



HAL
open science

Folk theorems for Bayesian (public good) games

Françoise Forges

► **To cite this version:**

Françoise Forges. Folk theorems for Bayesian (public good) games. Thirteenth annual conference (PET12) of the Association for Public Economic Theory (APET), Jun 2012, Taipei, China. pp.30. hal-02447604

HAL Id: hal-02447604

<https://hal.science/hal-02447604>

Submitted on 21 Jan 2020

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.

Folk theorems for Bayesian (public good) games*

Françoise Forges[†]

Abstract

Two characterization results are behind the intimate relationship between “repetition” and “cooperation” under complete information: the standard Folk theorem for infinitely repeated games and the “commitment Folk theorem” for one-shot games. We propose extensions of the previous characterization results in Bayesian games, with independent private values, which satisfy a further property, “uniform punishment strategies”. Public good games fall in this class. We show that the Nash equilibria of the Bayesian infinitely repeated game are payoff equivalent to separating (*i.e.*, completely revealing) equilibria and can be achieved as *interim* cooperative solutions of the Bayesian game. We also show that the reverse of the latter result is not true: unlike the set of *interim* cooperative solutions of the Bayesian game, the set of Nash equilibrium payoffs of the infinitely repeated game can be empty.

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

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

*Notes for a lecture at PET 12 (Taipei, June 13, 2012). The list of references is likely incomplete. Conversations with Gorkem Celik, Gaël Giraud, Vincent Iehlé, Frédéric Koessler, Antoine Salomon and Péter Vida are gratefully acknowledged. Bernard Lebrun made a number of precise and helpful comments on an earlier draft.

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

1 Introduction

Under complete information, the standard Folk theorem states that the Nash equilibrium payoffs of an infinitely repeated game coincide with the feasible and individually rational payoffs of the one-shot game. Recently, Kalai *et al.* (2010) established that the latter set can be interpreted as the set of cooperative solutions of the one-shot game, namely, as the set of all equilibrium payoffs that the players can achieve when they can voluntarily commit themselves. Kalai *et al.* (2010) refer to this result as the “commitment Folk theorem”, for two reasons: it is stated in terms of the set of feasible and individually rational payoffs and it has been known for years, under more or less precise versions (see, *e.g.*, Tennenholtz (1984) and Myerson (1991)).

Put together, the previous two Folk theorems state that the noncooperative solutions of the infinitely repeated game *coincide* with the cooperative solutions of the one-shot game. They thus yield more than a precise content to the idea that “repetition enables cooperation”: they also say that “repetition does as well as cooperation”.

In these notes, I gather some results which shed some light on the possible extension of the latter Folk theorems in Bayesian games, *i.e.*, when the players do not share the same information over the fundamentals of the game.

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 a very tractable way: they are payoff equivalent to *completely revealing* (also called “sepa-

rating”) equilibria.¹ Once such a tractable characterization is available, one can ask whether 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, Koren (1992)’s result holds true, namely, the Nash equilibrium of the (undiscounted) infinitely repeated game are all payoff equivalent to *completely revealing* equilibria.

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.

On the other hand, in Forges (2011), I propose an extension of Kalai *et al.* (2010)’s “commitment Folk theorem” to arbitrary n -person Bayesian games. More precisely, I consider all meta-Bayesian games in which, at the *interim* stage, the players give their instructions on how to play the original game. As in Kalai *et al.* (2010), the meta-games are formulated in order to avoid circular reasoning. I show that the set of all Nash equilibrium payoffs of all such meta-games coincides with the set of incentive compatible, *interim* individually rational, feasible payoffs identified by Myerson (1991, chapter 6). In particular, all payoffs in the latter set can be achieved as the Nash equilibrium payoffs of a single meta-Bayesian game. This result is recalled below as proposition 2.

Equipped with the characterizations of propositions 1 and 2, I show that, for any n -person Bayesian game with independent, private values and uni-

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 non-zero sum infinitely repeated games with incomplete information.

form punishments, the set of Nash equilibrium payoffs of the infinitely repeated game is contained in the set of *interim* cooperative solutions of the Bayesian game (proposition 3). In other words, the repetition of the game enables the players to cooperate, as in the case of 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 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)). 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 characterizations show 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 3. As explained in section 5, 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 3 is no longer true, while Koren’s characterization still holds in the two-person case. Section 5 further discusses our underlying assumptions, in particular, the undiscounted payoffs.

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 T_i
- a probability distribution q_i over T_i
- a finite set of actions A_i , with $|A_i| \geq |T_i|$
- a utility function $u_i : T_i \times A \rightarrow \mathbb{R}$, 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 $q = (q_i)_{1 \leq i \leq n}$.² Without loss of generality, we assume that $q_i(t_i) > 0$ for every $t_i \in T_i$. The interpretation is that types t_i , $i = 1, \dots, n$, are first chosen in T , independently of each other, according to q . At the *interim* stage, player i is only informed of his own type t_i . The players then choose simultaneously an action.

