



Working Papers

www.cesifo.org/wp

Bayesian Repeated Games and Reputation

Françoise Forges
Antoine Salomon

CESIFO WORKING PAPER NO. 4700
CATEGORY 2: PUBLIC CHOICE
MARCH 2014

An electronic version of the paper may be downloaded

- *from the SSRN website:* www.SSRN.com
- *from the RePEc website:* www.RePEc.org
- *from the CESifo website:* www.CESifo-group.org/wp

Bayesian Repeated Games and Reputation

Abstract

Under appropriate assumptions (private values and uniform punishments), the Nash equilibria of a Bayesian repeated game without discounting are payoff-equivalent to tractable, completely revealing, equilibria and can be achieved as interim cooperative solutions of the initial Bayesian game. This characterization does not apply to discounted games with patient players. In a class of public good games, the set of Nash equilibrium payoffs of the undiscounted game can be empty, while limit (perfect Bayesian) Nash equilibrium payoffs of the discounted game, as players become infinitely patient, do exist. These equilibria share some features with the ones of multi-sided reputation models.

JEL-Code: C730, C720, C710, D820, H410.

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

Françoise Forges
Department of Economics
University Paris Dauphine
Place du Maréchal de Tassigny
France - 75116 Paris
francoise.forges@gmail.com

Antoine Salomon
LEDa
University Paris Dauphine
Place du Maréchal de Tassigny
France - 75116 Paris
antoine.salomon@dauphine.fr

February 21, 2014

Conversations with Gorkem Celik, Martin Cripps, Gaël Giraud, Johannes Hörner, Vincent Iehlé, Frédéric Koessler, Marcin Peski and Péter Vida are gratefully acknowledged. Bernard Lebrun made a number of precise and helpful comments on an earlier draft.

1 Introduction

We consider infinitely repeated games in which the players privately know their own payoffs but are uncertain about the payoffs (i.e., the types) of the other players. We assume that the players start with given beliefs over each other's types and observe each other's decisions at every stage of the game. We are looking for a characterization of the perfect Bayesian equilibrium payoffs of the game, when the players are very patient. In other words, our goal is to understand how the well-known folk theorem extends in repeated games with private values and full monitoring.

The folk theorem goes back to the 1970's and the extension that we seek is very natural. So one would expect it to be already part of the literature. However, as we shall show below, only partial answers are available. Before surveying them, we first briefly describe our own contribution without reference to any earlier work.

We make the further assumption that, in the one-shot Bayesian game, *uniform punishment strategies* are available against every player. This property means that there is no need to know a player's type to punish him in the harshest possible way, i.e., at his ex post individually rational level. The assumption holds in a number of economic applications, including auctions and oligopoly.

Both properties, private values and uniform punishment strategies, are also satisfied in a class of familiar public good games, in which every player is endowed with some private good. Every player can either enjoy his endowment or devote it to a public good, which is produced if and only if enough players contribute. Individuals differ in their relative preferences for the private and the public goods. Such games will be our reference model.

Hoping for some intuition on how the equilibrium payoffs of the Bayesian discounted infinitely repeated game look like when players become very patient, we first characterize the set $\mathcal{N}[\Gamma_\infty]$ of equilibrium payoffs of the game Γ_∞ without discounting, in which the payoffs are evaluated by the limit of means criterion. The characterization (stated as Proposition 1) turns out to be amazingly tractable: a vector payoff is an equilibrium of the Bayesian undiscounted infinitely repeated game Γ_∞ if and only if it is a particular, completely revealing equilibrium payoff, which can be entirely described in terms of the one-shot game. A partial analog to the standard folk theorem (stated as Proposition 3) follows: the equilibrium payoffs of the Bayesian undiscounted infinitely repeated game can be interpreted as feasible, incen-

tive compatible and individually rational payoffs in the one-shot game.

The similarity with the complete information case stops essentially there. Applying the previous characterization to our reference public good games, we find that Bayesian undiscounted infinitely repeated games *may not have any equilibrium*, i.e., it may happen that $\mathcal{N}[\Gamma_\infty] = \emptyset$. This result is stated as Proposition 2. By contrast, Nash’s theorem guarantees existence of an equilibrium in every δ -discounted infinitely repeated game Γ_δ , for every fixed discount factor δ . But the problem is then to describe the solutions as players become very patient, namely as the discount factor δ goes to 1. The fact that the undiscounted game does not have any equilibrium makes useless the techniques that are often applied to check the non-emptiness of relevant limit sets like $\liminf_{\delta \rightarrow 1} \mathcal{N}[\Gamma_\delta]$.

