



Bayesian Repeated Games

Francoise Forges, Antoine Salomon

► To cite this version:

| Francoise Forges, Antoine Salomon. Bayesian Repeated Games. 2013. hal-00803919v2

HAL Id: hal-00803919

<https://hal.science/hal-00803919v2>

Preprint submitted on 27 Mar 2013 (v2), last revised 21 Feb 2014 (v5)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Bayesian repeated games*

Françoise Forges[†] and Antoine Salomon[‡]

Preliminary (March 2013)

Please do not circulate without the authors' permission

Abstract

We consider Bayesian games, with independent private values, in which uniform punishment strategies are available. We establish that the Nash equilibria of the Bayesian infinitely repeated game without discounting are payoff equivalent to tractable separating (i.e., completely revealing) equilibria and can be achieved as interim cooperative solutions of the initial Bayesian game. We also show, on a public good example, that the set of Nash equilibrium payoffs of the undiscounted game can be empty, while limit Nash equilibrium payoffs of the discounted game, as players become infinitely patient, do exist.

Keywords: Bayesian game, incentive compatibility, independent private values, individual rationality, infinitely repeated game, public good.

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

*Some of the results of this paper were first presented by Françoise Forges at PET 12 (Taipei, June 13, 2012). Conversations with Gorkem Celik, Gaël Giraud, Vincent Iehlé, Frédéric Koessler and Péter Vida are gratefully acknowledged. Bernard Lebrun made a number of precise and helpful comments on an earlier draft.

[†]CEREMADE and LEDa, Université Paris-Dauphine. Member of Institut Universitaire de France. E-mail: francoise.forges@gmail.com

[‡]LEM, Université Paris 2. E-mail: salomona@imagine.enpc.fr

1 Introduction

We consider Bayesian games with independent private values. We make the further assumption that “uniform punishment strategies” are available. We first characterize the Nash equilibrium payoffs of the Bayesian infinitely repeated game without discounting as tractable, separating (i.e., completely revealing) equilibrium payoffs. Using this characterization, we show, on a public good example, that the latter set of Nash equilibrium payoffs can be empty. We then establish a partial version of the folk theorem: the previous equilibrium payoffs can be achieved as interim cooperative solutions of the initial (one-shot) Bayesian game. We finally turn to the set of Nash equilibrium payoffs of the discounted game. In this case, we show, that, in the public good example, a limit Nash equilibrium payoff of the discounted game, as players become infinitely patient, exists and differs significantly from previous equilibria identified in the recent literature (e.g., in Peski (2012))

Relationships with the previous literature

Aumann and Maschler started to study the Nash equilibrium payoffs of infinitely repeated games with incomplete information in the mid-sixties (see Aumann and Maschler (1995)). At the very same time, Harsanyi (1967) proposed the formal definition of games with incomplete information as Bayesian games. Building on the work of Aumann, Maschler and Stearns (1968), S. Hart (1985) characterized the set of Nash equilibrium payoffs of any two-person (undiscounted) infinitely repeated game in which only one of the players has private information. This looks like an extremely particular class of games but the characterization is already quite intricate: it involves a description of the dynamic process followed by the equilibrium and so, does not give much hope to be related to solutions of the one-shot game.

As S. Hart (1985), Koren (1992) considers two-person games but, instead of assuming that only one player is privately informed, he assumes that every player “knows his own payoff”. According to a more usual terminology in microeconomics, he makes the assumption that “values are private and independent”. In this case, he shows that the Nash equilibrium of the (undiscounted) infinitely repeated game can be characterized in an elegant way: they are payoff equivalent to *completely revealing* (also called “separating”) equilibria.¹ Once such a characterization is available, one can ask whether

¹See Shalev (1994) for a similar characterization of Nash equilibrium payoffs in Hart (1985)’s model with “known own payoffs” and Forges (1992) for a survey of results on

it confirms that repeating a game has the same effect as commitment in the one-shot game.

The assumption of independent, private values is satisfied in many economic applications, *e.g.*, in some public good games (see Palfrey and Rosenthal (1994) and Fudenberg and Tirole (1991, example 6.1, p. 211). These games typically involve more than two players but satisfy a further assumption, which we call “uniform punishments”. We show (in proposition 1) that in these public good games, and more generally, in any n –person Bayesian game with independent, private values and uniform punishments, the Nash equilibrium of the undiscounted infinitely repeated game are all payoff equivalent to *completely revealing* equilibria. Furthermore, thanks to uniform punishments, our characterization is more tractable than Koren (1992)’s one and immediately goes through in the case of n players.²

Our tractable characterization facilitates the comparison with the cooperative solutions of the initial Bayesian game. Furthermore, it tells us how incentives to reveal private information can differ in the short and the long run. In a finitely repeated game, players may benefit from hiding their type, *e.g.*, their willingness to contribute to a public good, at an early stage of the game (see Fudenberg and Tirole (1991), example 8.3, p. 333). Proposition 1 tells us that, in an undiscounted infinitely repeated game, players cannot benefit from concealing their private information.

Equipped with the characterization of proposition 1, we show in proposition 2 that, for any n –person Bayesian game with independent, private values and uniform punishments, the set of Nash equilibrium payoffs of the undiscounted infinitely repeated game is contained in the set of *interim* cooperative solutions of the Bayesian game, as defined in Myerson (1991) and Forges (2013). In other words, the repetition of the game enables the players to cooperate, as in the folk theorem with complete information. However, the previous inclusion can be strict. More surprisingly, unlike the set of *interim* cooperative solutions of the one-shot game, the set of Nash equilibria of the infinitely repeated game can be empty. This is illustrated on a public good game (example 1).

The latter finding, which tells us that, without discounting, the folk theorem does not hold for Bayesian games in which several players have private

non-zero sum infinitely repeated games with incomplete information.

²If uniform punishments are available, there is no need to appeal to Blackwell (1956)’s approachability theorem. By extending the latter result to n players, as, in *e.g.*, in Hörner, Lovo and Tomala (2011), one should be able to prove Koren’s result for n players.

information, must be contrasted with the results on “reputation effects”. There is an extensive literature on this topic. Most papers concentrate on two-person games with a single informed player, who tries to establish a reputation, but allow for discounted payoffs (see Mailath and Samuelson (2006); Sorin (1999) gives a synthetic presentation of various related models, including infinitely repeated games with known own payoffs; as a sample of references, let us mention Kreps *et al.* (1982), Fudenberg and Maskin (1986), Schmidt (1993), Cripps and Thomas (1995, 1997, 2003), Cripps *et al.* (1996), Israeli (1999), Chan (2000), Cripps *et al.* (2005), Atakan and Ekmekci (2012, 2013)).³ An important difference between the models designed to study reputation effects and the one that we consider in these notes is that, rather than perturbing a Bayesian game with complete information, we start with given sets of types for every player and arbitrary beliefs over these types.

Our characterization shows that, under incomplete information, the cooperative solutions of the one-shot game and the non-cooperative solutions of the undiscounted repeated game mostly differ in the individual rationality levels of the players. Under the assumptions of independent private values and uniform punishments, the *ex post* individual rationality level of a player, namely the level at which the other players can punish him when they know his type, is relevant in the infinitely repeated game. *Interim* individually rational payoffs in the sense of Myerson (1991) are always *ex post* individually rational. When there exist uniform punishment strategies, the reverse also holds: this is the key of proposition 2. However, if the assumption of uniform punishments is relaxed, individual rationality in the infinitely repeated game relies on Blackwell (1956)’s approachability strategies. As a consequence, proposition 2 is no longer true, while Koren’s characterization still holds, at least in the two-person case. Section 6 discusses our underlying assumptions in details.

Until section 5, as in Hart (1985) and Koren (1992), we assume that payoffs in the infinitely repeated games are evaluated as Banach limits of the expected average payoffs, namely, without discounting. The main justification for proceeding in this way is well-known: we are looking for robust results, which do not depend on the precise discount factor of the players. Even more, we are interested in robust equilibrium strategies, which can be

³Chan (2000) devotes a section to two-sided reputation effects, while Atakan and Ekmekci (2013) deal exclusively with this model. Our paper shares some features with the latter article (see section 5).

used for a reasonable range of discount factors. Hence, it is natural to start with the study of the equilibrium payoffs of the undiscounted infinitely repeated game, in order to “guess” how the robust equilibrium payoffs look like. This is how game theorists derived the folk theorem for infinitely repeated games with complete information.

While the study of the undiscounted infinitely repeated game seems a mandatory first step, a careful analysis of the consequences of this assumption is in order (see Bergin (1989) for an early reference). As we illustrate in section 5, the phenomena described by proposition 1 (all equilibria are payoff equivalent to completely revealing ones) and example 1 (the slightest doubt on the players’ types can lead to non existence of equilibrium) are not to be expected in discounted games, even with patient players.

As we already noticed, many papers devoted to reputation effects consider repeated games with discounting but these papers impose restrictions on the players’ beliefs. Before that, in their seminal papers, Aumann and Maschler already showed that undiscounted zero-sum infinitely games in which both players are uninformed could fail to have a value. However, Mertens and Zamir (1971) proved that the value of the discounted game, as players become more and more patient, always converges. To the best of our knowledge, there is hardly any analog of this result in the non-zero-sum case. Peski (2008) and Peski (2012) study discounted repeated games with specific forms of incomplete information. Peski (2008) characterizes the limit of the sets of Nash equilibrium payoffs of two-person discounted repeated games with lack of information on one side and known own payoffs, when the informed player has two types (namely, the discounted version of Shalev (1994) in the case of two types). Peski (2008)’s characterization results show that more payoffs can be achieved in the limit discounted case than in the undiscounted case but still finitely revealing equilibria are the only ones that need to be considered. Peski (2012) extends these results to the class of discounted repeated games with known-own payoffs which satisfy an “open thread assumption”, namely, at least in the two-person case, in which there exist belief-free equilibria in the sense of Hörner and Lovo (2009) (see also Hörner, Lovo and Tomala (2011)).⁴ There are no belief free equilibria in the public good example of the current paper. Nonetheless, in every repeated version of this game with

⁴As already in Peski (2008), existence and non emptiness of limit sets of Nash equilibrium payoffs as players become increasingly patient is not an issue in Peski (2012), at least in the two player case, because belief free equilibria are assumed to exist. The main result is a full characterization of the limit set, in terms of finitely revealing equilibria.

a sufficiently high discount factor, we construct a Nash equilibrium payoff which converges when players become increasingly patient. We thus show in particular that Peski (2012)’s open thread assumption is not necessary for equilibrium payoffs convergence.

Wiseman (2012) establishes a partial folk theorem in discounted repeated games where the players have the same initial information and get private and public signals along the play. While his model captures in particular “known own payoffs”, as in multisided reputation models (see his example 3), his assumption 1 ensures “gradual public learning” which has no counterpart in infinitely repeated games like the ones considered here. As a consequence, Wiseman (2012)’s folk theorem can be formulated in terms of feasible, *ex post* individually rational payoffs, *without any requirement of incentive compatibility*. By contrast, incentive compatibility is crucial in this paper and in Peski (2008, 2012).

2 Basic Bayesian game

2.1 Definition and main assumptions

Let us fix n players and, for every player i , $i = 1, \dots, n$,

- a finite set of types Θ_i
- a probability distribution q_i over Θ_i
- a finite set of actions A_i , with $|A_i| \geq |\Theta_i|$
- a utility function $u_i : \Theta_i \times A \rightarrow \mathbb{R}$, where $A = \prod_{1 \leq i \leq n} A_i$.

This defines a (one-shot) Bayesian game with independent, private values, which we denote as $\Gamma(q)$, with $q = (q_i)_{1 \leq i \leq n}$.⁵ Without loss of generality, we assume that $q_i(\theta_i) > 0$ for every $\theta_i \in \Theta_i$. The interpretation is that types θ_i , $i = 1, \dots, n$, are first chosen in Θ , independently of each other, according to q . At the *interim* stage, player i is only informed of his own type θ_i . The players then choose simultaneously an action.