For any finite set E , let us denote as $\Delta(E)$ the set of probability distributions over E . A mixed strategy³ of player i in $\Gamma(q)$ is a mapping from T_i to $\Delta(A_i)$. Similarly, a correlated strategy for players $j \neq i$ is a mapping from $T_{-i} = \prod_{j \neq i} T_j$ 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 $i = 1, \dots, n$, $t_i \in T_i$, $\delta \in \Delta(A)$,

$$u_i(t_i, \delta) = \sum_a \delta(a) u_i(t_i, a)$$

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

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

²We 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.

³More correctly, “behavior strategy”.

For every player i , $i = 1, \dots, n$, and $t_i \in T_i$, let $v_i(t_i)$ be the value⁴ of the (complete information, zero-sum) game in which player i maximizes the payoff $u_i(t_i, \cdot)$, namely

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

Observe that, in the previous expression, the probability distribution τ_{-i} achieving the “min” possibly depends on t_i , which is fixed in the underlying optimization problem. $v_i(t_i)$ 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 t_i .

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

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

When (2) holds, τ_{-i} defines independent⁵ punishment strategies which enable players $j \neq i$ to punish player i uniformly, *i.e.*, whatever his type t_i is, but even more, to keep player i ’s payoff below his *ex post* individual rationality level.

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)).⁶ 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 5.

⁴If 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.

⁵Independent punishment strategies are important for proposition 1.

⁶As a slight weakening, $v_i(t_i)$ could just be defined as

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

2.2 Application: contribution to a public good

The private information of every player i , $i = 1, \dots, n$, is the value t_i that he attributes to his endowment of a single unit of the private good. The private endowment values t_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

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

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 (t_1, t_2) is

	c	d
c	1, 1	1, 1 + t_2
d	1 + t_1 , 1	t_1, t_2

where we always assume $t_i \geq 0$ but can have $t_i < 1$ or $t_i > 1$. Fudenberg and Tirole (1991, example 6.1, p. 211) propose the following interpretation: player 1 and player 2 belong to a group (say, the “economists” of some university) 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 t_i whatever the other players choose. By playing τ_{-i} , the players

$j \neq i$ guarantee that player i 's payoff does not exceed t_i . Hence, if $m > 1$, $v_i(t_i) = t_i$ and τ_{-i} is a uniform punishment strategy. Assume now that $m = 1$. Again, by playing d , player i guarantees himself t_i ; but now, by playing c , player i guarantees himself 1. Hence, by playing according to his type, player i can guarantee himself $\max\{t_i, 1\}$. By playing τ_{-i} , the players $j \neq i$ guarantee that player i 's payoff does not exceed $\max\{t_i, 1\}$. Hence, if $m = 1$, $v_i(t_i) = \max\{t_i, 1\}$ and τ_{-i} is a uniform punishment strategy.

3 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 the infinitely repeated version of the previous game, 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:

- stage 0: the types t_i , $i = 1, \dots, n$, are chosen in $T = \prod_{1 \leq i \leq n} T_i$ independently of each other, according to q . Player i is only informed of his own type t_i .
- stage k ($k = 1, 2, \dots$): every player i chooses an action in A_i . The choices are made simultaneously and revealed publicly right after stage k .

Payoffs in $\Gamma_\infty(q)$ are evaluated as (Banach) limits of arithmetic averages (see Hart (1985), Forges (1992)), thus without discounting. As the other assumptions, this one is discussed in section 5.

3.1 Characterization of Nash equilibrium payoffs

By proceeding as in Koren (1992), we can characterize the Nash equilibrium payoffs of $\Gamma_\infty(q)$. We write $q_{-i}(t_{-i})$ for $\prod_{j \neq i} q_j(t_j)$.

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

$$\begin{aligned} x_i(t_i) &= \sum_{t_{-i} \in T_{-i}} q_{-i}(t_{-i}) u_i(t_i, \delta(t_i, t_{-i})) \\ &\geq \sum_{t_{-i} \in T_{-i}} q_{-i}(t_{-i}) \max \{u_i(t_i, \delta(r_i, t_{-i})), v_i(t_i)\} \end{aligned} \quad (3)$$

In the case of *complete information*, namely if the prior probability distribution q is degenerated, 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 $\delta \in \Delta(A)$) 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 1, every player i truthfully reveals his type t_i
- at every stage $k \geq 2$, given the reported types $r = (r_i)_{1 \leq i \leq n}$, every player i plays according to $\delta(r) \in \Delta(A)$ provided that $\delta(r)$ has been followed at every previous stage $2, \dots, k-1$. Otherwise, if player i does not follow $\delta(r)$ at some stage $k \geq 2$, players $j \neq i$ punish player i by using uniform independent punishment strategies τ_{-i} holding player i at $v_i(t_i)$ at every stage $k+1, k+2, \dots$ whatever his type t_i and action are.