Given this difficulty, we pursue with the study of a particular model, namely, the two-person Bayesian δ -discounted infinitely repeated version of the public good games mentioned above. It turns out that, with discounting, Proposition 1 (all equilibria are payoff-equivalent to completely revealing ones) and Proposition 2 (the slightest doubt on the players’ types can lead to non existence of equilibrium) no longer hold, even when players become increasingly patient. More importantly, for every discount factor δ , we explicitly construct a family of perfect Bayesian equilibria in specific strategies, in which the players behave as in a war of attrition. We prove that the payoffs associated with these war of attrition equilibria converge as the discount factor δ goes to 1. As a consequence, the (inferior) limit of the corresponding sets of perfect Bayesian equilibrium payoffs, i.e., $\liminf_{\delta \rightarrow 1} \mathcal{N}[\Gamma_\delta]$, is non-empty (and even has a nonempty interior). These results are summarized in Proposition 7.

The war of attrition equilibria are first constructed in an auxiliary reputation game, in which normal players play a game with strictly competitive interests (SCI) and possibly face automata. In this auxiliary game, we allow the players to have different discount factors. Propositions 4, 5 and 6 propose three different sets of sufficient conditions (on the parameters of the SCI game, the strategies of the automata and the discount factors) that guarantee the existence of perfect Bayesian equilibria in war of attrition strategies in the auxiliary game. We believe that this analysis, which is entirely performed in the discrete time, discounted reputation game, deserves interest for its own sake.

Time has come to describe how our results are connected to three popular themes in the **repeated games literature**: information transmission,

cooperation and reputation.

Information transmission

Aumann and Maschler started to analyze two-person infinitely repeated games with incomplete information in the 1960's (see Aumann and Maschler (1995)). They mostly studied the zero-sum case, in which the assumption of known-own payoffs is obviously not meaningful. They provided a characterization of the value v_∞ of the undiscounted zero-sum infinitely repeated game and showed that v_∞ might not exist when both players have private information. Mertens and Zamir (1971) proved that the values v_δ of the δ -discounted games always converge as δ goes to 1; they characterized $\lim_{\delta \rightarrow 1} v_\delta$ and showed that it coincides with v_∞ when the latter exists. Our main results (Propositions 1, 2 and 7) partially extend the previous ones to non-zero-sum games.

Together with Stearns, Aumann and Maschler also started to investigate two-person non-zero-sum undiscounted infinitely repeated games in which only one of the players has private information (see Aumann et al. (1968) and Aumann and Maschler (1995)). In this framework, a full characterization of Nash equilibrium payoffs was obtained by Hart (1985) and existence was established by Sorin (1983) and Simon et al. (1995). As the previous authors, 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. He shows that the Nash equilibrium of the undiscounted infinitely repeated game Γ_∞ are payoff-equivalent to completely revealing equilibria of the infinitely repeated game Γ_∞ .¹

In Proposition 1, we add an assumption to Koren (1992)'s ones, namely, that uniform punishment strategies are available. This enables us to formulate the characterization by simple inequalities, entirely in terms of the one-shot game. By contrast, Koren (1992)'s characterization makes use of Blackwell (1956)'s approachability, i.e., of punishments that are typical to the infinitely repeated game. Our new assumption enables us to define individual rationality in the infinitely repeated game exactly as in the one-shot game, so that we can dispense with approachability. As a by-product, our re-

¹Koren (1992)'s result generalizes Shalev (1994)'s characterization of Nash equilibrium payoffs in undiscounted repeated games with lack of information on one side and known-own payoffs (see Forges (1992) for a survey of results on non-zero sum infinitely repeated games with incomplete information).

sult immediately goes through in the case of n players.² Known-own payoffs are crucial in Koren (1992) and in Proposition 1. Without this assumption, the Nash equilibria of Γ_∞ 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)). In the appendix, we illustrate the further role of uniform punishment strategies in Proposition 1.

