtain is a concave function of p . By playing nonrevealingly player I wins at each stage the value $u(p)$ of the nonrevealing game $D(p)$; hence $Cav u(p)$ is a lower bound on his payoff (where the concavification operator Cav , applied to a function f , denotes the smallest concave function greater than f). Consider now $G_n(p)$, the n -stage repeated game. This being basically a finite game, it has a value $v_n(p)$ and the

minmax theorem may be invoked to assume that player II knows the strategy of player I. The extra gain per stage due to the information of player I (compared to his payoff in $D(p)$) may be shown to be bounded by the variation (in L_1 -norm) of the martingale of posteriors (up to some constant). This finally implies that $v_n(p)$ is less than $Cav u(p) + C/\sqrt{n}$ for some constant C depending on the game only. Consider now the infinitely repeated game $G_\infty(p)$: the same bound applies to player I, while the analysis from the point of view of player II requires new tools. To make player II immune against any bluffing strategy of player I, Aumann and Maschler consider the incomplete information game as a game with vector payoffs in the sense of Blackwell (1956), where each component of the payoff corresponds to a specific state. Blackwell's theorem is used to characterize the payoffs that player II can approach and this leads to the proof of the existence of a value $v_\infty(p)$ for $G_\infty(p)$, which $v_\infty(p) = Cav u(p)$ (see also Postscript I.d).

Chapter II studies the situation with lack of information on both sides, where each of the players has some private information on the true state of the game. The independent case is obtained when the state space is a product space $L \times M$ with a product probability $p \otimes q$: with probability $p^l q^m$ the payoff's matrix is $A^{l,m}$, player I is informed of l , and player II is informed of m . As in the previous chapter the nonrevealing game $D(p, q)$ with value $u(p, q)$ is defined. Then by forgetting his initial information player I is in a situation similar to the one studied above: in fact, now he is uninformed (about m) and player II plays the role of the informed player (and because of the independence one can take the average payoff w.r.t. p). Thus player I can get $Vex_q u(p, q)$ (where Vex_q is the Convexification operator in the variable q) in the infinitely repeated game and the "splitting lemma" gives a lower bound of $Cav_p Vex_q u(p, q)$ on his payoff. The proof that he cannot get more relies on a new concept introduced by Stearns in Chapter III. Given a strategy σ of player I in the infinitely repeated game, the amount of information contained in it, $V(\sigma)$, is the maximum of the total variation (in L_2 norm) of the martingale of posterior probabilities $\{p_n\}$ on L , taken over all strategies τ of player II. Explicitly, $V(\sigma) = \sup_\tau E_{\sigma, \tau} (\sum_{m=1}^\infty \sum_l (p_{m+1}^l - p_m^l)^2)$. Note that the supremum can be taken on the sequences of moves of player II; hence given any σ , player II can first play nonrevealingly up to some stage N and can extract almost all the information from σ . Then by the definition of $V(\sigma)$, player I is essentially playing nonrevealingly from stage N on, and using the results of Chapter I again, player II can prevent him from getting more than $Vex_q u(p_N, q)$, hence $Cav_p Vex_q u(p, q)$ on the average. This shows that $Cav_p Vex_q u(p, q)$ is the maxmin of $G_\infty(p, q)$ (and dually $Vex_q Cav_p u(p, q)$ is the minmax). Several examples in Chapter II show that these quantities can differ, so in opposition to the one-sided information case, the infinitely repeated game may have no value. In this framework Mertens and Zamir (1971–1972) have proved that the sequence of values of the n -stage game converges.

Chapter IV returns to the one-sided information case but develops a more complex model (already introduced in Chapter II) motivated by the case where

each stage is a game in extensive form. While it may be natural here to assume that the moves of the players are public, usually the one-stage pure strategies are not. The authors then define, for each state k and for each player $x = \text{I, II}$, in addition to the payoff matrix, a signaling matrix H_k^x with values in some alphabet and assume that after each stage t the letter $H_k^x(i_t, j_t)$ is announced to player x , if the moves were (i_t, j_t) . (This style of analysis is typical of the authors: any refinement in the model is precisely and carefully justified by examples, but then established in the greatest generality.) There are two main changes compared to the analysis of Chapter I. First, the two notions of “not using one’s own information” and of “not revealing it” are no longer equivalent and it appears that the latter is the correct one. This leads to the definition of nonrevealing strategies (that may be state dependent) and of the corresponding nonrevealing game with a value still denoted by u . The fact that player I can always win $Cav u$ is now as in Chapter I. In the proof that asymptotically he cannot obtain more in the finite games (i.e., $\lim v_n = Cav u$), the main steps mirror those of Chapter I, but the analysis is really much more intricate because of the different ways the information may be released. This result shows also that the infinitely repeated game has a value, using the following neat trick (which has a much larger domain of application): player II plays once optimally in the one-stage game $G_1(p)$, then during two stages optimally in $G_2(p)$, then during three stages in $G_3(p)$, and so on, ignoring at each stage the nonrelevant information. It is clear that the average payoff he can guarantee in this way is bounded above by $\lim v_n(p)$, hence the result. However, and this is the second new aspect, the “approachability” strategy of player II cannot be applied since he does not know player I’s move, not even the stage vector payoff. These difficulties were finally overcome by Kohlberg (1975).