The nondeviation condition (3) expresses that, assuming that players $j \neq i$ follow the equilibrium strategies, player i of type t_i can report a type r_i possibly different from t_i . At the end of stage 1, player i learns the true types t_{-i} of the other players and can then either follow $\delta(r_i, t_{-i})$ or not. In

the former case, he fully mimics the equilibrium strategy of type r_i . In the latter case, he is punished at the level $v_i(t_i)$.⁷

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 $r_i = t_i$, (3) is equivalent to

$$\text{For every } i \text{ and } t = (t_i, t_{-i}) \in T : \quad u_i(t_i, \delta(t)) \geq v_i(t_i)$$

which implies that

$$\text{For every } i \text{ and } t_i \in T_i : \quad x_i(t_i) \geq v_i(t_i) \tag{4}$$

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 symmetric two-person game $PG(2, 1, q)$

		$t_2 = \omega$		$t_2 = z$	
		c	d	c	d
$t_1 = \omega$	c	1, 1	1, 1 + ω	1, 1	1, 1 + z
	d	1 + ω , 1	ω , ω	1 + ω , 1	ω , z
$t_1 = z$	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: $T_i = \{\omega, z\}$, $i = 1, 2$. We assume that $0 < \omega < 1$ and $z > 2$: ω represents a “normal” type, who values the public

⁷Koren (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 as (3) and the extension to n players.

good more than his initial endowment, while z represents a “greedy” type. We also assume 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 $\omega = \frac{2}{3}$ and $z = 3$. Consider the following distributions, which yield feasible, ex post individually rational payoffs:

		$t_2 = 2/3$ (<i>prob.</i> p_2)		$t_2 = 3$ (<i>prob.</i> $1 - p_2$)	
		c	d	c	d
$t_1 = 2/3$ (<i>prob.</i> p_1)	c	0	1/2	0	7/10
	d	1/2	0	3/10	0
$t_1 = 3$ (<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 $\Gamma_\infty(p)$ of example 0, with the beliefs $q_i = (p, 1-p)$, $i = 1, 2$, with $0 < p < 1$ and let us set $k = \frac{\omega}{1-\omega}$. We will show that

$$\text{If } z > k + 4 \text{ and } p > \frac{2}{k + 4}, \quad \mathcal{N}[\Gamma_\infty(p)] = \emptyset \quad (5)$$

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 p 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.⁸

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 $\delta(t)$ over $\{c, d\} \times \{c, d\}$, one for every pair of types t . If $t_1 = t_2 = z$, *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* equilibrium⁹. We thus focus on $\delta(t)$ ’s of the form:

		$t_2 = \omega$		$t_2 = z$	
		c	d	c	d
$t_1 = \omega$	c	δ_c	δ	β_c	ρ
	d	δ	δ_d	α	β_d
$t_1 = z$	c	β_c	α	0	0
	d	ρ	β_d	0	1

where all parameters are nonnegative and $2\delta + \delta_c + \delta_d = 1$, $\alpha + \rho + \beta_c + \beta_d = 1$.

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

$$\delta_d \leq k\delta, \beta_d \leq k\alpha \text{ and } \rho \geq \left(1 - \frac{1}{z}\right)(1 - \beta_d) \quad (6)$$

In the right hand sight of the equilibrium condition (3) for $t_1 = \omega$ and $r_1 = z$,

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

⁸Koren (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.

⁹If $\Gamma_\infty(p)$ has an equilibrium, there exist probability distributions $\delta(t)$, $t \in T$, over $\{c, d\} \times \{c, d\}$ satisfying (3). If $\delta(t)$, $t \in T$, satisfies (3), the probability distributions $\delta'(t)$, $t \in T$, in which player 1 and player 2 are permuted, also satisfy (3). The same holds for the symmetric distributions $\frac{1}{2}(\delta(t) + \delta'(t))$, $t \in T$, thanks to the linearity of u and the convexity of “max”.

namely, $\beta_d \leq k\rho$ because, from (6), $\beta_d \leq k\alpha$ and $\alpha \leq 1 - \beta_d - \rho \leq \frac{1}{z-1}\rho \leq \rho$. We can thus write (3) for $t_1 = \omega$ and $r_1 = z$ as

$$p(k\delta - \delta_d) + (1-p)(k\alpha - \beta_d) \geq p(k\rho - \beta_d) \quad (7)$$

This condition is not compatible with (6) if p is close enough to 1. In order to get some intuition for this, let us try $\delta_d = \beta_d = 0$, *i.e.*, an *ex post* efficient equilibrium. (6) reduces to $\rho \geq 1 - \frac{1}{z}$. (7) is $p\delta + (1-p)\alpha \geq p\rho$. Since $\delta \leq \frac{1}{2}$ and $\alpha \leq 1 - \rho$, (7) implies that $p \leq \frac{2}{z}(1 - \rho) \leq \frac{2}{z}$, 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 (5) for arbitrary δ_d, β_d satisfying (6).■

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 $t_1 = z$ and $r_1 = \omega$ 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 provide an answer to the question raised in the introduction: can we interpret the characterization in proposition 1 in the same way as under complete information, namely, as a precise relationship between “repetition” and “cooperation”?