Koren (1992) already proposed an example of an undiscounted infinitely repeated game Γ_∞ with known-own payoffs that did not have any Nash equilibrium. We recall it as Example 1. Proposition 2, together with a complement in the appendix, further shows that the emptiness of $\mathcal{N}[\Gamma_\infty]$ is quite robust. The phenomenon arises in a large class of public good games (see, e.g., Fudenberg and Tirole (1991, example 6.1) and Palfrey and Rosenthal (1994)). It does not rely on the fact that some players' types have a dominant strategy in the repeated game (in other words, it does not rely on the fact that for some given types, the set of feasible and individually rational payoffs has an empty interior).

The papers on non-zero-sum games surveyed up to now all deal with undiscounted infinitely repeated games. A natural next step is to check whether Mertens and Zamir (1971)'s result for zero-sum games survives, namely, whether an appropriate limit of the sets of equilibrium payoffs $\mathcal{N}[\Gamma_\delta]$ as δ goes to 1 is not empty and can be characterized.³ Restricting to public good games, Proposition 7 gives a positive answer to the first question. Peski (2008) focuses on two-person discounted repeated games Γ_δ 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). In this very particular model, he shows that $\lim_{\delta \rightarrow 1} \mathcal{N}[\Gamma_\delta]$ exists (in the sense that $\liminf_{\delta \rightarrow 1} \mathcal{N}[\Gamma_\delta] = \limsup_{\delta \rightarrow 1} \mathcal{N}[\Gamma_\delta]$). Peski (2008) characterizes the set $\lim_{\delta \rightarrow 1} \mathcal{N}[\Gamma_\delta]$ and confirms a phenomenon that was already suggested in Cripps and Thomas (2003), namely, that more equilibrium payoffs can be achieved in the limit discounted case than in the undiscounted case, namely, that $\lim_{\delta \rightarrow 1} \mathcal{N}[\Gamma_\delta] \supsetneq \mathcal{N}[\Gamma_\infty]$ (where $\mathcal{N}[\Gamma_\infty] \neq \emptyset$, as shown by Shalev (1994)).

Peski (2013) extends the previous results to a larger class of n -person

²See Hörner et al. (2011) for a generalization of Blackwell (1956) to an arbitrary number of players.

³Bergin (1989) proposes a characterization of equilibrium strategies in δ -discounted infinitely repeated games Γ_δ but does not consider $\lim_{\delta \rightarrow 1} \mathcal{N}[\Gamma_\delta]$.

discounted repeated games with known-own payoffs, in which several players may possess private information. More precisely, Peski (2013) makes an “open thread assumption” which, in the two-person case, amounts to the existence of a belief-free equilibrium in the sense of Hörner and Lovo (2009). As already in Peski (2008), the non emptiness of limit sets of Bayesian perfect Nash equilibrium payoffs as players become increasingly patient is not an issue in Peski (2013), at least in the two player case, since belief free equilibria are assumed to exist at the outset. The main result in Peski (2013) is the full characterization of $\lim_{\delta \rightarrow 1} \mathcal{N}[\Gamma_\delta]$ as the set of payoffs of special equilibria, the finitely revealing ones, which have the property that the players update their beliefs finitely many times.

As shown by Hörner and Lovo (2009), two-person discounted repeated games with lack of information on one side always have a belief-free equilibrium when players are patient enough.⁴ Hence Peski (2013)’s open thread assumption is always satisfied in this particular case. However, the assumption easily fails in familiar applications like our public good games, in which both players have private information. Indeed, it appears that, as soon as the *undiscounted* infinitely repeated good game has no equilibrium for *some* specification of the players’ beliefs, the discounted infinitely repeated game cannot have any belief free equilibrium.

Our Proposition 7 shows that nonetheless, in every discounted repeated version of our public good game with a sufficiently high discount factor, there exist perfect Bayesian equilibrium payoffs which converge as the players become increasingly patient. We thus show in particular that Peski (2013)’s open thread assumption is not necessary for the convergence of equilibrium payoffs, i.e., for $\liminf_{\delta \rightarrow 1} \mathcal{N}[\Gamma_\delta]$ to be nonempty. As already mentioned above, we establish Proposition 7 by explicitly constructing equilibria, in which the players’s strategies follow a war of attrition. These equilibria differ from the ones that are used in Peski (2013)’s characterization, as they are not finitely revealing. However, our construction can be adapted to show that, in the public good game, $\liminf_{\delta \rightarrow 1} \mathcal{N}[\Gamma_\delta]$ contains finitely revealing equilibrium payoffs.