⁵We only recall the parameter q in the notation $\Gamma(q)$ for the Bayesian game, because it will often happen, *e.g.*, in the examples, that the beliefs q vary while all other parameters are fixed.

For any finite set E , let us denote as $\Delta(E)$ the set of probability distributions over E . A mixed strategy⁶ of player i in $\Gamma(q)$ is a mapping from Θ_i to $\Delta(A_i)$. Similarly, a correlated strategy for players $j \neq i$ is a mapping from $\Theta_{-i} = \prod_{j \neq i} \Theta_j$ to $\Delta(A_{-i})$, where $A_{-i} = \prod_{j \neq i} A_j$. We keep the notation u_i for the (multi)linear extension of utility functions over mixed and/or correlated strategies. Hence we write, for every $i = 1, \dots, n$, $\theta_i \in \Theta_i$, $\pi \in \Delta(A)$,

$$u_i(\theta_i, \pi) = \sum_a \pi(a) u_i(\theta_i, a)$$

In particular, for every $i = 1, \dots, n$, $\theta_i \in \Theta_i$, $\sigma_i \in \Delta(A_i)$, $\tau_{-i} \in \Delta(A_{-i})$,

$$u_i(\theta_i, \sigma_i, \tau_{-i}) = \sum_{a_i, a_{-i}} \sigma_i(a_i) \tau_{-i}(a_{-i}) u_i(\theta_i, a_i, a_{-i})$$

For every player i , $i = 1, \dots, n$, and $\theta_i \in \Theta_i$, let $v_i(\theta_i)$ be the value⁷ of the (complete information, zero-sum) game in which player i maximizes the payoff $u_i(\theta_i, \cdot)$, namely

$$v_i(\theta_i) = \min_{\tau_{-i} \in \Delta(A_{-i})} \max_{\sigma_i \in \Delta(A_i)} u_i(\theta_i, \sigma_i, \tau_{-i}) = \min_{\tau_{-i} \in \Delta(A_{-i})} \max_{a_i \in A_i} u_i(\theta_i, a_i, \tau_{-i}) \quad (1)$$

Observe that, in the previous expression, the probability distribution τ_{-i} achieving the “min” possibly depends on θ_i , which is fixed in the underlying optimization problem. $v_i(\theta_i)$ can thus be interpreted as the *ex post* individual rationality level of player i , namely, the best amount that player i can guarantee to himself if the other players know his type θ_i .

We consider the following assumption (“uniform punishment strategies”):

$$\forall i \exists \tau_{-i} \in \prod_{j \neq i} \Delta(A_j) \text{ s.t. } \forall \theta_i \in \Theta_i \forall a_i \in A_i \quad u_i(\theta_i, a_i, \tau_{-i}) \leq v_i(\theta_i) \quad (2)$$

When (2) holds, τ_{-i} defines independent⁸ punishment strategies which enable players $j \neq i$ to punish player i uniformly, *i.e.*, whatever his type θ_i is, but even more, to keep player i ’s payoff below his *ex post* individual rationality level.

⁶More correctly, “behavior strategy”.

⁷If we allow for correlated mixed strategies, the value exists and can be expressed as a minmax or as a maxmin. We will nevertheless consider independent mixed strategies below.

⁸Independent punishment strategies are important for proposition 1.

Assumption (2) is quite strong but, as illustrated below, it is satisfied in a class of public good games (see, *e.g.*, Palfrey and Rosenthal (1994)).⁹ In these games, the independent private values assumption also holds. Peters and Szentes (2012)'s assumption 1 (p. 397) takes exactly the form of (2) if values are private and independent and mixed strategies are allowed. We will make a more precise comparison in section 4. We will discuss the role of our various assumptions in section 6.

2.2 Application: contribution to a public good

The private information of every player i , $i = 1, \dots, n$, is the value θ_i that he attributes to his endowment of a single unit of the private good. The private endowment values θ_i are chosen independently of each other, according to a probability distribution q_i . Player i has two possible actions a_i : “contribute” (c) and “do not contribute” (d). A public good is produced if, and only if, at least m players contribute. The value of the public good is normalized to 1 for all players. For every $a \in A$, let $M(a)$ be number of contributors, namely

$$M(a) = M((a_i)_{1 \leq i \leq n}) = |\{i : a_i = c\}|$$

The utility function of player i is described by

$$u_i(\theta_i, a_i, a_{-i}) = \begin{cases} 1 & \text{if } a_i = c \text{ and } M(a_i, a_{-i}) \geq m \\ 0 & \text{if } a_i = c \text{ and } M(a_i, a_{-i}) < m \\ 1 + \theta_i & \text{if } a_i = d \text{ and } M(a_i, a_{-i}) \geq m \\ \theta_i & \text{if } a_i = d \text{ and } M(a_i, a_{-i}) < m \end{cases}$$

We refer to the game as $PG(n, m, q)$, $1 \leq m \leq n$. For instance, in $PG(2, 1, q)$, the payoff matrix associated with the pair of types (θ_1, θ_2) is

	c	d
c	$1, 1$	$1, 1 + \theta_2$
d	$1 + \theta_1, 1$	θ_1, θ_2

where we always assume $\theta_i \geq 0$ but can have $\theta_i < 1$ or $\theta_i > 1$. Fudenberg and Tirole (1991, example 6.1, p. 211) propose the following interpretation:

⁹As a slight weakening, $v_i(\theta_i)$ could just be defined as

$$\min_{\tau_{-i} \in \prod_{j \neq i} \Delta(A_j)} \max_{\sigma_i \in \Delta(A_i)} u_i(\theta_i, \sigma_i, \tau_{-i}).$$

player 1 and player 2 belong to a group (say, the members of some university department) and each of them can represent the group at a committee (say, the scientific board of the university). To attend the committee is time consuming and it is enough that one player attends the committee meeting to defend the interests of the group. The whole problem is to decide which one of the players will go to the meeting, given that the value of time for each player is private information.

In $PG(n, m, q)$, a uniform punishment against player i is easily derived: the other players just have to decide not to contribute. More precisely, let $\tau_{-i} = (a_j)_{j \neq i}$ be the $(n - 1)$ -uple of actions in which $a_j = d$ for all players $j \neq i$. Assume first that $m > 1$. Then, by playing d , player i guarantees himself θ_i whatever the other players choose. By playing τ_{-i} , the players $j \neq i$ guarantee that player i 's payoff does not exceed θ_i . Hence, if $m > 1$, $v_i(\theta_i) = \theta_i$ and τ_{-i} is a uniform punishment strategy. Assume now that $m = 1$. Again, by playing d , player i guarantees himself θ_i ; but now, by playing c , player i guarantees himself 1. Hence, by playing according to his type, player i can guarantee himself $\max\{\theta_i, 1\}$. By playing τ_{-i} , the players $j \neq i$ guarantee that player i 's payoff does not exceed $\max\{\theta_i, 1\}$. Hence, if $m = 1$, $v_i(\theta_i) = \max\{\theta_i, 1\}$ and τ_{-i} is a uniform punishment strategy.

3 Undiscounted infinitely repeated Bayesian game

Nash equilibria always exist in the one-shot game $\Gamma(q)$, but fail to reflect the fact that the players may care about the future consequences of their present behavior. In a Nash equilibrium of $\Gamma(q)$, players may reveal a lot of information, choose an individualistic action, etc. Hence we turn to infinitely repeated versions of the previous game, starting with the undiscounted one, which we denote as $\Gamma_\infty(q)$. According to Aumann and Maschler's original model (see Aumann and Maschler (1995)), the players' types are fixed throughout the game. More precisely, $\Gamma_\infty(q)$ is played as follows:

- at a virtual stage (stage -1): the types θ_i , $i = 1, \dots, n$, are chosen in $\Theta = \prod_{1 \leq i \leq n} \Theta_i$ independently of each other, according to q . Player i is only informed of his own type θ_i .
- at every stage t ($t = 0, 1, \dots$): every player i chooses an action in A_i . The

choices are made simultaneously and revealed publicly right after stage t .

Payoffs in $\Gamma_\infty(q)$ are evaluated as (Banach) limits of arithmetic averages (see Hart (1985), Forges (1992)). In section 5, we shall rather consider the discounted version of the infinitely repeated game.

3.1 Characterization of Nash equilibrium payoffs

Let us write $q_{-i}(\theta_{-i})$ for $\prod_{j \neq i} q_j(\theta_j)$. A version of the next characterization of the Nash equilibrium payoffs of $\Gamma_\infty(q)$ was established in Koren (1992).¹⁰

Proposition 1 *Let $\Gamma(q)$ be a Bayesian game with independent private values in which uniform punishment strategies are available. Let $x = (x_i)_{1 \leq i \leq n} = ((x_i(\theta_i))_{\theta_i \in \Theta_i})_{1 \leq i \leq n}$. x is a Nash equilibrium payoff in $\Gamma_\infty(q)$ if and only if there exist $\pi(\theta) \in \Delta(A)$, $\theta \in \Theta$, such that for every $i = 1, \dots, n$, $\theta_i, \theta'_i \in \Theta_i$*

$$\begin{aligned} x_i(\theta_i) &= \sum_{\theta_{-i} \in \Theta_{-i}} q_{-i}(\theta_{-i}) u_i(\theta_i, \pi(\theta_i, \theta_{-i})) \\ &\geq \sum_{\theta_{-i} \in \Theta_{-i}} q_{-i}(\theta_{-i}) \max \{u_i(\theta_i, \pi(\theta'_i, \theta_{-i})), v_i(\theta_i)\} \end{aligned} \quad (3)$$

In the case of *complete information*, namely if the prior probability distribution q is degenerate, proposition 1 reduces to the standard folk theorem: $x = (x_i)_{1 \leq i \leq n} \in \mathbb{R}^n$ is a Nash equilibrium payoff of the infinitely repeated game if and only if x is feasible (*i.e.*, achieved by means of a probability distribution $\pi \in \Delta(A)$) and individually rational (*i.e.*, x_i is larger than player i 's minmax level).

The interpretation of Proposition 1, under *incomplete information*, is that all Nash equilibria of $\Gamma_\infty(q)$ are payoff equivalent to *completely revealing* equilibria, in which

- at stage 0, every player i truthfully reveals his type θ_i

¹⁰Koren (1992) establishes that all Nash equilibrium payoffs of $\Gamma_\infty(q)$ are completely revealing in the case of only *two* players, but without assuming uniform punishments. The latter assumption greatly facilitates the formulation of the equilibrium conditions and the extension to n players.

- at every stage $t \geq 1$, given the reported types $\theta' = (\theta'_i)_{1 \leq i \leq n}$, every player i plays according to $\pi(\theta') \in \Delta(A)$ provided that $\pi(\theta')$ has been followed at every previous stage $1, \dots, t-1$. Otherwise, if player i does not follow $\pi(\theta')$ at some stage $t \geq 1$, players $j \neq i$ punish player i by using uniform independent punishment strategies τ_{-i} holding player i at $v_i(\theta_i)$ at every stage $t+1, t+2, \dots$ whatever his type θ_i and action are.

The nondeviation condition (3) expresses that, assuming that players $j \neq i$ follow the equilibrium strategies, player i of type θ_i can report a type θ'_i possibly different from θ_i . At the end of stage 0, player i learns the true types θ_{-i} of the other players and can then either follow $\pi(\theta'_i, \theta_{-i})$ or not. In the former case, he fully mimics the equilibrium strategy of type θ'_i . In the latter case, he is punished at the level $v_i(\theta_i)$.