We first give a description of the set of feasible, incentive compatible and *interim* individually rational payoffs in the (one-shot) Bayesian game $\Gamma(q)$. We denote this set as $\mathcal{F}[\Gamma(q)]$. We use Myerson (1991, sections 6.4 and 6.6)’s definitions but obtain simpler expressions in our framework of independent private values. Proposition 2 states that $\mathcal{F}[\Gamma(q)]$ can be interpreted as the set of *interim* cooperative solutions of $\Gamma(q)$. Then we check whether $\mathcal{N}[\Gamma_\infty(q)]$, the set of Nash equilibrium payoffs of the infinitely repeated game $\Gamma_\infty(q)$, coincides with $\mathcal{F}[\Gamma(q)]$.

A payoff $x = (x_i)_{1 \leq i \leq n} = ((x_i(t_i))_{t_i \in T_i})_{1 \leq i \leq n}$ is *feasible* in $\Gamma(q)$ if there exists a correlated strategy $\delta(t) \in \Delta(A)$, $t \in T$, achieving x , namely

$$x_i(t_i) = \sum_{t_{-i}} q_{-i}(t_{-i}) u_i(t_i, \delta(t_i, t_{-i})) \quad i = 1, \dots, n, \quad t_i \in T_i \quad (8)$$

A feasible payoff x achieved through δ (as in (8)) is *incentive compatible* if

$$x_i(t_i) \geq \sum_{t_{-i}} q_{-i}(t_{-i}) u_i(t_i, \delta(r_i, t_{-i})) \quad \text{for every } i, t_i, r_i \in T_i \quad (9)$$

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 that¹⁰

$$x_i(t_i) \geq \max_{a_i \in A_i} u_i(t_i, a_i, \tau_{-i}) \quad \text{for every } t_i \in T_i \quad (10)$$

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 $x_i(t_i) \geq w_i(t_i)$, $t_i \in T_i$, for some well-defined individually rational level $w_i(t_i)$). By contrast, *ex post* individually rational payoffs are described by a “corner set”, since $(v_i(t_i))_{t_i \in T_i}$ is defined without ambiguity by (1).

¹⁰Literally, 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, (10) is an equivalent formulation, since $u_i(t_i, \cdot)$ is linear.

The set $\mathcal{F}[\Gamma(q)]$ is formally defined as the set of payoffs satisfying (8), (9) and (10). $\mathcal{F}[\Gamma(q)]$ contains the set of Nash equilibrium payoffs of $\Gamma(q)$ and is thus not empty.

Proposition 2 $\mathcal{F}[\Gamma(q)]$ is the set of all “interim cooperative solutions” of $\Gamma(q)$, i.e., the set of all payoffs that can be achieved when the players can commit themselves at the interim stage of $\Gamma(q)$, in any voluntary contract game extending $\Gamma(q)$.

This result extends Kalai *et al.* (2010)’s commitment Folk theorem and is established in Forges (2011). It holds without independent private values, even if the previous definitions (8), (9) and (10) take this assumption into account. Uniform punishment strategies are not used.

Proposition 2 can be decomposed into two parts. To see this, let us consider all meta-Bayesian games which extend $\Gamma(q)$ by allowing the players to voluntarily commit themselves by giving possibly conditional instructions at the *interim* stage. Proposition 2 first states that the union, over all these meta-Bayesian games, of the corresponding sets of all Nash equilibrium payoffs, is included in $\mathcal{F}[\Gamma(q)]$. Proposition 2 further says that the reverse inclusion also holds. To prove the latter property, Forges (2011) constructs a single voluntary contract game extending $\Gamma(q)$ in which *all* payoffs of $\mathcal{F}[\Gamma(q)]$ can be achieved as Nash equilibrium payoffs.

Peters and Szentes (2012) argue that the set of solutions that the players can achieve by signing a contract at the *interim* stage of a Bayesian game is in general *smaller* than $\mathcal{F}[\Gamma(q)]$. They insist on the fact that natural contracts should involve two stages (loosely speaking, the players first publicly announce the commitments that they are ready to sign before signing them). However, under an assumption which is similar to our uniform punishment strategies, Peters and Szentes (2012) recover a result similar to proposition 2 (see their proposition 1).