⁴This result can be deduced from Shalev (1994)’s characterization of the set of Nash equilibrium payoffs $\mathcal{N}[\Gamma_\infty]$ of any two-person undiscounted infinitely repeated game Γ_∞ with lack of information on one side and known-own payoffs, which implies that $\mathcal{N}[\Gamma_\infty]$ is always nonempty.

Cooperation

The tractable characterization in Proposition 1 facilitates the comparison between the solutions of the undiscounted repeated game and the cooperative solutions of the initial Bayesian game. More precisely, Proposition 1 leads to a partial version of the folk theorem, which we state as Proposition 3: for any n -person Bayesian game Γ with known-own payoffs and uniform punishments, the set $\mathcal{N}[\Gamma_\infty]$ of Nash equilibrium payoffs of the undiscounted infinitely repeated game Γ_∞ is contained in the set $\mathcal{F}[\Gamma]$ of *interim* cooperative solutions of the Bayesian game, as defined in Myerson (1991) and Forges (2013). By definition, the set $\mathcal{F}[\Gamma]$ consists of the payoffs that are feasible (i.e., achievable by means of a mechanism), incentive compatible and (interim) individually rational. The fact that $\mathcal{N}[\Gamma_\infty] \subseteq \mathcal{F}[\Gamma]$ confirms that the repetition of the game enables the players to cooperate, as in the folk theorem with complete information. However, it may happen that $\mathcal{N}[\Gamma_\infty] \subsetneq \mathcal{F}[\Gamma]$, even when $\mathcal{N}[\Gamma_\infty]$ is not empty. Hence, according to a common, strict interpretation, we could say that the folk theorem does not hold.

Our characterization shows that, under incomplete information, the cooperative solutions of the one-shot game (i.e., $\mathcal{F}[\Gamma]$) and the non-cooperative solutions of the undiscounted repeated game (i.e., $\mathcal{N}[\Gamma_\infty]$) mostly differ in the individual rationality levels of the players. Under our assumptions of private values and uniform punishments, the *ex post* individual rationality level $v_i(\theta_i)$ of player i , namely the level at which the other players can punish him when they know his type θ_i , 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. 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 3 is no longer true, while Koren (1992)'s characterization still holds, at least in the two-person case. To sum up, the inclusion $\mathcal{N}[\Gamma_\infty] \subseteq \mathcal{F}[\Gamma]$, which does not seem very surprising at first sight, crucially depends on the fact that uniform punishment strategies are available.

A recent paper of Wiseman (2012) offers a different perspective on the previous topic. He 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 (see his example 3), he makes an assumption that

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, 2013).

Reputation

An extensive literature shows how a player can use the other players’ uncertainty on his payoff to establish a reputation (see the references in Mailath and Samuelson (2006)).⁵ The repeated games of this paper share a property with reputation models: players know their own payoffs. However, our basic model is not necessarily generated by perturbing a game with complete information and may not involve any commitment type. Many other differences could be mentioned (e.g., most reputation models deal with a single informed player⁶ and imperfect monitoring) but the main one lies in the motivations. We are interested in the *characterization* and the *existence* of equilibrium payoffs while a reputation effect is identified by a form of equilibrium payoff *uniqueness*. Loosely speaking, a reputation effect obtains when all appropriately refined equilibria of the perturbed game guarantee a high payoff to the player who is possibly committed to some behavior (see, e.g., Schmidt (1993) for a more precise, yet informal, definition).

The previous comparison is relevant for characterization results like Proposition 1. More must be said for the results that are established for public good games. In these, there are two players, who can both be “normal” or “greedy” and have two actions (“contribute” or “do not contribute”). When both players are normal, they play a particular game with strictly competitive interests (SCI), a notion defined in Cripps et al. (2005) and Atakan and Ekmekci (2013). The difference with the latter papers is that our greedy players are not committed to a specific strategy. Proposition 2, which says that, for a substantial set of parameters, the undiscounted repeated game has no Nash equilibrium at all, implies that, without discounting, no two-sided reputation effect is achievable. While Cripps and Thomas (1995) and Israeli (1999) acknowledge this phenomenon (as a direct consequence of Ko-

⁵See also Sorin (1999) for a synthetic presentation of various related models, including infinitely repeated games with known own payoffs.

