

STOCHASTIC AND REPEATED GAMES

Françoise Forges

Universite Paris – Dauphine, , France.

Keywords: Bayesian game, cooperation, discounted game, finitely repeated game, infinitely repeated game, information, Nash equilibrium, non-zero-sum game, stochastic game, strategic game, strategy, supergame, zero-sum game, value

Contents

1. Introduction
 2. Supergames
 - 2.1 Standard Signals
 - 2.1.1 Finitely Repeated Games
 - 2.1.2 Infinitely Repeated Games
 - 2.1.3 Discounted Games
 - 2.1.4 Variants of the Model
 - 2.2 Imperfect Monitoring
 3. Repeated Games with Incomplete Information
 - 3.1 Bayesian Games
 - 3.2 Zero-Sum Games
 - 3.2.1 Lack of Information on One Side
 - 3.2.2 Incomplete Information on Both Sides
 - 3.2.3 Symmetric Information
 - 3.3 Non-Zero-Sum Games
 - 3.3.1 Lack of Information on One Side
 - 3.3.2 Other Models
 4. Stochastic Games
 - 4.1 Zero-Sum Games
 - 4.2 Non-Zero-Sum Games
- Acknowledgments
Glossary
Bibliography
Biographical Sketches

Summary

This article presents the basic models of discrete dynamic multiperson decision problems: repeated games with complete information, also called "supergames", repeated games with incomplete information and stochastic games. When relevant, zero-sum games are first investigated. The reference solution concept is the Nash equilibrium, hence the value in the zero-sum case. Finitely and infinitely repeated games are studied. The latter ones can be solved in an asymptotic way, by considering the limit, as the number of stages goes to infinity (respectively, as the discount factor goes to one) of the solutions of the finitely (respectively, discounted) repeated games. Another approach consists of solving directly the infinite horizon game by defining a limit payoff function in this game. The various methods are illustrated on classical results and examples.

1. Introduction

Most phenomena evolve over time and depend on random events. When they are governed by nature alone, they can be modeled as stochastic processes, as for instance Markov chains. When they are determined by nature and a single decision maker, they can be studied by dynamic optimization, in particular, dynamic programming. When the decisions of conflicting individuals matter, they become the topic of game theory. As soon as several agents interact, many new issues arise, in addition to the ones that are already faced by a single agent in front of nature. For instance, the information of every individual on nature's moves and on the other agents' actions becomes a key ingredient of the model. Given the possible complexity of dynamic games, one is led to identify particular frameworks in which only some of the difficulties are present.

The simplest model of a dynamic game consists of playing the same strategic form game for a very large number of stages, while observing all players' actions after every stage. One then refers to repeated games "with complete information and perfect monitoring" or simply, to "supergames". The main result in this framework, which is known as the "Folk theorem", characterizes the set of Nash equilibrium payoffs of the repeated game as the set of jointly feasible and individually rational outcomes of the one-shot game. The usual interpretation is that the repetition of a game favors cooperation between the players. The basic idea behind this result is that the observation of past moves makes it possible to punish any player who deviates from a prescribed sequence of actions.

The previous model can be made more realistic by relaxing the assumption of perfect observation of the past moves, which gives rise to games with "imperfect monitoring". In such a framework, deviations may not be detectable, and even when they are, the implementation of punishments can be problematic. The systematic study of this natural extension of the model started recently and is far from being complete.

Besides facilitating cooperative behavior, the repetition of a game may also enable the players to communicate: if the game lasts for a sufficiently long time, some stages can just be used to exchange information. Communication is particularly desirable in models where some individuals have more information than others do. Aumann and Maschler introduced such models in the sixties, in a series of reports to the United States Arms Control and Disarmament Agency. They mostly studied zero-sum two-person repeated games in which the players are uncertain about their stage payoffs. These are chosen once and for all by a move of nature at the very beginning of the game. The players receive some private information on nature's choice; during the course of the game, different forms of imperfect monitoring are possible. Obviously, in a zero-sum game, one does not expect any cooperation between the players. Hence, the analysis concentrates on the communication aspects, which can already be subtle in such games. This study turned out to be a necessary step toward the characterization of Nash equilibrium payoffs in non-zero-sum games. Indeed, Folk theorem-like results rely on the players' individually rational levels, which are determined in zero-sum games.

Repeated games with incomplete information are intimately related to another model of dynamic games that historically came before them: the stochastic games, which were introduced by Shapley in the fifties. They can be viewed as an extension of dynamic

programming to multiperson decision problems. In these games, a new state of nature, which determines the payoffs, is chosen at every stage, as a function of the state and the players' moves at the previous stage. The players are completely informed of the past history and the current state. The model was extensively studied in the zero-sum case. The existence of a value was established under very general mathematical assumptions on the state and action spaces. Substantial progress was made recently in proving the existence of Nash equilibrium payoffs in the non-zero-sum case.