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, (10)), x is ex post individually rational (namely, (4)): $INTIR_i[\Gamma(q)] \subseteq EXPIR_i[\Gamma(q)]$ 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 t_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 (10).

Proposition 3 *Let $\Gamma(q)$ be a Bayesian game with independent private values in which uniform punishment strategies are available: $\mathcal{N}[\Gamma_\infty(q)] \subseteq \mathcal{F}[\Gamma(q)]$.*

Proof: The proposition readily follows from the characterizations of $\mathcal{N}[\Gamma_\infty(q)]$ (in proposition 1) and $\mathcal{F}[\Gamma(q)]$ ((8), (9) and (10) above): the equality in (3) is (8), the inequality in (3) implies (9) and (4), which in turn implies (10) 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 inclusion¹¹.

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: $T_i = \{t, b\}$, $i = 1, 2$, 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 t (namely, $q_i = (p_i, 1 - p_i)$). Payoffs are described by the following matrices:

		$t_2 = t$		$t_2 = b$	
		c	d	c	d
$t_1 = t$	c	3, 1	0, 0	3, 1	0, 3
	d	0, 0	1, 3	0, 1	1, 3
$t_1 = b$	c	3, 1	3, 0	3, 1	3, 3
	d	1, 0	1, 3	1, 1	1, 3

When $t_1 = t$, player 1 prefers c to d , but also prefers to make the same choice as the other player. When $t_1 = b$, player 1 just prefers c to d , independently

¹¹For 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.

of the choice of the other player. The preferences of player 2 are similar. In this game, $v_i(t) = \frac{3}{4}$, $v_i(b) = 3$, $i = 1, 2$. 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 $\delta(t)$, $t \in T$, defined by

		$t_2 = t$		$t_2 = b$	
		c	d	c	d
$t_1 = t$	c	$\frac{1}{2}$	0	0	0
	d	0	$\frac{1}{2}$	0	1
$t_1 = b$	c	1	0	0	1
	d	0	0	0	0

(11)

It is easily checked that $\delta(t)$ satisfies (9) and (10), 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 $\mathcal{N}[\Gamma_\infty(q)]$, $\delta(t)$ must satisfy (3); in particular, player 1 of type $t_1 = t$ cannot gain by pretending to be of type $r_1 = b$, 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 t but pretends to be of type b at the first stage of $\Gamma_\infty(q)$. Then he learns player 2's type t_2 and faces $\delta(b, t_2)$. If $t_2 = t$, player 1 gets the best payoff 3 by playing according to $\delta(b, t)$. However, if $t_2 = b$, player 1 gets 0 by playing according to $\delta(b, b)$. In this case, he should not play according to $\delta(b, t)$ 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 $\delta(t)$ 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

		$t_2 = t$		$t_2 = b$	
		c	d	c	d
$t_1 = t$	c	0	0	0	0
	d	0	1	0	1
$t_1 = b$	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 $q \in \Delta(T)$ for which $\delta(t)$ defined by (11) 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 (11).■

5 Role of the assumptions

5.1 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)). As already mentioned, independent private values are not necessary in proposition 2.

5.2 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 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 3 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 3 still fails.

In both examples 3 and 4, there are two players and only player 1 has private information ($|T_2| = 1$, $A = A_1 \times A_2$), so that the conditions in proposition 1 reduce to: there exists $\delta(t_1) \in \Delta(A)$, $t_1 \in T_1$, such that, for

player 1,

$$\begin{aligned} x_1(t_1) &= u_1(t_1, \delta(t_1)) \\ &\geq u_1(t_1, \delta(r_1)) \quad \forall t_1, r_1 \in T_1 \quad \text{i.e., incentive compatibility} \quad (12) \end{aligned}$$

$$\geq v_1(t_1) \quad \forall t_1 \in T_1 \quad \text{i.e., ex post individual rationality} \quad (13)$$

and, for player 2,

$$x_2 = u_2 \left(\sum_{t_1 \in T_1} q_1(t_1) \delta(t_1) \right) \geq v_2 \quad \text{i.e., individual rationality} \quad (14)$$

As shown by Hart (1985), in order to characterize the equilibrium payoffs of $\Gamma_\infty(q)$, *ex post* individual rationality (namely, (4) or (13) 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.¹² Let $val_1[u]$ denote the value to player 1 of the one-shot game with payoff function u .