Condition (3) is thus both an incentive compatibility condition and an individual rationality condition. Even under our strong assumptions, it is not possible to separate these two aspects of player i 's nondeviating condition. Obviously, for $\theta'_i = \theta_i$, (3) is equivalent to

$$\text{For every } i \text{ and } \theta = (\theta_i, \theta_{-i}) \in \Theta : \quad u_i(\theta_i, \pi(\theta)) \geq v_i(\theta_i) \quad (4)$$

which implies that

$$\text{For every } i \text{ and } \theta_i \in \Theta_i : \quad x_i(\theta_i) \geq v_i(\theta_i) \quad (5)$$

With some abuse of language, we will refer to the latter property as x is *ex post* individually rational and will denote as $EXPIR_i[\Gamma(q)]$ the set of all vector payoffs which satisfy it for player i .

The previous equilibrium conditions are illustrated on examples 0, 1 and 2 below. Examples 0 and 1 belong to the class of public good games introduced in section 2.2. Example 2 is a variant of the battle of the sexes already proposed by Koren (1992) and is actually simpler.

Proposition 1 is established in an appendix.

Example 0:

Recalling section 2.2, let us consider the following two-person game $PG(2, 1, q)$

		$\theta_2 = \omega$		$\theta_2 = z$	
		c	d	c	d
$\theta_1 = \omega$	c	1, 1	1, 1 + ω	1, 1	1, 1 + z
	d	1 + ω , 1	ω , ω	1 + ω , 1	ω , z
$\theta_1 = z$	c	1, 1	1, 1 + ω	1, 1	1, 1 + z
	d	1 + z , 1	z , ω	1 + z , 1	z , z

Each player has two possible types: $\Theta_i = \{\omega, z\}$, $i = 1, 2$. We assume that $0 < \omega < 1$ and $z > 2$: ω represents a “normal” type, who values the public good more than his initial endowment, while z represents a “greedy” type.

Let $\omega = \frac{2}{3}$ and $z = 3$. Consider the following distributions, which yield feasible, ex post individually rational payoffs:

		$\theta_2 = 2/3$ (<i>prob.</i> p_2)		$\theta_2 = 3$ (<i>prob.</i> $1 - p_2$)	
		c	d	c	d
$\theta_1 = 2/3$ (<i>prob.</i> p_1)	c	0	1/2	0	7/10
	d	1/2	0	3/10	0
$\theta_1 = 3$ (<i>prob.</i> $1 - p_1$)	c	0	3/10	0	0
	d	7/10	0	0	1

Conditions (3) show that these distributions induce an equilibrium if and only if $p_1 \leq \frac{3}{5}$ and $p_2 \leq \frac{3}{5}$.

3.2 Existence of Nash equilibrium

Let us denote as $\mathcal{N}[\Gamma_\infty(q)]$ the set of all Nash equilibrium payoffs of $\Gamma_\infty(q)$. Thanks to proposition 1, the set $\mathcal{N}[\Gamma_\infty(q)]$ has a tractable representation so that it is relatively easy to check whether it is empty or not. Koren (1992) already proposes a two-player example in which there is no Nash equilibrium. The next example pertains to the class of public good games introduced in section 2.2.

Example 1: A public good game in which $\mathcal{N}[\Gamma_\infty(q)]$ is empty

Let us consider the game $PG(2, 1, q)$ of example 0 and assume now that the players hold the same beliefs: $q_i = (p, 1 - p)$, $i = 1, 2$, with $0 < p < 1$. We thus refer to the game as $\Gamma_\infty(p)$. Let us set $k = \frac{\omega}{1-\omega}$. We will show that

$$\text{If } z > k + 4 \text{ and } p > \frac{2}{k + 4}, \quad \mathcal{N}[\Gamma_\infty(p)] = \emptyset \quad (6)$$

In other words, if the “greedy” type z is sufficiently high, but has an *arbitrarily small probability* $1 - p$, the infinitely repeated game has *no equilibrium*. For instance, if $\omega = \frac{1}{3}$ and $z > 4.5$, the infinitely repeated game has no equilibrium as soon as the probability of the “greedy” type is smaller than $\frac{5}{9}$.

This finding should be contrasted with the results obtained in standard *reputation* models, in which a very small probability of a “crazy” type is enough to generate interesting equilibrium behavior in the incomplete information game (see Kreps et al. (1982), Fudenberg and Maskin (1986), etc.). Here, if both types are “normal” ($p = 1$), the infinitely repeated game has a plethora of equilibria, but as soon as there is an arbitrarily small doubt that the players could be (very) “greedy”, the game has no equilibrium at all.¹¹

Recalling again section 2.2, the individual levels in $\Gamma(p)$ are $v_i(\omega) = 1$ and $v_i(z) = z$, $i = 1, 2$. According to proposition 1, the equilibrium payoffs of $\Gamma_\infty(p)$ are characterized by four probability distributions $\pi(\theta)$ over $\{c, d\} \times \{c, d\}$, one for every pair of types θ . If $\theta_1 = \theta_2 = z$, *ex post* individual rationality implies that (d, d) must have probability 1. In order to show that $\Gamma_\infty(p)$ has no equilibrium, it is enough to show that $\Gamma_\infty(p)$ has no *symmetric* equilibrium¹². We thus focus on $\pi(\theta)$ ’s of the form:

¹¹Koren (1992) shows that Nash equilibrium payoffs may fail to exist in two-person repeated games in which both players are privately informed. Cripps and Thomas (1995) discuss the consequences of this phenomenon for reputation effects.

¹²If $\Gamma_\infty(p)$ has an equilibrium, there exist probability distributions $\pi(t)$, $t \in T$, over $\{c, d\} \times \{c, d\}$ satisfying (3). If $\pi(t)$, $t \in T$, satisfies (3), the probability distributions $\pi'(t)$, $t \in T$, in which player 1 and player 2 are permuted, also satisfy (3). The same holds for the symmetric distributions $\frac{1}{2}(\pi(t) + \pi'(t))$, $t \in T$, thanks to the linearity of u and the convexity of “max”.

		$\theta_2 = \omega$		$\theta_2 = z$	
		c	d	c	d
$\theta_1 = \omega$	c	γ_c	γ	β_c	ρ
	d	γ	γ_d	α	β_d
$\theta_1 = z$	c	β_c	α	0	0
	d	ρ	β_d	0	1

where all parameters are nonnegative and $2\gamma + \gamma_c + \gamma_d = 1$, $\alpha + \rho + \beta_c + \beta_d = 1$.

The *ex post* individual rationality conditions (4) can be written as

$$\gamma_d \leq k\gamma, \beta_d \leq k\alpha \text{ and } \rho \geq (1 - \frac{1}{z})(1 - \beta_d) \quad (7)$$

In the right hand sight of the equilibrium condition (3) for $\theta_1 = \omega$ and $\theta'_1 = z$,

$$\max \{1 + \rho\omega - (1 - \omega)\beta_d, 1\} = 1 + \rho\omega - (1 - \omega)\beta_d$$

namely, $\beta_d \leq k\rho$ because, from (7), $\beta_d \leq k\alpha$ and $\alpha \leq 1 - \beta_d - \rho \leq \frac{1}{z-1}\rho \leq \rho$. We can thus write (3) for $\theta_1 = \omega$ and $\theta'_1 = z$ as

$$p(k\gamma - \gamma_d) + (1 - p)(k\alpha - \beta_d) \geq p(k\rho - \beta_d) \quad (8)$$

This condition is not compatible with (7) if p is close enough to 1. In order to get some intuition for this, let us try $\gamma_d = \beta_d = 0$, *i.e.*, an *ex post* efficient equilibrium. (7) reduces to $\rho \geq 1 - \frac{1}{z}$. (8) is $p\gamma + (1 - p)\alpha \geq p\rho$. Since $\gamma \leq \frac{1}{2}$ and $\alpha \leq 1 - \rho$, (8) implies that $p \leq 2(1 - \rho) \leq \frac{2}{z}$, which imposes a constraint on p if $z > 2$. In the appendix, we show that the same kind of argument can be used to show (6) for arbitrary γ_d, β_d satisfying (7).■

Remarks:

1. If p is small enough in example 1 (with respect to z , which is kept fixed, as the other parameters), equilibria of $\Gamma_\infty(p)$ are easily constructed. For instance, if $p \leq \frac{2}{z}$, an *ex post* efficient equilibrium as above is achievable (*i.e.*, condition (3) for $\theta_1 = z$ and $\theta'_1 = \omega$ is no problem).
2. Proposition 1 tells us that, when an equilibrium exists in the infinitely repeated public good game of example 1, the associated payoff can as well be achieved at a completely revealing equilibrium; in particular,

the players cannot benefit from behaving as if they were “greedy” when their type is “normal”. Such a result does not hold in a finitely repeated game. For instance, Fudenberg and Tirole (1991) (example 8.3, p. 333) consider a two stage version of the public good game in which the players’ types belong to the unit interval. They show that, in any perfect Bayesian Nash equilibrium, the players contribute less in the first period than in the second one: “Each player gains by developing a reputation for not being willing to supply the public good”.

3. If there is uncertainty on the type of only one of the players, an equilibrium always exists (see Shalev (1994)).

4 Bayesian game with commitment

In this section, we show that the characterization in proposition 1 implies a relationship between “repetition” and “cooperation”. Under complete information, the standard folk theorem states that the set of Nash equilibrium payoffs of an infinitely repeated game coincides with the set of feasible and individually rational payoffs of the one-shot game, which in turn can be interpreted as the set of cooperative solutions of the one-shot game (see e.g. Kalai *et al.* (2010)). Under incomplete information, a natural candidate for the latter set is Myerson (1991)’s set of feasible, *incentive compatible* and *interim* individually rational payoffs in the (one-shot) Bayesian game $\Gamma(q)$ (see Forges (2013)). We denote this set as $\mathcal{F}[\Gamma(q)]$ and define it precisely below. Myerson (1991)’s definitions take a simpler form in our framework of independent private values. We then establish a partial folk theorem, namely that $\mathcal{F}[\Gamma(q)]$ contains $\mathcal{N}[\Gamma_\infty(q)]$, the set of Nash equilibrium payoffs of the infinitely repeated game $\Gamma_\infty(q)$.

A payoff $x = (x_i)_{1 \leq i \leq n} = ((x_i(\theta_i))_{\theta_i \in \Theta_i})_{1 \leq i \leq n}$ is *feasible* in $\Gamma(q)$ if there exists a correlated strategy $\pi(\theta) \in \Delta(A)$, $\theta \in \Theta$, achieving x , namely

$$x_i(\theta_i) = \sum_{\theta_{-i}} q_{-i}(\theta_{-i}) u_i(\theta_i, \pi(\theta_i, \theta_{-i})) \quad i = 1, \dots, n, \quad \theta_i \in \Theta_i \quad (9)$$

A feasible payoff x achieved through π (as in (9)) is *incentive compatible* if

$$x_i(\theta_i) \geq \sum_{\theta_{-i}} q_{-i}(\theta_{-i}) u_i(\theta_i, \pi(\theta'_i, \theta_{-i})) \quad \text{for every } i, \theta_i, \theta'_i \in \Theta_i \quad (10)$$

A payoff x is *interim individually rational* if, for every player i , there exists a correlated strategy $\tau_{-i} \in \Delta(A_{-i})$ of players $j \neq i$ such that¹³

$$x_i(\theta_i) \geq \max_{a_i \in A_i} u_i(\theta_i, a_i, \tau_{-i}) \quad \text{for every } \theta_i \in \Theta_i \quad (11)$$

Let $INTIR_i[\Gamma(q)]$ be the set of all vector payoffs satisfying the previous property for player i . Observe that the previous definition describes a *set* of *vector* payoffs which cannot be reduced to a “corner set” (of the form $x_i(\theta_i) \geq w_i(\theta_i)$, $\theta_i \in \Theta_i$, for some well-defined individually rational level $w_i(\theta_i)$). By contrast, *ex post* individually rational payoffs are described by a “corner set”, since $(v_i(\theta_i))_{\theta_i \in \Theta_i}$ is defined without ambiguity by (1).