In the sequel of the article, we survey classical results on supergames, repeated games with incomplete information and stochastic games. This order of presentation allows us to illustrate the basic methods of solution on simple examples.

2. Supergames

2.3 Standard Signals

2.3.1 Finitely Repeated Games

Let G be a finite strategic form game, namely, a (finite) set of players I , a (finite) set of actions A^i for every player i and payoff (or, better, utility) functions $u^i : A \rightarrow \mathbb{R}$, where $A = \prod_{i \in I} A^i$. Let n be a positive integer. In the n time repeated game, denoted as Γ_n , at every stage $t = 1, \dots, n$, every player i chooses a move $a_t^i \in A^i$; the vector of actions $a_t = (a_t^i)_{i \in I}$ is then revealed to all players. The next stage is played in a similar way. At the end, player i receives the payoff $\frac{1}{n} \sum_{t=1}^n u^i(a_t)$. As part of the "complete information" assumption, the description of the game is assumed to be common knowledge among the players. It is also understood that Γ_n is a game with perfect recall, so that (by Kuhn's theorem), one can focus on behavior strategies without loss of generality. Let H_t be the set of histories up to stage t : $H_t = A^{t-1}$; for any set E , let $\Delta(E)$ denote the set of all probability distributions over E ; a (behavior) strategy of player i in Γ_n is a sequence of mappings $\sigma_t^i : H_t \rightarrow \Delta(A^i)$, $t = 1, \dots, n$. Let $\sigma^i = (\sigma_t^i)_{t \geq 1}$ and $\sigma = (\sigma^i)_{i \in I}$. A vector of strategies σ induces a probability distribution over histories (with expectation E_σ) and hence over payoffs. Nash equilibria of Γ_n are defined in the standard way. Let \mathcal{E}_n be the set of Nash equilibrium payoffs of Γ_n .

2.3.2 Infinitely Repeated Games

Long repeated games are particularly interesting. They can be solved by means of the limit (with respect to the Hausdorff topology), as $n \rightarrow \infty$, of the \mathcal{E}_n or by introducing the infinitely repeated game Γ_∞ , which, at every stage t , is played as Γ_n but lasts forever. As suggested by the example below, the latter model may be more appropriate to capture the behavior of individuals in long games of unknown duration. In Γ_∞ , a strategy σ^i of player i is an infinite sequence, with σ_t^i as above. Given a strategy vector σ in Γ_∞ , let us denote as $\gamma_n^i(\sigma)$ the expected n stage payoff of player i , i.e., $\gamma_n^i(\sigma) = E_\sigma[\frac{1}{n} \sum_{t=1}^n u^i(a_t)]$. Player i 's payoff in Γ_∞ can be defined as some (Banach) limit (or the "lim inf" or the "lim sup"),

as $n \rightarrow \infty$, of $\gamma_n^i(\sigma)$. Some care must indeed be taken, since the payoffs associated with arbitrary strategies are not necessarily convergent. A priori, the set of equilibrium payoffs depends on the limit that is chosen. Alternatively, σ is a uniform equilibrium in Γ_∞ if the $\gamma_n^i(\sigma)$ converge and for every $\varepsilon > 0$, there exists n_0 such that for every $n \geq n_0$, σ is an ε -equilibrium of Γ_n , i.e., for all i and all strategies τ^i , $\gamma_n^i(\sigma) \geq \gamma_n^i(\tau^i, \sigma^{-i}) - \varepsilon$ (where (τ^i, σ^{-i}) is obtained by replacing σ^i by τ^i in σ). The different approaches generate the same set \mathcal{E}_∞ of equilibrium payoffs in Γ_∞ .