Definition A vector payoff $x_1 = (x_1(t_1))_{t_1 \in T_1}$ is individually rational for player 1 in the *infinitely repeated game* $\Gamma_\infty(q)$ if and only if

$$\forall p_1 \in \Delta(T_1), \quad \sum_{t_1} p_1(t_1) x_1(t_1) \geq val_1 \left[\sum_{t_1} p_1(t_1) u_1(t_1, \cdot) \right] \quad (15)$$

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 (15) 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 $x_1(t_1)$ when he is of type t_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

¹²The same condition holds as well in two-person games with independent private values (see Koren (1992)).

consequence of Blackwell (1956)'s characterization, (10) implies (15). Furthermore, (15) holds in particular at $p_1(t_1) = 1$ for every t_1 , so that it implies *ex post* individual rationality (*i.e.*, (13)). 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 $n = 2$, $T_1 = \{h, l\}$, $|T_2| = 1$: 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

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

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 (15)). Player 2's payoff $x_2 = 1$ is individually rational since the value of player 2's game is $v_2 = 0$. δ is clearly incentive compatible since it is nonrevealing. According to (15), 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 (10): 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

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

δ satisfies the equilibrium conditions of proposition 1 (namely (12), (13) and (14) 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 (15). Hence δ 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 (15) if and only if it is *ex post* individually rational (namely, (13): $x_1(h) \geq 1$ 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 (12), (13) and (14).¹³

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

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

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 (10): 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)).

5.3 Undiscounted payoffs

Throughout this paper, as in Hart (1985) and Koren (1992), we assumed 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 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.

¹³The 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).

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). 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 briefly mentioned in the introduction, many papers devoted to reputation effects deal with repeated games with discounting. In addition, a number of rather recent papers study discounted repeated games with specific forms of incomplete information. The papers that look most relevant to these notes are Peski (2008) and Peski (2012). Peski (2008) characterizes the 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 are much more complex than Shalev (1994)'s ones but show that equilibria keep an elegant structure. Peski (2012) extends these results to the class of discounted repeated games with known-own payoffs which satisfy an "open thread assumption", namely, in which there exists an open set of belief-free equilibria in the sense of Horner and Lovo (2009).

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 in these notes. 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 these notes and in Peski (2008). Peski (2008) concludes by saying that the problem of characterizing equilibrium payoffs in discounted repeated games with multisided incomplete information is open. As mentioned above, Peski (2012) gives an answer in a class of such games.¹⁴

¹⁴The first versions of these notes were deliberately written without any attempt to

6 Appendix

6.1 Proof of proposition 1

6.1.1 Strategies and payoff functions

A strategy of player i in $\Gamma_\infty(q)$ is a sequence of mappings $\sigma_i = (\sigma_i^k)_{k \geq 1}$, $\sigma_i^k : T_i \times A^{k-1} \rightarrow \Delta(A_i)$. 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 $T \times A^\mathbb{N}$, where $A^\mathbb{N}$ is the set of all infinite sequence of moves. We denote as $E_{q,\sigma}$ the corresponding expectation. Given $\mathbf{a} = (a^k)_{k \geq 1} \in A^\mathbb{N}$, let us define

$$\bar{U}_i^K(t_i, \mathbf{a}) = \frac{1}{K} \sum_{k=1}^K u_i(t_i, a^k) \quad \text{for every } i, t_i \text{ and } K = 1, 2, \dots$$

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

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

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

6.1.2 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 $x_i(t_i)$ (namely, such that $x_i(t_i) = U_i(t_i, \sigma)$ for every i, t_i) and which define a Nash equilibrium of $\Gamma_\infty(q)$. For every player i , σ_i is described as follows:

at stage $k = 1$: choose a_i so as to reveal type t_i (which is possible since $|A_i| \geq |T_i|$)

survey the literature. The current list is still likely incomplete. My initial goal was to convey the main message of Koren (1992) in a simple model, which could be useful in applications but did not require the use of Blackwell's strategies.

at every stage $k \geq 2$: given the n -tuple of reported types r , play according to $\delta(r)$ if $\delta(r)$ was chosen at every previous stage; otherwise, play a punishment strategy in order to keep the first player j who did not follow $\delta(r)$ below his *ex post* individually rational level $v_j(t_j)$.

6.1.3 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 $\sigma_i(t_i)$ be the associated strategy of player i of type t_i , namely, $\sigma_i(t_i) = (\sigma_i^k(t_i))_{k \geq 1}$, with $\sigma_i^k(t_i) : A^{k-1} \rightarrow \Delta(A_i)$. Let $x_i(t_i) = U_i(t_i, \sigma)$ be the associated *interim* equilibrium payoff of player i of type t_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 k_0) and that afterwards (thus, at stages $k_0 + 1, k_0 + 2, \dots$), the players play independently of their types. Since σ is an equilibrium, player i of type t_i cannot benefit from playing according to $\sigma_i(r_i)$, with r_i possibly different from t_i , until stage k_0 and then, from stage $k_0 + 1$ on, by either continuing to play $\sigma_i(r_i)$ or just guaranteeing himself $v_i(t_i)$ (*i.e.*, by playing optimally in “his true one-shot game”, with payoffs $u_i(t_i, \cdot)$, at every stage $k_0 + 1, k_0 + 2, \dots$).¹⁵