The set $\mathcal{F}[\Gamma(q)]$ is formally defined as the set of payoffs satisfying (9), (10) and (11).¹⁴ $\mathcal{F}[\Gamma(q)]$ contains the set of Nash equilibrium payoffs of $\Gamma(q)$ and is thus not empty.

In the next two statements, we make use of uniform punishment strategies, which were not assumed to exist earlier in this section.

Lemma 1 *Let $\Gamma(q)$ be a Bayesian game with independent private values and let x be a feasible payoff in $\Gamma(q)$. If x is interim individually rational (namely, (11)), x is ex post individually rational (namely, (5)): $INTIR_i[\Gamma(q)] \subseteq EXPIR_i[\Gamma(q)]$ for every player i . If there exist uniform punishment strategies, namely (2), then the reverse also holds: $INTIR_i[\Gamma(q)] = EXPIR_i[\Gamma(q)]$ for every player i .*

The proof of lemma 1 is straightforward and therefore omitted. The intuition behind the first part is that players $j \neq i$ can impose a harder punishment to player i if they know player i 's type θ_i (*i.e.*, *ex post*). For the second part, a uniform punishment strategy of players $j \neq i$ against player i provides an appropriate correlated strategy τ_{-i} in (11).

Proposition 2 *Let $\Gamma(q)$ be a Bayesian game with independent private values in which uniform punishment strategies are available: $\mathcal{N}[\Gamma_\infty(q)] \subseteq \mathcal{F}[\Gamma(q)]$.*

¹³Literally, Myerson (1991)'s *interim* individual rationality condition requires that there exists a type dependent correlated strategy of players $j \neq i$, $\tau_{-i}(t_{-i}) \in \Delta(A_{-i})$, $t_{-i} \in T_{-i}$, such that $x_i(t_i) \geq \max_{a_i \in A_i} \sum_{t_{-i}} q_{-i}(t_{-i}) u_i(t_i, a_i, \tau_{-i}(t_{-i}))$ for every $t_i \in T_i$. But, with independent private values, (11) is an equivalent formulation, since $u_i(t_i, \cdot)$ is linear.

¹⁴Peters and Szentes (2012) argue that the set of solutions that can be achieved under interim commitment in $\Gamma(q)$ should be smaller set than $\mathcal{F}[\Gamma(q)]$ but, under an assumption which is similar to our uniform punishment strategies, they recover $\mathcal{F}[\Gamma(q)]$.

Proof: The proposition readily follows from the characterizations of $\mathcal{N}[\Gamma_\infty(q)]$ (in proposition 1) and $\mathcal{F}[\Gamma(q)]$ ((9), (10) and (11) above): the equality in (3) is (9), the inequality in (3) implies (10) and (5), which in turn implies (11) by lemma 1. ■

As illustrated on example 1, unlike $\mathcal{F}[\Gamma(q)]$, $\mathcal{N}[\Gamma_\infty(q)]$ can be empty. Hence, $\mathcal{N}[\Gamma_\infty(q)]$ can be strictly included in $\mathcal{F}[\Gamma(q)]$. In other words, it may happen that repetition makes some form of cooperation possible, but does not exhaust the players' cooperation possibilities. Example 2 below, taken from Koren (1992), further illustrates the possible strict inclusion¹⁵.

Example 2: A game in which $\mathcal{N}[\Gamma_\infty(q)]$ is not empty and strictly included in $\mathcal{F}[\Gamma(q)]$

We will study a variant of the well-known battle of the sexes. Each player has two possible types: $\Theta_i = \{n, g\}$, $i = 1, 2$, and two possible actions: $A_i = \{c, d\}$, $i = 1, 2$. We denote as $p_i \in [0, 1]$ the probability that player i 's type is n (namely, $q_i = (p_i, 1 - p_i)$). Payoffs are described by the following matrices:

		$\theta_2 = n$		$\theta_2 = g$	
		c	d	c	d
$\theta_1 = n$	c	3, 1	0, 0	3, 1	0, 3
	d	0, 0	1, 3	0, 1	1, 3
$\theta_1 = g$	c	3, 1	3, 0	3, 1	3, 3
	d	1, 0	1, 3	1, 1	1, 3

When $\theta_1 = n$, player 1 prefers c to d , but also prefers to make the same choice as the other player. When $\theta_1 = g$, player 1 just prefers c to d , independently of the choice of the other player. The preferences of player 2 are similar. In this game, $v_i(n) = \frac{3}{4}$, $v_i(g) = 3$, $i = 1, 2$. A uniform punishment strategy of player 1 (*resp.*, 2) is to play c with probability $\frac{3}{4}$ (*resp.*, $\frac{1}{4}$).

Let us consider the (*ex post* efficient) correlated strategy $\pi(\theta)$, $\theta \in \Theta$,

¹⁵For appropriate values of q , it also happens in the public good games of example 1 that $\mathcal{N}[\Gamma_\infty(q)]$ is not empty and is strictly included in $\mathcal{F}[\Gamma(q)]$. However, a full characterization of $\mathcal{N}[\Gamma_\infty(q)]$ seems much harder in example 1 than in Koren (1992)'s example.

defined by

		$\theta_2 = n$		$\theta_2 = b$	
		c	d	c	d
$\theta_1 = n$	c	$\frac{1}{2}$	0	0	0
	d	0	$\frac{1}{2}$	0	1
$\theta_1 = g$	c	1	0	0	1
	d	0	0	0	0

(12)

It is easily checked that $\pi(\theta)$ satisfies (10) and (11), namely, induces a payoff in $\mathcal{F}[\Gamma(q)]$, if and only if $p_i \leq \frac{1}{2}$, $i = 1, 2$. Similarly, in order to induce a payoff in $\mathcal{N}[\Gamma_\infty(q)]$, $\pi(\theta)$ must satisfy (3); in particular, player 1 of type $\theta_1 = n$ cannot gain by pretending to be of type $\theta'_1 = g$, namely,

$$p_2 + 1 \geq p_2 \max \left\{ 3, \frac{3}{4} \right\} + (1 - p_2) \max \left\{ 0, \frac{3}{4} \right\} \Leftrightarrow p_2 \leq \frac{1}{5}$$

The previous condition illustrates that, as expected, player 1 has more deviation possibilities at a (completely revealing) Nash equilibrium of $\Gamma_\infty(q)$ than at an *interim* cooperative solution of $\Gamma(q)$. Imagine that player 1 is of type n but pretends to be of type g at the first stage of $\Gamma_\infty(q)$. Then he learns player 2's type θ_2 and faces $\pi(g, \theta_2)$. If $\theta_2 = n$, player 1 gets the best payoff 3 by playing according to $\pi(g, n)$. However, if $\theta_2 = g$, player 1 gets 0 by playing according to $\pi(g, g)$. In this case, he should not play according to $\pi(g, g)$ but rather play c with probability $\frac{3}{4}$ at every stage in order to guarantee himself $\frac{3}{4}$. By checking the other equilibrium conditions in (3), we get that $\pi(\theta)$ induces a payoff in $\mathcal{N}[\Gamma_\infty(q)]$ if and only if $p_i \leq \frac{1}{5}$, $i = 1, 2$.

On the other hand, as already pointed out in Koren (1992), the correlated strategy defined by

		$\theta_2 = n$		$\theta_2 = g$	
		c	d	c	d
$\theta_1 = n$	c	0	0	0	0
	d	0	1	0	1
$\theta_1 = g$	c	$\frac{3}{4}$	$\frac{1}{4}$	0	1
	d	0	0	0	0

induces a payoff in $\mathcal{N}[\Gamma_\infty(q)]$ if and only if $p_2 \leq \frac{1}{6}$. There are thus many probability distributions $q \in \Delta(\Theta)$ for which $\pi(\theta)$ defined by (12) induces a payoff in $\mathcal{F}[\Gamma(q)]$, and at the same time, $\mathcal{N}[\Gamma_\infty(q)]$ is not empty but does not contain the payoff defined by (12). ■

5 Discounted infinitely repeated Bayesian game

In this section, we turn to the δ -discounted version $\Gamma_\delta(q)$ of the infinitely repeated game, for a given discount factor $\delta \in (0, 1)$. $\Gamma_\delta(q)$ is played as in section 3, but the payoffs associated with a sequence of actions $\mathbf{a} = (a^t)_{t \geq 0} \in A^\mathbb{N}$ are now evaluated as

$$U_i^\delta(\theta_i, \mathbf{a}) = (1 - \delta) \sum_{t=0}^{\infty} \delta^t u_i(\theta_i, a^t) \quad \text{for every } i, \theta_i$$

Let us denote as $\mathcal{N}[\Gamma_\delta(q)]$ the set of all (interim expected) Nash equilibrium payoffs of $\Gamma_\delta(q)$. By the same arguments as under complete information, $\mathcal{N}[\Gamma_\delta(q)]$ is nonempty and compact, for every $\delta \in (0, 1)$. An interesting question is whether $\lim_{\delta \rightarrow 1} \mathcal{N}[\Gamma_\delta(q)]$ is also nonempty, for some appropriate definition of the limit of a sequence of sets. If there exist belief free equilibria or more generally, if $\mathcal{N}[\Gamma_\infty(q)]$ is nonempty¹⁶, one has a candidate that could belong to $\lim_{\delta \rightarrow 1} \mathcal{N}[\Gamma_\delta(q)]$. But the question seems especially hard (and also more important) when $\mathcal{N}[\Gamma_\infty(q)] = \emptyset$, as may happen when $\Gamma_\infty(q)$ is associated to the one-shot public good game $PG(2, 1, q)$ (recall example 1 in section 3.2). A corollary of the proposition below is that, in this game, for every q , $\liminf_{\delta \rightarrow 1} \mathcal{N}[\Gamma_\delta(q)]$ is nonempty.¹⁷ We denote as $\lfloor x \rfloor$ the largest integer not greater than x .

Proposition 3 *Let $\Gamma_\delta(q)$ be the δ -discounted infinitely repeated game associated with the public good game $PG(2, 1, q)$, with beliefs $q = ((p_i, 1 - p_i)(p_j, 1 - p_j))$, such that $0 < p_j \leq p_i < 1$. For every $\delta \in (1 - \omega, 1)$, $\Gamma_\delta(q)$ has an equilibrium of the following form:*

- *at every stage, independently of the past history, a “greedy” type player does not contribute.*
- *before stage 0, a “normal” type player i (resp., player j) chooses $\tau_i \in \{0, 2, \dots, 2T\}$ (resp., $\tau_j \in \{1, 3, \dots, 2T + 1\}$) according to an appropriate probability distribution, where $T = \left\lfloor \frac{\log(1-p_j)}{\log\left(\frac{\delta-1+\omega}{\delta-\delta^2(1-\omega)}\right)} \right\rfloor$.*

¹⁶Or, more carefully, if $\mathcal{N}[\Gamma_\infty(q)]$ has a nonempty interior.

¹⁷With the definition $x \in \liminf_{\delta \rightarrow 1} \mathcal{N}[\Gamma_\delta(q)] \Leftrightarrow \forall \delta_n \rightarrow 1 \exists x_n \in \mathcal{N}[\Gamma_{\delta_n}(q)]$ such that $x_n \rightarrow x$.

- at every stage t , if the other player contributed at least once in the past, a “normal” type player does not contribute. Otherwise, namely, if the other player played d at stages $0, \dots, t-1$, and if $t \geq \tau_i$ (resp., $t \geq \tau_j$), a “normal” type player i (resp., player j) contributes.

When the discount factor $\delta \rightarrow 1$, the sequence of interim payoffs $x(\delta)$ associated to this equilibrium converges to

$$\begin{aligned} x_i(n) &= x_j(n) = 1 \\ x_i(g) &= x_j(g) = z + (1 - \omega) \left(1 - (1 - p_j)^{\frac{1}{1-\omega}} \right) \end{aligned}$$

As expected, the payoff x does depend on the players’ beliefs (but only through $1 - p_j \geq 1 - p_i$, namely through the probability of being “greedy” for the player who is more likely to be so).