The basic result on repeated games, the "Folk theorem", characterizes \mathcal{E}_∞ as the set V of all feasible, individually rational payoffs of the one-shot game G . More precisely, the set of feasible vector payoffs in G is the convex hull of the $(u^i(a))_{i \in I}$, $a \in A$. A payoff is thus feasible in G if it can be achieved by means of a correlated strategy. In order to define player i 's minimax level in G , recall that $\Delta(A^j)$ is the set of mixed strategies of player j in G , and let $\Delta^{-i} = \prod_{j \neq i} \Delta(A^j)$. Let $v^i = \min_{\alpha^{-i} \in \Delta^{-i}} \max_{\alpha^i \in \Delta(A^i)} u^i(\alpha^i, \alpha^{-i})$ (where u^i also denotes the extension of u^i to mixed strategies). A vector payoff z , in \mathbb{R}^I , is individually rational if $z^i \geq v^i$, for every $i \in I$. Obviously, an equilibrium payoff of Γ_∞ must belong to V . To establish the converse, fix a payoff in V . Since it is feasible, it can be achieved by means of an infinite sequence of moves, say $(a_t)_{t \geq 1}$. In order to generate this sequence in equilibrium, the opponents of a possible defector make the threat of punishing him (for a sufficiently large number of stages) at his minimax level. Since the payoff is individually rational, no deviation can be profitable. The argument is illustrated on the next example.

Example 1

The following two-person game is known as the "prisoner's dilemma":

$$\begin{pmatrix} 2,2 & 0,3 \\ 3,0 & 1,1 \end{pmatrix}$$

In the one-shot game G , both players have a dominant strategy; $\mathcal{E}_1 = \{(1, 1)\}$. For every n , $\mathcal{E}_n = \{(1, 1)\}$ as well; a rigorous proof of this is tedious, but the argument is straightforward for subgame perfect equilibrium payoffs. Thus, the n stage repeated game does not reflect the intuition that repetition favors cooperation. By the Folk theorem, \mathcal{E}_∞ is the set of all payoffs, with both components greater than 1, which can be written as a convex combination of the four payoffs in the above matrix. In particular, the cooperative payoff $(2, 2) \in \mathcal{E}_\infty$. To get this payoff, each player plays his first ("cooperative") action as long as the other player does. If a player observes a deviation, he switches to his other action ("non-cooperative"), so that the deviator cannot expect more than 1. It is sufficient to apply the "punishment" for a finite number of stages to prevent any deviation. (For further account of this example, see *Experimental Game Theory*).

In the previous example, \mathcal{E}_n does not converge (with respect to the Hausdorff topology)

to \mathcal{E}_∞ . However, if for every player i , G (or some Γ_{n_i}) has an equilibrium payoff $e(i)$ that is strictly individually rational for i (i.e., such that $e^i(i) > v^i$), the Folk theorem characterization also applies to $\lim_{n \rightarrow \infty} \mathcal{E}_n$.

Furthermore, an analog to the Folk theorem holds for the set \mathcal{E}'_∞ of subgame perfect equilibrium payoffs of Γ_∞ (by the same argument as above). Under suitable assumptions on the one-shot game G , the convergence result also extends to the sets \mathcal{E}'_n of subgame perfect equilibrium payoffs in Γ_n .

2.3.3 Discounted Games

In the previous analysis, it is understood that players are infinitely patient. This assumption is useful in that it provides a benchmark: the Folk theorem characterization. One takes account of the possible impatience of the players by introducing a discount factor in the payoff function. Another interpretation is that the game can stop, with some positive probability, at every stage. Let the stages $n = 1, 2, \dots$, and the strategies $\sigma^i, i \in I$, be described as above. Let $\lambda \in (0, 1)$; in the discounted game Γ_λ , player i evaluates his expected payoff from σ as $E_\sigma[\sum_{t=1}^{\infty} \lambda(1-\lambda)^{t-1} u^i(a_t)]$. i.e. player i uses a geometric average. λ is best interpreted as the probability that the game ends; the discount factor is $1 - \lambda$. We make the simplifying assumption that all players use the same discount factor throughout the article. The previous payoff function makes Γ_λ a well-defined game, exactly as Γ_n . Let \mathcal{E}_λ (respectively, \mathcal{E}'_λ) be the set of equilibrium (respectively, subgame perfect equilibrium) payoffs in Γ_λ . As $\lambda \rightarrow 0$, \mathcal{E}'_λ converges to V , provided that the latter set has a non-empty interior. The proof uses the same basic ideas as above, but punishments are more delicate.

-
-
-

TO ACCESS ALL THE 20 PAGES OF THIS CHAPTER,
Visit: <http://www.eolss.net/Eolss-sampleAllChapter.aspx>

Bibliography

Aumann R. (1985). Survey of Repeated Games. *Issues in Contemporary Microeconomics and Welfare* (ed. G. Feiwel), 209-242. Macmillan. [The survey discusses a number of results and suggests new themes of research in the area, e.g., dealing with bounded rationality].