More precisely, the equilibrium strategies $\sigma_i(t_i)$ generate probability distributions $\mu_\sigma(\cdot | t_1, \dots, t_n)$ 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 $T \times \Delta(A)$ such that

$$x_i(t_i) = U_i(t_i, \sigma) = E_{q, \mu_\sigma} \left(u_i(t_i, \tilde{\delta}) | t_i \right) \quad \text{for every } i, t_i \quad (16)$$

where E_{q, μ_σ} is the expectation with respect to P_{q, μ_σ} and $\tilde{\delta}$ stands for the frequency of move as a random variable.¹⁶

By considering the previous specific deviations of player i of type t_i (namely, mimic type r_i and/or play optimally in the one-shot game), we

¹⁵Note 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.

¹⁶If information transmission ends up after finitely many stages k_0 , $\tilde{\delta}$ can be interpreted as the frequency of moves from stage $k_0 + 1$ on.

obtain that

$$x_i(t_i) \geq E_{q, \mu_\sigma} \left(\max \left\{ u_i(t_i, \tilde{\delta}), v_i(t_i) \right\} \mid r_i \right) \quad \text{for every } i, t_i, r_i \quad (17)$$

We can also rely on a variant of the revelation principle to see that (16) and (17) must be satisfied as soon as σ is an equilibrium. Let us imagine that a fully reliable mediator asks the players to report their types and then given the n -tuple of reported types $r \in T$, chooses a frequency of moves $\delta \in \Delta(A)$ according to $\mu_\sigma(\cdot \mid r)$ and recommends δ to all players¹⁷. In other words, when the players report $r = (r_i)_{1 \leq i \leq n}$, the mediator selects δ exactly as the players themselves do at the equilibrium σ . (16) says that by telling the truth and following the recommendation of the mediator, the players get the same *interim* payoff than by playing σ . (17) says that if players $j \neq i$ tell the truth to the mediator, follow the recommendation δ as long as every player follows δ and punish any deviator at his *ex post* minmax level, then player i of type t_i cannot benefit from reporting type r_i to the mediator and/or not following δ .

Conditions (16) and (17) differ from (3) in two respects. (16) and (17) involve (type dependent) probability distributions over $\Delta(A)$, while (3) is formulated in terms of deterministic distributions $\delta(t)$, $t \in T$. Moreover, in (17), the probability distribution μ_σ is not necessarily completely revealing¹⁸.

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

$$E_{q, \mu_\sigma} \left(f(\tilde{\delta}) \mid t_i \right) = \sum_{t_{-i}} q_{-i}(t_{-i}) E_{\mu_\sigma} \left(f(\tilde{\delta}) \mid t_i, t_{-i} \right) \quad \text{for every } i, t_i$$

Hence, for every i, t_i , (16) can be rewritten as

$$x_i(t_i) = \sum_{t_{-i}} q_{-i}(t_{-i}) E_{\mu_\sigma} \left(u_i(t_i, \tilde{\delta}) \mid t_i, t_{-i} \right)$$

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

$$x_i(t_i) = \sum_{t_{-i}} q_{-i}(t_{-i}) u_i \left(t_i, E_{\mu_\sigma}(\tilde{\delta} \mid t_i, t_{-i}) \right)$$

¹⁷As in the standard proof of the Folk theorem under complete information, we interpret a distribution of moves δ as a deterministic sequence of moves (in A) which achieves the frequency of δ . This interpretation is straightforward if the components of δ are rational (in \mathbb{Q}).

¹⁸The above reliable mediator selects δ as a random function of the players' reported types but does not reveal these reported types.

which is the first part of (3) if we set $\delta(t) = E_{\mu_\sigma}(\tilde{\delta} | t)$.

By proceeding similarly and using in addition that “max” is convex, for every i, t_i, r_i , (17) can be rewritten as

$$\begin{aligned}
x_i(t_i) &\geq \sum_{t_{-i}} q_{-i}(t_{-i}) E_{\mu_\sigma} \left(\max \left\{ u_i(t_i, \tilde{\delta}), v_i(t_i) \right\} \mid r_i, t_{-i} \right) \\
&\geq \sum_{t_{-i}} q_{-i}(t_{-i}) \max \left\{ E_{\mu_\sigma} \left(u_i(t_i, \tilde{\delta}) \mid r_i, t_{-i} \right), v_i(t_i) \right\} \\
&\geq \sum_{t_{-i}} q_{-i}(t_{-i}) \max \left\{ u_i \left(t_i, E_{\mu_\sigma}(\tilde{\delta} \mid r_i, t_{-i}) \right), v_i(t_i) \right\} \\
&\geq \sum_{t_{-i}} q_{-i}(t_{-i}) \max \left\{ u_i(t_i, \delta(r_i, t_{-i})), v_i(t_i) \right\}
\end{aligned}$$