At equilibrium, the players behave as in a war of attrition: the players do not contribute until one of them gives in to his opponent and contributes forever.¹⁸ The contributing player thus reveals that his type is normal, whereas the opponent never contributes and keeps his type unknown. Note that, once one of the players has revealed that his type is normal, the players agree on a fixed sequence of moves, as in the standard proof of the Folk theorem with complete information. Under our assumption of private values, given such an agreement, the player who has not revealed his type would not mind revealing it (provided that the agreed sequence of moves is not modified, of course). However, the previous equilibrium cannot be reduced to a completely revealing one. Indeed, in the discounted game, the time before revelation is costly and matters at equilibrium. The fact that payoffs are discounted is thus critical in the equilibrium described by the previous proposition. In

¹⁸Atakan and Ekmekci (2013) make use of strategies involving a war of attrition to show the existence of a unique equilibrium in a class of two-sided reputation models. Our example fits in that class because (i) when both players are normal, the game has strictly conflicting interests and (ii) the greedy player behaves as a commitment type. Leaving aside the differences of scope and treatment in reputation models and Bayesian repeated games, a main difference between the construction in proposition 3 and Atakan and Ekmekci (2013)’s theorem 2 is that they use a standard war of attrition with *continuous* time, while we stick to the *discrete* time that is standard in repeated games. Hence we have to introduce some coordination between the players, who alternate in contributing with positive probability. To the best of our knowledge, this feature of the equilibrium strategies is completely new.

the undiscounted game, a normal player is always better off waiting for his opponent to give in, because waiting is free.

Proof of proposition 3

In order to show that the strategies described in the statement define an equilibrium, we first construct appropriate distributions for τ_i and τ_j . We adopt the convention that $\tau_i = +\infty$ if player i is greedy. If player i is normal, τ_i has the value that player i had chosen before stage 0. This means that player i will start to contribute at stage τ_i if player does not contribute before him. We adopt the same convention for player j .

Let us compute the distribution \mathbb{P} that an outside observer will assign to τ_i and τ_j , prior to stage 0. Note that player i (resp. j) is an outside observer of τ_j (resp. τ_i).

Probability distribution of τ_j

We have $\mathbb{P}(\tau_j = t) = 0$ if t is an even number or if $t > 2T + 1$. Also, $\mathbb{P}(\tau_j = +\infty) = 1 - p_j$. Observe that if i never contributes and if j is normal, then j is right to start to contribute at least at stage $2T + 1$. Indeed, his stage payoff will become 1 instead of ω from then on.

The law of τ_j is such that player i cannot be tempted to start to contribute at odd stages or at stages $t \geq 2T + 1$: if i starts to contribute at such a stage t and if j plays as assumed, it is then strictly better for i to start to contribute at stage $t - 1$, as his stage payoffs will switch from ω to 1 earlier. Yet, it remains to show that i can not increase his overall payoff by starting to contribute at stage $2T + 2$. Also, player i of normal type has to be indifferent between starting contributing at any stage $t = 0, 2, \dots, 2T$. Let us denote $G_i(2t)$ the overall payoff of player i when his type is normal and when he is willing to start to contribute at stage $2t$, the law of τ_j being fixed. We have, for any $t \geq 0$:

$$G_i(2t) = \sum_{s=0}^{t-1} \mathbb{P}(\tau_j = 2s + 1)(\omega + \delta^{2s+1}) + \mathbb{P}(\tau_j > 2t - 1)(\omega + (1 - \omega)\delta^{2t}).$$

In particular, $G_i(0) = 1$.

So, the law of τ_j has to be such that $1 = G_i(2) = \dots = G_i(2T)$, and $1 \geq G_i(2T + 2)$.

From equality $G_i(2) = 1$, we get

$$\mathbb{P}(\tau_j = 1) = \frac{(1 - \delta^2)(1 - \omega)}{\delta - \delta^2(1 - \omega)}.$$

Now, let us show by induction that:

$$\forall t \in \{0, 1, \dots, T-1\}, \mathbb{P}(\tau_j = 2t+1) = (1 - \mathbb{P}(\tau_j = 1))^t \mathbb{P}(\tau_j = 1).$$

The property is obvious for $t = 0$. Then, if the property holds for each $s \in \{0, 1, \dots, t-1\}$ ($t \leq T-1$), let us show that it also holds for $s = t$. The fact that $G_i(2t+2) = G_i(2t)$ implies that $G_i(2t+2) - G_i(2t) = 0$, which gives:

$$\begin{aligned} & \mathbb{P}(\tau_j = 2t+1) (\omega + \delta^{2t+1}) + \mathbb{P}(\tau_j > 2t+1) (\omega + (1-\omega)\delta^{2t+2}) \\ & - \mathbb{P}(\tau_j > 2t-1) (\omega + (1-\omega)\delta^{2t}) = 0. \end{aligned}$$

As $\mathbb{P}(\tau_j > 2t+1) = \mathbb{P}(\tau_j > 2t-1) - \mathbb{P}(\tau_j = 2t+1)$, we get:

$$\begin{aligned} \mathbb{P}(\tau_j = 2t+1) &= \mathbb{P}(\tau_j > 2t-1) \frac{(1-\delta^2)(1-\omega)}{\delta - \delta^2(1-\omega)} \\ &= \mathbb{P}(\tau_j > 2t-1) \mathbb{P}(\tau_j = 1) \\ &= \left(1 - \sum_{s=0}^{t-1} \mathbb{P}(\tau_j = 2s+1)\right) \mathbb{P}(\tau_j = 1) \\ &= \left(1 - \sum_{s=0}^{t-1} (1 - \mathbb{P}(\tau_j = 1))^s \mathbb{P}(\tau_j = 1)\right) \mathbb{P}(\tau_j = 1) \\ &= \left(1 - \mathbb{P}(\tau_j = 1) \frac{1 - (1 - \mathbb{P}(\tau_j = 1))^t}{1 - (1 - \mathbb{P}(\tau_j = 1))}\right) \mathbb{P}(\tau_j = 1) \\ &= (1 - \mathbb{P}(\tau_j = 1))^t \mathbb{P}(\tau_j = 1), \end{aligned} \tag{13}$$

which is the property for $s = t$.

$$\text{As } T = \left\lfloor \frac{\log(1-p_j)}{\log\left(\frac{\delta-1+\omega}{\delta-\delta^2(1-\omega)}\right)} \right\rfloor = \left\lfloor \frac{\log(1-p_j)}{\log(1-\mathbb{P}(\tau_j=1))} \right\rfloor, \text{ this implies that}$$

$$\begin{aligned} \sum_{t=0}^{T-1} \mathbb{P}(\tau_j = 2t+1) &= \sum_{t=0}^{T-1} (1 - \mathbb{P}(\tau_j = 1))^t \mathbb{P}(\tau_j = 1) \\ &= 1 - (1 - \mathbb{P}(\tau_j = 1))^T \leq p_j \end{aligned}$$

The definition of T also implies that

$$\sum_{t=0}^T (1 - \mathbb{P}(\tau_j = 1))^t \mathbb{P}(\tau_j = 1) > p_j.$$

Since $\sum_{t=0}^T \mathbb{P}(\tau_j = 2t + 1) = p_j$, we have

$$\mathbb{P}(\tau_j = 2T + 1) < (1 - \mathbb{P}(\tau_j = 1))^T \mathbb{P}(\tau_j = 1),$$

and this shows that $G_i(2T + 2) < G_i(2T) = 1$. Indeed if we assume by contradiction that $G_i(2T + 2) - G_i(2T) \geq 0$, then, as in equation (13), we have that $\mathbb{P}(\tau_j = 2T + 1) \geq (1 - \mathbb{P}(\tau_j = 1))^T \mathbb{P}(\tau_j = 1)$.

Probability distribution of τ_i

We have $\mathbb{P}(\tau_i = t) = 0$ if t is an odd number or if $t > 2T$. Also, note that $\mathbb{P}(\tau_i = +\infty) = 1 - p_i$. Similar arguments as the ones used for the law of τ_j show that j should not start contributing at any stage $2t$, $t \geq 1$, or at any stage $t > 2T + 1$. Yet, it remains to show that he should not start contributing at stage 0. Let us denote $G_j(2t + 1)$ the overall payoff of player j when he is normal and is willing to start contributing at stage $2t + 1$, the law of τ_i being fixed. We have, for any $t \geq 0$:

$$G_j(2t + 1) = \sum_{s=0}^t \mathbb{P}(\tau_i = 2s)(\omega + \delta^{2s}) + \mathbb{P}(\tau_i > 2t)(\omega + (1 - \omega)\delta^{2t+1}).$$

The law of τ_i has to be such that $G_j(1) = \dots = G_j(2T + 1)$, and such that $1 \leq G_j(1)$, for j not to be tempted to start contributing at stage 0. The fact that $G_j(1) = G_j(3)$ implies that

$$\mathbb{P}(\tau_i = 2) = \frac{(1 - \delta^2)(1 - \omega)(1 - \mathbb{P}(\tau_i = 0))}{\delta - \delta^2(1 - \omega)} = \mathbb{P}(\tau_j = 1)(1 - \mathbb{P}(\tau_i = 0)),$$

and then, as before, one can prove by induction that:

$$\begin{aligned} \forall t \in \{1, \dots, T\}, \mathbb{P}(\tau_i = 2t) &= \left(\frac{\delta - 1 + \omega}{\delta - \delta^2(1 - \omega)} \right)^{t-1} \mathbb{P}(\tau_i = 2) \\ &= (1 - \mathbb{P}(\tau_j = 1))^{t-1} \mathbb{P}(\tau_j = 1)(1 - \mathbb{P}(\tau_i = 0)) \end{aligned}$$

The fact that $\mathbb{P}(\tau_i = +\infty) = 1 - p_i$ is equivalent to:

$$p_i = \sum_{t=0}^T \mathbb{P}(\tau_i = 2t),$$

which is equivalent to

$$\begin{aligned}
p_i &= \mathbb{P}(\tau_i = 0) + \sum_{t=1}^T \mathbb{P}(\tau_i = 2t) \\
&= \mathbb{P}(\tau_i = 0) + \mathbb{P}(\tau_j = 1)(1 - \mathbb{P}(\tau_i = 0)) \sum_{t=1}^T (1 - \mathbb{P}(\tau_j = 1))^{t-1} \\
&= \mathbb{P}(\tau_i = 0) + (1 - \mathbb{P}(\tau_i = 0)) \left(1 - (1 - \mathbb{P}(\tau_j = 1))^{T-1}\right)
\end{aligned}$$

so that

$$\begin{aligned}
\mathbb{P}(\tau_i = 0) &= 1 - \frac{1 - p_i}{(1 - \mathbb{P}(\tau_j = 1))^{T-1}} \\
&= 1 - (1 - \mathbb{P}(\tau_j = 1)) \frac{1 - p_i}{(1 - \mathbb{P}(\tau_j = 1))^T} \\
&\geq 1 - (1 - \mathbb{P}(\tau_j = 1)) \frac{1 - p_i}{1 - p_j} \\
&\geq 1 - (1 - \mathbb{P}(\tau_j = 1)) \\
&\geq \mathbb{P}(\tau_j = 1).
\end{aligned}$$

We then have:

$$\begin{aligned}
G_j(1) &= \mathbb{P}(\tau_i = 0)(1 + \omega) + \mathbb{P}(\tau_i > 0)(\omega + (1 - \omega)\delta) \\
&= \mathbb{P}(\tau_i = 0)(1 + \omega) + (1 - \mathbb{P}(\tau_i = 0))(\omega + (1 - \omega)\delta) \\
&= \mathbb{P}(\tau_i = 0)(1 - \delta(1 - \omega)) + \omega + (1 - \omega)\delta \\
&\geq \mathbb{P}(\tau_j = 1)(1 - \delta(1 - \omega)) + \omega + (1 - \omega)\delta \\
&= 1 + \frac{1 - \delta}{\delta}(1 - \omega),
\end{aligned}$$

so that $G_j(1) \geq 1$.