⁶There are exceptions, though, e.g., Kreps and Wilson (1982), Chan (2000) and Atakan and Ekmekci (2013).

ren (1992)’s example), they do not demonstrate its robustness, as we do here (especially in the three action version of the public good game given in the appendix).

We also construct sequences of equilibrium payoffs in every δ -discounted public good game which all converge as δ goes to 1 (Proposition 7). These equilibria are constructed in two steps. First, we assume that the greedy players are committed to some strategies, while the normal players have complete freedom to pick theirs. This defines an auxiliary reputation game. Then we show that it is a best reply for the greedy players to behave as automata. The auxiliary game is at the same time more restrictive and more general than the SCI model of Atakan and Ekmekci (2013). We assume that the players have only two actions⁷ but, to account for the strategic possibilities of our greedy players, we are led to consider a richer class of automata.

We derive explicit conditions (on the parameters of the SCI game, the strategies of the automata and the normal players’ discount factors⁸) which guarantee that the auxiliary game has a perfect Bayesian equilibrium in war of attrition strategies. Atakan and Ekmekci (2013) make use of a war of attrition game in continuous time that was proposed in Abreu and Gul (2000) to study reputation in bargaining. We perform our analysis in the repeated (i.e., discrete time) game with a fixed discount factor and generate a war of attrition through appropriate players’ strategies. By proceeding in this way, we can compute the exact expression of a family of equilibrium payoffs (each equilibrium corresponding to a specification of the automata) and show that they do converge as players become increasingly patient.

Atakan and Ekmekci (2013) are rather interested in a reputation result. Having identified relevant reputation payoffs for each normal player in the SCI game, they allow the normal players to face particular, elementary automata and they prove that, when the stage game is played arbitrarily frequently, the normal players get approximately their reputation payoffs at any Bayesian perfect equilibrium. Our results are consistent with theirs, in the sense that, when we restrict to elementary automata, we recover their reputation payoffs as limit payoffs.

As suggested in the previous paragraphs, the paper develops as follows: Section 2 describes the one-shot Bayesian game and the reference public good

⁷We allow for arbitrary payoffs as long as the SCI property is satisfied.

⁸The normal players may have different discount factors in this part of the paper.

games; Section 3 proposes a characterization of the set of equilibrium payoffs in the Bayesian undiscounted repeated game and demonstrates its possible emptiness in public good games; Section 4 deduces a partial folk theorem; Section 5 goes on with the analysis of equilibrium payoffs in Bayesian discounted repeated games; Section 5.1 is devoted to auxiliary reputation games, which are based on SCI games with two actions and rich automata; Section 5.2 concentrates on the public good games; Section 6 is an appendix which contains supplementary material and proofs. Sections 5.1 and 6.4 (resp., 5.2 and 6.5) contain precise elements of comparison with Atakan and Ekmekci (2013) (resp., Peski (2013)).

2 Basic Bayesian game

2.1 Definition

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 values⁹, which we denote as $\Gamma(q)$, with $q = (q_i)_{1 \leq i \leq n}$.¹⁰ 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.

For any finite set E , let us denote as $\Delta(E)$ the set of probability distributions over E . A mixed strategy¹¹ of player i in $\Gamma(q)$ is a mapping from Θ_i

⁹While private values (i.e., “known own payoffs”) are crucial for our results, the independence assumption can be relaxed in most of them.

¹⁰We 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.

¹¹More correctly, “behavior strategy”.

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

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

We introduce a new property, which we call “uniform punishment strategies”. To define it, consider the (type-dependent) individually rational level $v_i(\theta_i)$ of player i , for every possible type θ_i . We assume that the other players, $j \neq i$, have a strategy profile (of course independent of player i ’s type) which enable them to keep player i ’s payoff below $v_i(\theta_i)$, for every possible type θ_i . So, while player i can make use of his information θ_i to guarantee himself the level $v_i(\theta_i)$, the other players can guarantee that his payoff does not exceed this level, even if they do not know his type θ_i .