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

6.2 Computation of example 1

Recall from section 3.2. that the *ex post* individual rationality conditions are (6), namely,

$$\delta_d \leq k\delta, \beta_d \leq k\alpha \text{ and } \rho \geq \left(1 - \frac{1}{z}\right)(1 - \beta_d).$$

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 (7), namely,

$$p(k\delta - \delta_d) + (1 - p)(k\alpha - \beta_d) \geq p(k\rho - \beta_d).$$

Since $k\delta - \delta_d \leq \frac{k}{2}$ 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] Aumann, R. J. and M. Maschler (1995), *Repeated Games of Incomplete Information*, Cambridge: M.I.T. Press.
- [3] 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.
- [4] Bergin, J., (1989), “A characterization of sequential equilibrium strategies in infinitely repeated incomplete information games,” *Journal of Economic Theory* 47, 51-65.
- [5] Blackwell, D. (1956), “An analog of the minimax theorem for vector payoffs”, *Pacific Journal of Mathematics* 10, 1-8.
- [6] Chan, J., (2000), “On the non-existence of reputation effects in two-person infinitely repeated games”, mimeo, Johns Hopkins University.
- [7] Cripps, M. and J. Thomas (1995), “Reputation and commitment in two-person repeated games without discounting”, *Econometrica* 63, 1401-1419.
- [8] Cripps, M. and J. Thomas (1997), “Reputation and perfection in repeated common interest games”, *Games and Economic Behavior* 18, 141-158.
- [9] Cripps, M. and J. Thomas (2005), “Some asymptotic results in discounted repeated games of one-sided incomplete information”, *Mathematics of Operations Research* 28, 433-462.
- [10] Cripps, M., J. Thomas and K. Schmidt (1996), “Reputation in perturbed repeated games,” *Journal of Economic Theory* 69, 387-410.
- [11] 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.

- [12] Forges, F. (1988), “Communication equilibria in repeated games with incomplete information”, *Mathematics of Operations Research* 13, 191-231.
- [13] 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.
- [14] Forges, F. (2011, revised 2012), “A Folk theorem for Bayesian games with commitment”, *mimeo*, Université Paris-Dauphine.
- [15] Fudenberg, D. and E. Maskin (1986), “The Folk theorem in repeated games with discounting or with incomplete information”, *Econometrica* 54, 533-554.
- [16] Fudenberg, D. and J. Tirole (1991), *Game theory*, M.I.T. Press.
- [17] Hart S. (1985), “Nonzero-sum two-person repeated games with incomplete information”, *Mathematics of Operations Research* 10, 117–153.
- [18] Horner, J. and S. Lovo (2009), “Belief-free equilibria in games with incomplete information”, *Econometrica* 77, 453–487.
- [19] Israeli, E. (1999), “Sowing doubt optimally in two-person repeated games,” *Games and Economic Behavior* 28, 203-216.
- [20] Koren, G. (1992), “Two-person repeated games where players know their own payoffs”, *mimeo*, Courant Institute of Mathematical Sciences, New York University.
- [21] 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.
- [22] Kalai, A., E. Kalai, E. Lehrer and D. Samet (2010), “A commitment Folk theorem”, *Games and Economic Behavior* 69, 127-137.
- [23] Mailath, G. and L. Samuelson (2006), *Repeated Games and Reputations*, *Oxford University Press*.

- [24] Myerson, R. (1991), *Game theory: analysis of conflict*, Harvard University Press.
- [25] Palfrey, T. and H. Rosenthal (1994), “Repeated play, cooperation and coordination: an experimental study”, *Review of Economic Studies* 61, 545-565.
- [26] Peski, M. (2008), “Repeated games with incomplete information on one side”, *Theoretical Economics* 3, 29-84.
- [27] Peski, M. (2012), “Repeated games with incomplete information and discounting”, *mimeo*, University of Toronto.
- [28] Peters, M. and B. Szentes (2012), “Definable and contractible contracts”, *Econometrica* 80, 363-411.
- [29] Shalev, J. (1994), “Nonzero-sum two-person repeated games with incomplete information and known-own payoffs”, *Games and Economic Behavior* 7, 246-259.
- [30] Schmidt, K. (1993), “Reputation and equilibrium characterization in repeated games with conflicting interests”, *Econometrica* 61, 325-351.
- [31] Sorin, S. (1999), “Merging, reputation, and repeated games with incomplete information”, *Games and Economic Behavior* 29, 274-308.
- [32] Tennenholtz, M. (2004), “Program equilibrium”, *Games and Economic Behavior* 49, 363-373.
- [33] Wiseman, T. (2012), “A partial Folk theorem for games with private learning”, *Theoretical Economics* 7, 217-239.