Type dependent probabilities defining the equilibrium strategies

The values of $\mathbb{P}(\tau_j = 2t + 1)$ and of $\mathbb{P}(\tau_i = 2t)$ enables us to specify player i and j 's strategies: a “normal” player i chooses stage 0 with probability

$$\alpha_0^i = \frac{\mathbb{P}(\tau_i = 0)}{p_i} = \frac{1}{p_i} \left(1 - \frac{1 - p_i}{(1 - \mathbb{P}(\tau_j = 1))^{T-1}}\right) = \frac{1}{p_i} \left(1 - \frac{1 - p_i}{\left(\frac{\delta - 1 + \omega}{\delta - \delta^2(1 - \omega)}\right)^{T-1}}\right)$$

and stage $2t$ ($1 \leq t \leq T$) with probability

$$\begin{aligned}
\alpha_{2t}^i &= \frac{\mathbb{P}(\tau_i = 2t)}{p_i} = \frac{1}{p_i} (1 - \mathbb{P}(\tau_j = 1))^{t-1} \mathbb{P}(\tau_j = 1) (1 - \mathbb{P}(\tau_i = 0)) \\
&= \frac{1}{p_i} \left(\frac{\delta - 1 + \omega}{\delta - \delta^2(1 - \omega)} \right)^{t-1} \frac{(1 - \delta^2)(1 - \omega)}{\delta - \delta^2(1 - \omega)} \frac{1 - p_i}{\left(\frac{\delta - 1 + \omega}{\delta - \delta^2(1 - \omega)} \right)^{T-1}} \\
&= \frac{1 - p_i}{p_i} \left(\frac{\delta - 1 + \omega}{\delta - \delta^2(1 - \omega)} \right)^{t-T} \frac{(1 - \delta^2)(1 - \omega)}{\delta - \delta^2(1 - \omega)},
\end{aligned}$$

A “normal” player j chooses stages $2t + 1$ ($0 \leq t \leq T - 1$) with probability

$$\begin{aligned}
\alpha_{2t+1}^j &= \frac{\mathbb{P}(\tau_j = 2t + 1)}{p_j} = \frac{1}{p_j} (1 - \mathbb{P}(\tau_j = 1))^t \mathbb{P}(\tau_j = 1) \\
&= \frac{1}{p_j} \left(\frac{\delta - 1 + \omega}{\delta - \delta^2(1 - \omega)} \right)^t \frac{(1 - \delta^2)(1 - \omega)}{\delta - \delta^2(1 - \omega)}
\end{aligned}$$

and stage $2T + 1$ with probability

$$\begin{aligned}
\alpha_{2T+1}^j &= \frac{\mathbb{P}(\tau_j = 2T + 1)}{p_j} = \frac{1}{p_j} \left(p_j - \sum_{t=0}^{T-1} \mathbb{P}(\tau_j = 2t + 1) \right) \\
&= \frac{1}{p_j} \left(p_j - 1 + (1 - \mathbb{P}(\tau_j = 1))^T \right) \\
&= \frac{1}{p_j} \left(p_j - 1 + \left(\frac{\delta - 1 + \omega}{\delta - \delta^2(1 - \omega)} \right)^T \right).
\end{aligned}$$

Our construction ensures that this strategic profile is an equilibrium.

Limit payoffs

Let us now give limit payoffs as δ goes to 1.

For any $\delta \in (1 - \omega, 1)$, by construction player i 's overall expected payoff is 1

if his type is normal. If his type is greedy, his payoff is:

$$\begin{aligned}
& z + \sum_{t=0}^T \mathbb{P}(\tau_j = 2t + 1) \delta^{2t+1} \\
&= z + \sum_{t=0}^{T-1} (1 - \mathbb{P}(\tau_j = 1))^t \mathbb{P}(\tau_j = 1) \delta^{2t+1} + o(1) \\
&= z + \frac{\delta \mathbb{P}(\tau_j = 1)}{1 - \delta^2 (1 - \mathbb{P}(\tau_j = 1))} \left(1 - \delta^{2T} (1 - \mathbb{P}(\tau_j = 1))^T \right) + o(1)
\end{aligned}$$

One can check that

$$(1 - \mathbb{P}(\tau_j = 1))^T \xrightarrow{\delta \rightarrow 1} 1 - p_j.$$

and that

$$\frac{\delta \mathbb{P}(\tau_j = 1)}{1 - \delta^2 (1 - \mathbb{P}(\tau_j = 1))} = 1 - \omega.$$

Moreover,

$$\delta^{2T} \underset{\delta \rightarrow 1}{\sim} \delta^{2 \frac{\log(1-p_j)}{\log\left(\frac{\delta-1+\omega}{\delta-\delta^2(1-\omega)}\right)}} = \exp\left(2 \log \delta \frac{\log(1-p_j)}{\log\left(\frac{\delta-1+\omega}{\delta-\delta^2(1-\omega)}\right)}\right),$$

so that, by a straightforward calculation:

$$\delta^{2T} \underset{\delta \rightarrow 1}{\sim} (1 - p_j)^{\frac{\omega}{1-\omega}}.$$

Finally, limit equilibrium payoff is 1 for a normal player i and

$$z + (1 - \omega) \left(1 - (1 - p_j)^{\frac{1}{1-\omega}} \right)$$

for a greedy player i . By similar means, the same holds for player j .

Limit probabilities over joint actions

Let us define the probability

$$\pi_\delta(\theta_i, \theta_j)(a) = \sum_{t=0}^{+\infty} (1 - \delta) \delta^t \mathbb{P}(a_t = a | \theta_i, \theta_j).$$

The values of $\lim_{\delta \rightarrow 1} \pi_\delta$ are given by¹⁹

		$\theta_i = \omega$		$\theta_i = z$	
		c	d	c	d
$\theta_j = \omega$	c	0	$\frac{1-P}{2}$	0	$\frac{\gamma}{p_j}$
	d	$\frac{1-P}{2}$	P	0	$1 - \frac{\gamma}{p_j}$
$\theta_j = z$	c	0	0	0	0
	d	$\frac{\gamma}{p_i}$	$1 - \frac{\gamma}{p_i}$	0	1

where

$$\gamma = (1 - \omega) \left(1 - (1 - p_j)^{\frac{1}{1-\omega}} \right)$$

and

$$\begin{aligned} P = & \left(1 - \frac{1}{p_i} \frac{1 - p_i}{1 - p_j} \right) \left(1 - \frac{1}{p_j} \right) \left(1 - (1 - p_i)^{\frac{\omega}{1-\omega}} \right) \\ & + \left(\frac{1}{p_i} \frac{1 - p_i}{1 - p_j} \left(1 - \frac{1}{p_j} \right) + \frac{1}{p_j} \left(1 - \frac{1}{p_i} \frac{1 - p_i}{1 - p_j} \right) \right) \omega \left(1 - (1 - p_j)^{\frac{1}{1-\omega}} \right) \\ & + \frac{1}{p_i p_j} \frac{1 - p_i}{1 - p_j} \frac{\omega}{2 - \omega} \left(1 - (1 - p_j)^{\frac{2-\omega}{1-\omega}} \right) \end{aligned} \quad \blacksquare$$

6 Role of the assumptions

6.1 Independent private values

Independent private values are crucial in proposition 1. Without this assumption, the Nash equilibria of $\Gamma_\infty(q)$ are no longer payoff equivalent to completely revealing equilibria, even if there are two players and only one of them has private information (see Hart (1985) and Aumann and Maschler (1995)).

6.2 Uniform punishments

In the case of two players, if values are private and independent in $\Gamma(q)$, Koren (1992) proves that the Nash equilibria of $\Gamma_\infty(q)$ are payoff equivalent

¹⁹The type dependent distributions described by $\lim_{\delta \rightarrow 1} \pi_\delta$ are not ex post individually rational (recall (4)) when one of the player is normal and the other is greedy. Hence, as expected, they do not satisfy (3) and do not define an equilibrium of the undiscounted game.

to completely revealing equilibria without assuming uniform punishments (*i.e.*, (2)). However, in this more general case, the equilibrium conditions can take a more complex form than (3). Examples 3 and 4 below illustrate how the absence of uniform punishments modifies the results.

In example 3, the conditions (3) of proposition 1 are no longer sufficient for an equilibrium. Proposition 2 does not hold either: we construct an equilibrium payoff in $\Gamma_\infty(q)$ which does not belong to $\mathcal{F}[\Gamma(q)]$, *i.e.*, cannot be achieved through commitment in $\Gamma(q)$.

In example 4, an assumption weaker than uniform punishments holds, which guarantees that the Nash equilibrium payoffs of $\Gamma_\infty(q)$ can be characterized *exactly* as in proposition 1, by (3). However, proposition 2 still fails.

In both examples 3 and 4, there are two players and only player 1 has private information ($|\Theta_2| = 1$, $A = A_1 \times A_2$), so that the conditions in proposition 1 reduce to: there exists $\pi(\theta_1) \in \Delta(A)$, $\theta_1 \in \Theta_1$, such that, for player 1,

$$\begin{aligned} x_1(\theta_1) &= u_1(\theta_1, \pi(\theta_1)) \\ &\geq u_1(\theta_1, \pi(\theta'_1)) \quad \forall \theta_1, \theta'_1 \in \Theta_1 \quad \text{i.e., incentive compatibility} \end{aligned} \quad (14)$$

$$\geq v_1(\theta_1) \quad \forall \theta_1 \in \Theta_1 \quad \text{i.e., ex post individual rationality} \quad (15)$$

and, for player 2,

$$x_2 = u_2(\pi(\theta_1)) \geq v_2 \quad \forall \theta_1 \in \Theta_1 \quad \text{i.e., ex post individual rationality} \quad (16)$$

As shown by Hart (1985), in order to characterize the equilibrium payoffs of $\Gamma_\infty(q)$, *ex post* individual rationality (namely, (5) or (15) above) is not sufficient. A stronger condition, which makes full use of the fact that $\Gamma_\infty(q)$ is an *infinitely repeated* game, is needed. This condition is formally stated below, in the current framework of lack of information on one side.²⁰ Let $val_1[u]$ denote the value to player 1 of the one-shot game with payoff function u .

Definition A vector payoff $x_1 = (x_1(\theta_1))_{\theta_1 \in \Theta_1}$ is individually rational for player 1 in the *infinitely repeated game* $\Gamma_\infty(q)$ if and only if

$$\forall p_1 \in \Delta(\Theta_1), \quad \sum_{\theta_1} p_1(\theta_1) x_1(\theta_1) \geq val_1 \left[\sum_{\theta_1} p_1(\theta_1) u_1(\theta_1, \cdot) \right] \quad (17)$$

²⁰The same condition holds as well in two-person games with independent private values (see Koren (1992)).

Let $INTIR_1[\Gamma_\infty(q)]$ be the set of vector payoff that are individually rational for player 1 in the *infinitely repeated game* $\Gamma_\infty(q)$. The previous definition is justified by Blackwell (1956)'s approachability theorem: condition (17) is necessary and sufficient for player 2 to have a strategy in the *infinitely repeated game* $\Gamma_\infty(q)$ such that player 1's payoff cannot exceed $x_1(\theta_1)$ when he is of type θ_1 .