Finally, Chapter V is devoted to the non-zero-sum case, and once more several new and fundamental ideas and concepts are introduced. First it is observed that, even if cooperation is possible and better for both players, there are situations where it cannot occur in equilibrium because of incomplete information, lack of trust, and the possibility of cheating. Then, based on the zero-sum analysis, and in a quite similar way to the Folk Theorem, equilibrium payoffs generated by “simple nonrevealing agreements” are defined: they basically consist of a plan (a history to follow) and of threats that support the incentive compatibility conditions for both players. The next step is to show that one can actually construct a whole hierarchy of such agreements (yielding new equilibrium payoffs) by adding two devices: one is signals sent by Player I and related to the “splitting lemma” (in addition, a “no cheating” condition has to be satisfied); the other is the “jointly controlled lottery” and corresponds in modern terminology to a public correlated device generated by the players themselves through their moves. One can look at the complexity of an agreement as being the number of stages of signals and/or joint lottery before reaching a simple nonrevealing agreement. It took almost 20 years before Hart (1985) proved that in fact all equilibrium

strategies do generate a process similar to such a complex agreement, including infinite complexity. The detailed and crystal-clear story of this construction may be found in the 17 pages of Postscripts V.c to V.f. In addition, V.g carefully presents Forges's famous example (1990) requiring unbounded complexity.

There is no way to cover here all the topics discussed in the Postscripts such as, among others, discounting (one regret: the report by Mayberry, 1967, is not included), continuous time, and games without a recursive structure or with identical information. I would just like to point out some of them. The analysis of the conceptual distinction between the large finite game and the infinite game (II.c and II.d) is very profound and enlightening. The study of the value of information and of the monotonicity of v_n (including a very illuminating presentation of Lehrer's counterexample, 1987) (I.e and IV.a) gives deep insights into the meaning of the recursive structure. The case of incomplete information on "one and a half sides" with a nice presentation of my joint work with Zamir (1985) and its connections with the dependent case (II.i and II.j) shows the difference between the information about the true game and the information about the information of the other players. Finally, one finds an extensive discussion of the error term (which corresponds to the speed of convergence of v_n to its limit) (I.c, IV.e), first in the standard signaling case, with a subtle study of the time and size of revelation of information, then in more general situations. One last comment on this topic: the result reported on p. 215, I.11, is proved in Mertens *et al.* (1994, Chap. IV, Corollary 4.9) and the exact bound in the case of lack of information on one side and general signals has been very recently proved by Mertens (1995) to be $\log(n)/n^{1/3}$.

In conclusion, this is a wonderful book and is highly recommended; the talent of its authors makes it enjoyable to newcomers in the field, unforgettable to graduate students, and precious to confirmed game theorists.

REFERENCES

- Blackwell, D. (1956). "An Analogue of the Minmax Theorem for Vector Payoffs," *Pacific J. Math.* **6**, 1–8.
- Forges, F. (1990). "Equilibria with Communication in a Job Market Example," *Quart. J. Econ.* **105**, 375–398.
- Harsanyi, J. C. (1967–1968). "Games with Incomplete Information Played by Bayesian Players, I, II, III," *Management Sci.* **14**, 159–182, 320–334, 486–502.
- Hart, S. (1985). "Nonzero-Sum Two-Person Repeated Games with Incomplete Information," *Math. Oper. Res.* **10**, 117–153.
- Kohlberg, E. (1975). "Optimal Strategies in Repeated Games with Incomplete Information," *Int. J. Game Theory* **4**, 7–24.
- Lehrer, E. (1987). "A Note on the Monotonicity of v_n ," *Econ. Lett.* **23**, 341–342.
- Mayberry, J. P. (1967). "Discounted Repeated Games with Incomplete Information," in *Report of the U.S. Arms Control and Disarmament Agency*, ST-116, Chapter V, pp. 435–461.

- Mertens, J.-F. (1995). "The Speed of Convergence in Repeated Games with Incomplete Information on One Side," *CORE D.P.*, 9506.
- Mertens, J.-F., and Zamir, S. (1971–1972). "The Value of Two-Person Zero-Sum Repeated Games with Lack of Information on Both Sides," *Int. J. Game Theory* **1**, 39–64.
- Mertens, J.-F., Sorin, S., and Zamir, S. (1994). "Repeated Games," *CORE D.P.*, 9420, 9421, 9422.
- Sorin, S. (1979). "Une introduction aux jeux à deux joueurs, à somme nulle et à information incomplète," *Cahiers du Groupe de Mathématiques Economiques*, Vol. 1, Paris. (English version: (1980) IMSSS, Stanford, TR 312)
- Sorin, S., and Zamir, S. (1985). "A 2-Person Game with Lack of Information on 1 1/2 Sides," *Math. Oper. Res.* **10**, 17–23.

*Sylvain Sorin**

*Laboratoire d'Econometrie
Ecole Polytechnique
1 rue Descartes
75005 Paris
France*

*E-mail: sorin@poly.polytechnique.fr.