Aumann R. and Hart S., eds. (1992). *Handbook of Game Theory with Economic Applications*, Volume I. North-Holland: Elsevier Science Publishers. [Three chapters of this book are devoted to repeated games: Chapter 4, "Repeated Games with Complete Information", by S. Sorin (pp. 69-107), Chapter 5, "Repeated Games of Incomplete Information: Zero-Sum", by S. Zamir (pp. 110-154), and Chapter 6, "Repeated Games of Incomplete Information : Non-Zero-Sum", by F. Forges (pp. 155-177). They contain a description of the state of the art at the time, rigorous statements and hints of the proofs.]

Aumann R. and Hart S., eds. (forthcoming). *Handbook of Game Theory with Economic Applications*, Volume III. North-Holland: Elsevier Science Publishers. [This volume contains a chapter on "Stochastic Games", by J.-F. Mertens and N. Vieille. Same comments as for the chapters on repeated games.]

Aumann R. and Maschler M., with the collaboration of Stearns R. (1995). *Repeated Games with Incomplete Information*, 342 pp. Cambridge, Mass: MIT Press. [The book collects the seminal papers on repeated games with incomplete information, which were originally done under contract to the United States Arms Control and Disarmament Agency; they contain the basic results and enlightening examples; they also mention still unsolved problems. Postscripts have been added to the articles; the bibliography has been updated.]

Fudenberg D. and Tirole J. (1991). *Game Theory*, 579 pp. Cambridge, Mass: MIT Press. [This textbook contains a detailed chapter on repeated games with complete information, possibly with imperfect monitoring. The authors discuss a number of interesting variants of the model which could not be covered in details in the present article (among others, varying opponents, reputation effects, see subsection 2.1.4). The book also contains some sections on stochastic games.]

Ichiishi T., Neyman A. and Tauman Y., eds. (1990). *Game Theory and Applications*, 435 pp. San Diego: Academic Press. [The book collects papers that were presented at the International Conference on Game Theory and Applications held at the Ohio State University in 1987. It contains four surveys related with the present article: "Supergames", by S. Sorin (pp. 465-63), "Repeated games with incomplete information" by F. Forges (pp. 64-76), "Repeated Games", by J.-F. Mertens (pp. 77- 130, reprinted from *Proceedings of the International Congress of Mathematicians*, Berkeley, 1986, pp. 1528-1577), "Bounded Rationality and Strategic Complexity in Repeated Games" by E. Kalai (pp. 131-157).]

Mertens J.-F. (1982). Repeated Games: an Overview of the Zero-Sum Case. *Advances in Economic Theory* (ed. W. Hildenbrand), 175-182. Cambridge University Press. [This survey connects repeated games and stochastic games.]

Mertens J.-F., Sorin S. and Zamir S. (1994). *Repeated Games*. CORE Discussion Paper 9420, 9421, 9422. [This forthcoming advanced textbook presents fundamental game theoretical tools in a highly general mathematical framework in order to state the finest possible results on repeated and stochastic games.]

Neyman A. and Sorin S., eds. (1999). *Stochastic Games, Proceedings of the NATO Advanced Study Institute, Stony Brook*. [The volume contains classical as well as recent results on stochastic games].

Osborne M. and Rubinstein A. (1994). *A Course on Game Theory*, 352 pp. Cambridge, Mass: MIT Press. [A chapter of this textbook covers repeated games with complete information and perfect monitoring. Various models, e.g. generated by different payoff evaluation criteria, are carefully analyzed. Another chapter deals with complexity considerations in repeated games.]

Raghavan T.E.S., T Ferguson., T. Parthasarathy and Vrieze O.J., eds. (1991). *Stochastic Games and Related Topics: Essays in Honor of Lloyd Shapley*, 233 pp. Boston: Kluwer Academic Publishers. [This volume contains specialized articles on stochastic games]

Biographical Sketch

Françoise Forges is Professor of Economics at Université Paris – Dauphine, France. She received master's and doctoral degrees in Mathematics from Université Catholique de Louvain, Belgium, and "habilitation à diriger des recherches" in Applied Mathematics from Université de Paris I, France. Her research is mostly in Game Theory; she developed the concept of communication equilibrium in non-cooperative, possibly repeated, games with incomplete information and studied the strategic foundations of general equilibrium. Subsequently, she was interested in cooperative games with asymmetric information. Dr. Forges published articles in scientific journals such as *Econometrica*, *Economic Theory*, *Games and Economic Behavior*, *International Journal of Game Theory*, *Journal of Economic Theory* and *Mathematics of Operations Research*. She is on the editorial board of *Economic Theory*, *Games and Economic Behavior* and *Mathematics of Operations Research*. Dr. Forges is a member of the *Institut Universitaire de France*; and is a fellow of the *Econometric Society* and a member of the executive committee of the *Game Theory Society*.