Let us compare $INTIR_1[\Gamma_\infty(q)]$ with the two sets of individually rational payoffs introduced for the one-shot game $\Gamma(q)$, namely, $EXPIR_1[\Gamma(q)]$ and $INTIR_1[\Gamma(q)]$. First of all, player 2 can use a punishment strategy of the one-shot game at every stage of the infinitely repeated game: as a consequence of Blackwell (1956)'s characterization, (11) implies (17). Furthermore, (17) holds in particular at the extreme points of $\Delta(\Theta_1)$, so that it implies *ex post* individual rationality (*i.e.*, (15)). To sum up,

$$INTIR_1[\Gamma(q)] \subseteq INTIR_1[\Gamma_\infty(q)] \subseteq EXPIR_1[\Gamma(q)]$$

These inclusions hold in two-person games with independent private values, even if player 2 also has private information (see Koren (1992)). From Lemma 1, under the assumption of uniform punishments, the three sets coincide. In examples 3 and 4 below, this assumption does not hold. In example 3, the two inclusions are strict. In example 4, the first inclusion is strict but $INTIR_1[\Gamma_\infty(q)] = EXPIR_1[\Gamma(q)]$.

Example 3

Let $n = 2$, $\Theta_1 = \{h, l\}$, $|\Theta_2| = 1$: only player 1 has private information. Here, the prior probability distribution is fully described by the probability that player 1's type is h , which we still denote as $q \in [0, 1]$. Let $|A_1| = |A_2| = 2$ and the utility functions be described by

$$u_1(h, \cdot) = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \quad u_1(l, \cdot) = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}$$

$$u_2(\cdot) = \begin{pmatrix} 0 & 2 \\ 0 & 0 \end{pmatrix}$$

The assumption of uniform punishments is clearly not satisfied: player 2 must play right in order to hold player 1 of type h at his value level $v_1(h) = 0$ and must play left to hold him at $v_1(l) = 0$. Consider the probability distribution

$$\pi(h) = \pi(l) = \pi = \begin{pmatrix} \frac{1}{4} & \frac{1}{2} \\ 0 & \frac{1}{4} \end{pmatrix} \in \Delta(A_1 \times A_2)$$

Let us check that it defines an equilibrium of $\Gamma_\infty(q)$, for every $p \in (0, 1)$, namely that the associated payoffs, $x_1(h) = x_1(l) = \frac{1}{4}$, $x_2 = 1$, verify the above conditions (including (17)). Player 2's payoff $x_2 = 1$ is individually rational since the value of player 2's game is $v_2 = 0$. π is clearly incentive compatible since it is nonrevealing. According to (17), a vector payoff $(x_1(h), x_1(l))$ is individually rational for player 1 in $\Gamma_\infty(q)$ if and only if

$$\forall p \in [0, 1], \quad px_1(h) + (1-p)x_1(l) \geq val_1 \begin{pmatrix} p & 0 \\ 0 & 1-p \end{pmatrix} = p(1-p)$$

so that $(\frac{1}{4}, \frac{1}{4})$ is indeed individually rational for player 1 in $\Gamma_\infty(p)$, for every $p \in (0, 1)$. Hence $((\frac{1}{4}, \frac{1}{4}), 1) \in \mathcal{N}[\Gamma_\infty(q)]$ for every $q \in (0, 1)$. However, $((\frac{1}{4}, \frac{1}{4}), 1) \notin \mathcal{F}[\Gamma(q)]$ because $(\frac{1}{4}, \frac{1}{4})$ is not *interim* individually rational in the sense of (11): let $\tau = (\beta, 1 - \beta)$; $\max_{a_1} u_1(h, a_1, \tau) = \beta \leq \frac{1}{4}$ is incompatible with $\max_{a_1} u_1(l, a_1, \tau) = 1 - \beta \leq \frac{1}{4}$.

Consider now the probability distribution

$$\pi(h) = \pi(l) = \pi = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$$

π satisfies the equilibrium conditions of proposition 1 (namely (14), (15) and (16) above) but the vector payoff of player 1 is $(0, 0)$ and is not individually rational for player 1 in $\Gamma_\infty(q)$, namely does not satisfy (17). Hence π does not define an equilibrium of $\Gamma_\infty(q)$.

Example 3 illustrates that player 1 can benefit from not revealing his information to player 2, if player 2 intends to punish him. Of course, when uniform punishments are available, the revelation of information does not matter. ■

Example 4

The framework is the same as in example 3 but the utility functions are described by

$$u_1(h, \cdot) = \begin{pmatrix} 1 & 2 \\ 0 & -1 \end{pmatrix} \quad u_1(l, \cdot) = \begin{pmatrix} -1 & 0 \\ 2 & 1 \end{pmatrix}$$

$$u_2(\cdot) = \begin{pmatrix} 2 & 0 \\ 0 & 2 \end{pmatrix}$$

$v_1(h) = v_1(l) = 1$. As in the previous example, the assumption of uniform punishments is not satisfied. Let $p \in [0, 1]$.

$$\begin{aligned}
val_1 [pu_1(h, \cdot) + (1-p)u_1(l, \cdot)] &= val_1 \begin{pmatrix} 2p-1 & 2p \\ 2-2p & 1-2p \end{pmatrix} \\
&= 1-2p \quad \text{if } p \leq \frac{1}{4} \\
&= \frac{1}{2} \quad \text{if } \frac{1}{4} \leq p \leq \frac{3}{4} \\
&= 2p-1 \quad \text{if } p \geq \frac{3}{4}
\end{aligned}$$

This function is convex so that a vector payoff $(x_1(h), x_1(l))$ is individually rational for player 1 in the sense of (17) if and only if it is *ex post* individually rational (namely, (15): $x_1(h) \geq 1$ and $x_1(l) \geq 1$): $INTIR_1[\Gamma_\infty(q)] = EXPIR_1[\Gamma(q)]$. In particular, in this example, the equilibrium conditions in $\Gamma_\infty(q)$ are correctly described in proposition 1, namely by (14), (15) and (16).²¹

In spite of the previous property, proposition 3 fails. The probability distributions

$$\pi(h) = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \quad \pi(l) = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}$$

lead to an equilibrium in $\Gamma_\infty(q)$, with payoff $((1, 1), 2)$, but $(1, 1)$ is not *interim* individually rational for player 1 in the sense of (11): let $\tau = (\beta, 1-\beta)$; $\max_{a_1} u_1(h, a_1, \tau) = 2 - \beta \leq 1$ is incompatible with $\max_{a_1} u_1(l, a_1, \tau) = \beta + 1 \leq 1$. ■

In both examples 3 and 4, *interim* individual rationality takes a different form in the one-shot game and in the infinitely repeated game. In example 3, in order to defend himself, player 1 must play in a non-revealing way in the repeated game. In example 4, player 1 benefits from revealing his information to player 2.

The phenomena described in the previous examples were first identified in the study of zero-sum infinitely repeated games with incomplete information (see Aumann and Maschler (1995)).

²¹The simplification of the equilibrium conditions in the case of convex value functions (which give rise to a linear concavification) is acknowledged in Koren (1992), remark 4. A similar condition is considered in Forges (1988).

7 Appendix

7.1 Proof of proposition 1

7.1.1 Strategies and payoff functions

A strategy of player i in $\Gamma_\infty(q)$ is a sequence of mappings $\sigma_i = (\sigma_i^t)_{t \geq 0}$, $\sigma_i^t : \Theta_i \times A^{t-1} \rightarrow \Delta(A_i)$. The n -tuple of prior probability distributions $q = (q_i)_{1 \leq i \leq n}$ and an n -tuple of strategies $\sigma = (\sigma_i)_{1 \leq i \leq n}$ induce a probability distribution over $\Theta \times A^\mathbb{N}$, where $A^\mathbb{N}$ is the set of all infinite sequence of moves. We denote as $E_{q,\sigma}$ the corresponding expectation. Given $\mathbf{a} = (a^t)_{t \geq 0} \in A^\mathbb{N}$, let us define

$$\bar{U}_i^{T+1}(\theta_i, \mathbf{a}) = \frac{1}{T+1} \sum_{t=0}^T u_i(\theta_i, a^t) \quad \text{for every } i, \theta_i \text{ and } T = 0, 1, \dots$$

As in Hart (1985) (see also Forges (1992), Koren (1992), Shalev (1994)), we define the *interim* payoffs associated with an n -tuple of strategies σ as

$$U_i(\theta_i, \sigma) = \mathcal{L} \left[E_{q,\sigma}(\bar{U}_i^T(\theta_i, \tilde{\mathbf{a}}) \mid \theta_i) \right]$$

where \mathcal{L} is a Banach limit and $\tilde{\mathbf{a}}$ denotes the sequence of moves as a random variable.

7.1.2 Sufficient conditions for an equilibrium

Let us assume that the conditions (3) hold. Then we can construct an n -tuple of strategies $\sigma = (\sigma_i)_{1 \leq i \leq n}$ in $\Gamma_\infty(q)$ which achieve the *interim* payoffs $x_i(\theta_i)$ (namely, such that $x_i(\theta_i) = U_i(\theta_i, \sigma)$ for every i, θ_i) and which define a Nash equilibrium of $\Gamma_\infty(q)$. For every player i , σ_i is described as follows:

at the first stage ($t = 0$): choose a_i so as to reveal type θ_i (which is possible since $|A_i| \geq |\Theta_i|$)

at every stage $t \geq 1$: given the n -tuple of reported types θ' , play according to $\pi(\theta')$ if $\pi(\theta')$ was chosen at every previous stage; otherwise, play a punishment strategy in order to keep the first player j who did not follow $\pi(\theta')$ below his *ex post* individually rational level $v_j(\theta_j)$.

7.1.3 Necessary conditions for an equilibrium

Let us start with an arbitrary Nash equilibrium $\sigma = (\sigma_i)_{1 \leq i \leq n}$ in $\Gamma_\infty(q)$. Let $\sigma_i(\theta_i)$ be the associated strategy of player i of type θ_i , namely, $\sigma_i(\theta_i) = (\sigma_i^t(\theta_i))_{t \geq 0}$, with $\sigma_i^t(\theta_i) : A^{t-1} \rightarrow \Delta(A_i)$. Let $x_i(\theta_i) = U_i(\theta_i, \sigma)$ be the associated *interim* equilibrium payoff of player i of type θ_i . Let us show that the conditions (3) hold, namely, that the same payoffs can be achieved by a completely revealing equilibrium.

In order to get some intuition, let us assume that, at equilibrium, there is a finite, possibly very long, phase of information transmission (say, until stage t_0) and that afterwards (thus, at stages $t_0 + 1, t_0 + 2, \dots$), the players play independently of their types. Since σ is an equilibrium, player i of type θ_i cannot benefit from playing according to $\sigma_i(\theta'_i)$, with θ'_i possibly different from θ_i , until stage t_0 and then, from stage $t_0 + 1$ on, by either continuing to play $\sigma_i(\theta'_i)$ or just guaranteeing himself $v_i(\theta_i)$ (*i.e.*, by playing optimally in “his true one-shot game”, with payoffs $u_i(\theta_i, \cdot)$, at every stage $t_0 + 1, t_0 + 2, \dots$).²²

More precisely, the equilibrium strategies $\sigma_i(\theta_i)$ generate probability distributions $\mu_\sigma(\cdot \mid \theta_1, \dots, \theta_n)$ over the limit frequencies of moves, *i.e.*, over $\Delta(A)$ (see Hart (1985) or Koren (1992) for details). Together with the prior q , these probability distributions generate a probability distribution P_{q, μ_σ} over $\Theta \times \Delta(A)$ such that

$$x_i(\theta_i) = U_i(\theta_i, \sigma) = E_{q, \mu_\sigma}(u_i(\theta_i, \tilde{\pi}) \mid \theta_i) \quad \text{for every } i, \theta_i \quad (18)$$

where E_{q, μ_σ} is the expectation with respect to P_{q, μ_σ} and $\tilde{\pi}$ stands for the frequency of move as a random variable.²³

By considering the previous specific deviations of player i of type θ_i (namely, mimic type θ'_i and/or play optimally in the one-shot game), we obtain that

$$x_i(\theta_i) \geq E_{q, \mu_\sigma}(\max\{u_i(\theta_i, \tilde{\pi}), v_i(\theta_i)\} \mid \theta'_i) \quad \text{for every } i, \theta_i, \theta'_i \quad (19)$$

We can also rely on a variant of the revelation principle to see that (18) and (19) must be satisfied as soon as σ is an equilibrium. Let us imagine

²²Note that player i may reveal further information on his type by playing so as to guarantee himself $v_i(t_i)$. This typically happens out of equilibrium.

