Available online at www.sciencedirect.com

ScienceDirect

Journal of Economic Theory 159 (2015) 70–104

JOURNAL OF Economic Theory

www.elsevier.com/locate/jet

Bayesian repeated games and reputation ☆

Antoine Salomon [a], Françoise Forges [b,*]

[a] *PSL, Université Paris-Dauphine, LEDa, France*
[b] *PSL, Université Paris-Dauphine, CEREMADE and LEDa, France*

Received 12 February 2014; final version received 22 May 2015; accepted 25 May 2015

Available online 28 May 2015

Abstract

We consider two-person undiscounted and discounted infinitely repeated games in which every player privately knows his own payoffs (private values). Under a further assumption (existence of uniform punishment strategies), the Nash equilibria of the Bayesian infinitely repeated game without discounting are payoff-equivalent to tractable, completely revealing, equilibria. This characterization does not apply to discounted games with sufficiently patient players. We show that in a class of public good games, the set of Nash equilibrium payoffs of the undiscounted game can be empty, while limit (perfect Bayesian) Nash equilibrium payoffs of the discounted game, as players become increasingly patient, do exist. These equilibria share some features with the ones of two-sided reputation models.

JEL classification: C73; C72; C71; D82; H41

Keywords: Discounting; Incomplete information; Public good; Repeated game; Reputation

1. Introduction

We consider two-person infinitely repeated games with the following features: the players' types are selected once and for all at the beginning of the game, the players privately know

☆ Conversations with Gorkem Celik, Martin Cripps, Gaël Giraud, Johannes Hörner, Vincent Iehlé, Frédéric Koessler, Bernard Lebrun, Marcin Peski and Péter Vida are gratefully acknowledged. We also thank the Associate Editor and three anonymous referees for useful suggestions. F. Forges acknowledges support from Institut Universitaire de France.

* Corresponding author.
E-mail addresses: antoine.salomon@dauphine.fr (A. Salomon), francoise.forges@gmail.com (F. Forges).

http://dx.doi.org/10.1016/j.jet.2015.05.014
0022-0531/

their own payoffs (private values) and they observe each other's decisions at every stage of the game (full monitoring). Furthermore, in the one-shot Bayesian game, every player has a *uniform punishment strategy* against the other one. This property means that there is no need to know a player's type to punish him in the harshest possible way, i.e., at his ex post individually rational level.

Our main concern is whether there exist Nash equilibrium payoffs[1] that can be achieved, at least approximately, in the discounted infinitely repeated game, independently of the discount factor, provided that the players are patient enough. A natural way to address this question is to investigate the set $\mathcal{N}[\Gamma_\infty]$ of Nash equilibrium payoffs of the undiscounted repeated game Γ_∞, in which payoffs are evaluated by the limit of means criterion. Indeed, if $\mathcal{N}[\Gamma_\infty]$ is nonempty, one can hope that payoffs in $\mathcal{N}[\Gamma_\infty]$ become approximately achievable in the δ-discounted infinitely repeated game Γ_δ when δ is sufficiently close to 1. This approach has been used successfully to prove the folk theorem in games with complete information.

Koren (1992) established that, when players know their own payoffs, every Nash equilibrium of Γ_∞ is payoff-equivalent to a completely revealing equilibrium of Γ_∞. We show that, under the additional assumption that uniform punishment strategies are available, the set $\mathcal{N}[\Gamma_\infty]$ can be characterized in an even more tractable way, as a set of payoffs that are feasible, incentive compatible and individually rational in the *one-shot game*.

We apply the previous characterization to a class of familiar public good games, in which the players can each be "normal" or "greedy" (according to some prior probability distribution p) and have two actions: "contribute" or "do not contribute." The public good yields the normalized payoff 1 and is produced from the players' private good endowment if at least one of them contributes. A normal (resp., greedy) player gets the payoff $\omega < 1$ (resp., $z > 2$) from his private good endowment. We denote the game as $G(p, \omega, z)$. When both players are normal, it is a game with strictly conflicting interests (see Cripps et al., 2005 and Atakan and Ekmekci, 2013).

We find that, for a significant set of parameters p, ω, z, the undiscounted infinitely repeated public good game $G_\infty(p, \omega, z)$ has no equilibrium at all, i.e., $\mathcal{N}\left[G_\infty(p, \omega, z)\right] = \emptyset$. The nonexistence phenomenon, which was already illustrated in Koren (1992) on a particular example, appears to be robust.[2]

By contrast, for every given discount factor δ, Nash's theorem guarantees the existence of an equilibrium in every δ-discounted infinitely repeated game Γ_δ. The set $\mathcal{N}[\Gamma_\delta]$ of associated equilibrium payoffs typically depends on δ. To fix the ideas, imagine that $\mathcal{N}[\Gamma_\delta]$ contains a unique payoff $x(\delta)$. Then the question raised above reduces to the convergence of $x(\delta)$ when δ goes to 1. In general, when $\mathcal{N}[\Gamma_\delta]$ is possibly multivalued, the problem can be formalized as the nonemptiness of the set $\liminf_{\delta\to 1} \mathcal{N}[\Gamma_\delta]$. Our main result is that, in the public good game under consideration, $\liminf_{\delta\to 1} \mathcal{N}\left[G_\delta(p, \omega, z)\right]$ is not empty for every (p, ω, z), i.e., even when $\mathcal{N}\left[G_\infty(p, \omega, z)\right] = \emptyset$.

The proof of the result is constructive. For every discount factor δ, we identify a family of equilibria in specific strategies, in which at the first stages of the game, greedy players do not contribute and normal players behave as in a *war of attrition*, namely, contribute with a low probability in the hope that the other player will give up. We prove that the payoffs associated with each of these war of attrition equilibria converge as the discount factor δ goes to 1.

[1] Throughout the paper, equilibrium payoffs are computed conditionally on the players' private information, namely, at the interim stage.

[2] We show in particular that the emptiness phenomenon survives if the set of feasible individually rational payoffs has a nonempty interior under complete information, for every pair of types.