²³If information transmission ends up after finitely many stages t_0 , $\tilde{\pi}$ can be interpreted as the frequency of moves from stage $t_0 + 1$ on.

that a fully reliable mediator asks the players to report their types and then given the n -tuple of reported types $\theta' \in \Theta$, chooses a frequency of moves $\pi \in \Delta(A)$ according to $\mu_{\sigma}(\cdot \mid \theta')$ and recommends π to all players²⁴. In other words, when the players report $\theta' = (\theta'_i)_{1 \leq i \leq n}$, the mediator selects π exactly as the players themselves do at the equilibrium σ . (18) says that by telling the truth and following the recommendation of the mediator, the players get the same *interim* payoff as by playing σ . (19) says that if players $j \neq i$ tell the truth to the mediator, follow the recommendation π as long as every player follows π and punish any deviator at his *ex post* minmax level, then player i of type θ_i cannot benefit from reporting type θ'_i to the mediator and/or not following π .

Conditions (18) and (19) differ from (3) in two respects. (18) and (19) involve (type dependent) probability distributions over $\Delta(A)$, while (3) is formulated in terms of deterministic distributions $\pi(\theta)$, $\theta \in \Theta$. Moreover, in (19), the probability distribution μ_{σ} is not necessarily completely revealing²⁵.

By construction, and recalling that types are independent of each other, for any function f over $\Delta(A)$, the probability $P_{q, \mu_{\sigma}}$ satisfies

$$E_{q, \mu_{\sigma}}(f(\tilde{\pi}) \mid \theta_i) = \sum_{\theta_{-i}} q_{-i}(\theta_{-i}) E_{\mu_{\sigma}}(f(\tilde{\pi}) \mid \theta_i, \theta_{-i}) \quad \text{for every } i, \theta_i$$

Hence, for every i, θ_i , (18) can be rewritten as

$$x_i(\theta_i) = \sum_{\theta_{-i}} q_{-i}(\theta_{-i}) E_{\mu_{\sigma}}(u_i(\theta_i, \tilde{\pi}) \mid \theta_i, \theta_{-i})$$

Recalling that $u_i(\theta_i, \cdot)$ is linear, we get

$$x_i(\theta_i) = \sum_{\theta_{-i}} q_{-i}(\theta_{-i}) u_i(\theta_i, E_{\mu_{\sigma}}(\tilde{\pi} \mid \theta_i, \theta_{-i}))$$

which is the first part of (3) if we set $\pi(\theta) = E_{\mu_{\sigma}}(\tilde{\pi} \mid \theta)$.

²⁴As in the standard proof of the Folk theorem under complete information, we interpret a distribution of moves π as a deterministic sequence of moves (in A) which achieves the frequency of π . This interpretation is straightforward if the components of π are rational (in \mathbb{Q}).

²⁵The above reliable mediator selects π as a random function of the players' reported types but does not reveal these reported types.

By proceeding similarly and using in addition that “max” is convex, for every i, θ_i, θ'_i , (19) can be rewritten as

$$\begin{aligned}
x_i(\theta_i) &\geq \sum_{\theta_{-i}} q_{-i}(\theta_{-i}) E_{\mu_\sigma} (\max \{u_i(\theta_i, \tilde{\pi}), v_i(\theta_i)\} \mid \theta'_i, \theta_{-i}) \\
&\geq \sum_{\theta_{-i}} q_{-i}(\theta_{-i}) \max \{E_{\mu_\sigma} (u_i(\theta_i, \tilde{\pi}) \mid \theta'_i, \theta_{-i}), v_i(\theta_i)\} \\
&\geq \sum_{\theta_{-i}} q_{-i}(\theta_{-i}) \max \{u_i(\theta_i, E_{\mu_\sigma}(\tilde{\pi} \mid \theta'_i, \theta_{-i})), v_i(\theta_i)\} \\
&\geq \sum_{\theta_{-i}} q_{-i}(\theta_{-i}) \max \{u_i(\theta_i, \pi(\theta'_i, \theta_{-i})), v_i(\theta_i)\}
\end{aligned}$$

The last expression is the inequality in (3). ■

7.2 Computation of example 1

Recall from section 3.2. that the *ex post* individual rationality conditions are (7), namely,

$$\gamma_d \leq k\gamma, \beta_d \leq k\alpha \text{ and } \rho \geq (1 - \frac{1}{z})(1 - \beta_d).$$

Since $\rho \leq 1 - \beta_d$, we can set $\rho = (1 - \epsilon)(1 - \beta_d)$ with $\epsilon \leq \frac{1}{z}$.

Furthermore, $\alpha \leq (1 - \beta_d) - \rho = \epsilon(1 - \beta_d)$. Hence, $\beta_d \leq k\alpha \leq k\epsilon(1 - \beta_d)$, so that $\beta_d \leq \frac{k\epsilon}{1+k\epsilon}$.

A further equilibrium condition is (8), namely,

$$p(k\gamma - \gamma_d) + (1 - p)(k\alpha - \beta_d) \geq p(k\rho - \beta_d).$$

Since $k\gamma - \gamma_d \leq \frac{k}{2}$ and $k\alpha - \beta_d \leq k\epsilon - (k\epsilon + 1)\beta_d$, this condition implies

$$p \left[(2 + k)\beta_d - \frac{k}{2} \right] \geq (1 + k\epsilon)\beta_d - k\epsilon.$$

In the left hand side, $(2 + k)\beta_d \leq \frac{k}{2}$, because $\beta_d \leq \frac{k\epsilon}{1+k\epsilon}$ and $\epsilon \leq \frac{1}{z} < \frac{1}{k+4}$. We thus get the following condition on p

$$p \leq \frac{(1 + k\epsilon)\beta_d - k\epsilon}{(2 + k)\beta_d - \frac{k}{2}} \leq 2\epsilon < \frac{2}{k + 4},$$

where the second inequality comes from the fact that the expression is decreasing in β_d if $\epsilon < \frac{1}{k+4}$. ■

References

- [1] Atakan, E. and M. Ekmekci (2012), “Reputation in long-run relationships”, *Review of Economic Studies* 79, 451-480.
- [2] Atakan, E. and M. Ekmekci (2013), “A two-sided reputation result with long-run players”, *Journal of Economic Theory* 148, 376–392.
- [3] Aumann, R. J. and M. Maschler (1995), *Repeated Games of Incomplete Information*, Cambridge: M.I.T. Press.
- [4] Aumann, R. J., M. Maschler and R. Stearns (1968), “Repeated games with incomplete information: an approach to the nonzero sum case”, Reports to the U.S. Arms Control and Disarmament Agency, ST-143, Chapter IV, 117-216.
- [5] Bergin, J., (1989), “A characterization of sequential equilibrium strategies in infinitely repeated incomplete information games,” *Journal of Economic Theory* 47, 51-65.
- [6] Blackwell, D. (1956), “An analog of the minimax theorem for vector payoffs”, *Pacific Journal of Mathematics* 10, 1-8.
- [7] Chan, J., (2000), “On the non-existence of reputation effects in two-person infinitely repeated games”, mimeo, Johns Hopkins University.
- [8] Cripps, M. and J. Thomas (1995), “Reputation and commitment in two-person repeated games without discounting”, *Econometrica* 63, 1401-1419.
- [9] Cripps, M. and J. Thomas (1997), “Reputation and perfection in repeated common interest games”, *Games and Economic Behavior* 18, 141-158.
- [10] Cripps, M. and J. Thomas (2003), “Some asymptotic results in discounted repeated games of one-sided incomplete information”, *Mathematics of Operations Research* 28, 433-462.
- [11] Cripps, M., J. Thomas and K. Schmidt (1996), “Reputation in perturbed repeated games,” *Journal of Economic Theory* 69, 387-410.

- [12] Cripps, M., E. Dekel and W. Pesendorfer (2005), “Reputation with equal discounting in repeated games with strictly conflicting interests”, *Journal of Economic Theory* 121, 259-272.
- [13] Forges, F. (1988), “Communication equilibria in repeated games with incomplete information”, *Mathematics of Operations Research* 13, 191-231.
- [14] Forges, F. (1992), “Repeated games of incomplete information : non-zero-sum”, in: R. Aumann et S. Hart, *Handbook of Game Theory with Economic Applications*, Elsevier Science Publishers, North-Holland, chapter 6, 155-177.
- [15] Forges, F. (2013), “A folk theorem for Bayesian games with commitment”, *Games and Economic Behavior* 78, 64–71.
- [16] Fudenberg, D. and E. Maskin (1986), “The folk theorem in repeated games with discounting or with incomplete information”, *Econometrica* 54, 533-554.
- [17] Fudenberg, D. and J. Tirole (1991), *Game theory*, M.I.T. Press.
- [18] Hart S. (1985), “Nonzero-sum two-person repeated games with incomplete information”, *Mathematics of Operations Research* 10, 117–153.
- [19] Hörner, J. and S. Lovo (2009), “Belief-free equilibria in games with incomplete information”, *Econometrica* 77, 453–487.
- [20] Hörner, J., S. Lovo and T. Tomala (2011), “Belief-free equilibria in games with incomplete information: characterization and existence”, *Journal of Economic Theory* 146, 1770–1795.
- [21] Israeli, E. (1999), “Sowing doubt optimally in two-person repeated games,” *Games and Economic Behavior* 28, 203-216.
- [22] Koren, G. (1992), “Two-person repeated games where players know their own payoffs”, mimeo, Courant Institute of Mathematical Sciences, New York University.
- [23] Kreps, D., P. Milgrom, J. Roberts and R. Wilson (1982), “Rational cooperation in the finitely repeated prisoners dilemma”, *Journal of Economic Theory* 27, 245-252.

- [24] Kalai, A., E. Kalai, E. Lehrer and D. Samet (2010), “A commitment folk theorem”, *Games and Economic Behavior* 69, 127-137.
- [25] Mailath, G. and L. Samuelson (2006), *Repeated Games and Reputations*, *Oxford University Press*.
- [26] Mertens, J.-F and S. Zamir (1971), “The value of two-person zero-sum repeated games with lack of information on both sides”, *International Journal of Game Theory* 1, 39-64.
- [27] Myerson, R. (1991), *Game theory: analysis of conflict*, Harvard University Press.
- [28] Palfrey, T. and H. Rosenthal (1994), “Repeated play, cooperation and coordination: an experimental study”, *Review of Economic Studies* 61, 545-565.
- [29] Peski, M. (2008), “Repeated games with incomplete information on one side”, *Theoretical Economics* 3, 29-84.
- [30] Peski, M. (2012), “Repeated games with incomplete information and discounting”, *mimeo*, University of Toronto.
- [31] Peters, M. and B. Szentes (2012), “Definable and contractible contracts”, *Econometrica* 80, 363-411.
- [32] Shalev, J. (1994), “Nonzero-sum two-person repeated games with incomplete information and known-own payoffs”, *Games and Economic Behavior* 7, 246-259.
- [33] Schmidt, K. (1993), “Reputation and equilibrium characterization in repeated games with conflicting interests”, *Econometrica* 61, 325-351.
- [34] Sorin, S. (1999), “Merging, reputation, and repeated games with incomplete information”, *Games and Economic Behavior* 29, 274-308.
- [35] Wiseman, T. (2012), “A partial folk theorem for games with private learning”, *Theoretical Economics* 7, 217-239.