More precisely, for every player i , $i = 1, \dots, n$, and $\theta_i \in \Theta_i$, the individually rational level $v_i(\theta_i)$ is

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

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 lowest amount at which players $j \neq i$ can hold player i when they know his type θ_i and correlate their strategies.

Our **assumption** of “**uniform punishment strategies**” can be formalized as follows:

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

When (2) holds, τ_{-i} defines independent¹² punishment strategies which enable players $j \neq i$ to punish player i uniformly, *i.e.*, whatever his type θ_i is,

¹²Independent punishment strategies are important for proposition 1.

but even more, to keep player i 's payoff below his *ex post* individual rationality level.¹³ Assumption (2) is satisfied in many applications, in particular in the public good games below.¹⁴ We will clarify its role in the appendix.

2.2 A public good game

The example belongs to a class of public good games that was studied, *e.g.*, by Palfrey and Rosenthal (1994). There are two players. The private information of player i , $i = 1, 2$, 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 as soon as one of the players contributes. The value of the public good is normalized to 1 for both players.

The payoff matrix associated with the pair of types (θ_1, θ_2) is thus

$$\begin{array}{cc} & c & d \\ c & 1, 1 & 1, 1 + \theta_2 \\ d & 1 + \theta_1, 1 & \theta_1, \theta_2 \end{array}$$

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: 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 this public good game, by playing d , player i guarantees himself θ_i while by playing c , he guarantees himself 1. Hence, by playing according to

¹³As a slight weakening, $v_i(\theta_i)$ could just be defined as

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

¹⁴As a recent reference, 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.

his type, player i can guarantee himself $\max\{\theta_i, 1\}$. By not contributing, namely by playing $a_j = d$, player j guarantees that player i 's payoff does not exceed $\max\{\theta_i, 1\}$. Hence, $v_i(\theta_i) = \max\{\theta_i, 1\}$ and $a_j = d$ is a uniform punishment strategy of player j against player i .

We will specifically consider a Bayesian public good game in which each player has two possible types: $\Theta_i = \{\omega, z\}$, $i = 1, 2$, where $0 < \omega < 1$ and $z > 2$: ω represents a “normal” type, who values the public good more than his initial endowment, while z represents a “greedy” type. The possible payoff matrices are thus

$$\begin{array}{c|cc|cc|cc}
 & & \theta_2 = \omega & & \theta_2 = z & & \\
 & & c & d & c & d & \\
 \hline
 \theta_1 = \omega & c & 1, 1 & 1, 1 + \omega & 1, 1 & 1, 1 + z & \\
 & d & 1 + \omega, 1 & \omega, \omega & 1 + \omega, 1 & \omega, z & \\
 \hline
 \theta_1 = z & c & 1, 1 & 1, 1 + \omega & 1, 1 & 1, 1 + z & \\
 & d & 1 + z, 1 & z, \omega & 1 + z, 1 & z, z &
 \end{array} \tag{3}$$

We denote as p_i , $0 < p_i < 1$, the probability that player i is normal ($i = 1, 2$) and we refer to this game as $PG(p, \omega, z)$, with $p = (p_1, p_2)$.

3 Bayesian undiscounted repeated game

Let us turn to infinitely repeated versions of the Bayesian game $\Gamma(q)$, starting with the undiscounted one, which we denote as $\Gamma_\infty(q)$. The players' types are intrinsic characteristics, which 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 limits of means, namely 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 $\prod_{j \neq i} q_j(\theta_j)$ and let us denote as $\mathcal{N}[\Gamma_\infty(q)]$ the set of all Nash equilibrium interim expected payoffs of $\Gamma_\infty(q)$. The next Proposition provides a tractable characterization of $\mathcal{N}[\Gamma_\infty(q)]$.¹⁵

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

$$\begin{aligned} x_i(\theta_i) &= \sum_{\theta_{-i} \in \Theta_{-i}} q_{-i}(\theta_{-i}) u_i(\theta_i, \pi(\theta_i, \theta_{-i})) \\ &\geq \sum_{\theta_{-i} \in \Theta_{-i}} q_{-i}(\theta_{-i}) \max \{u_i(\theta_i, \pi(\theta'_i, \theta_{-i})), v_i(\theta_i)\} \end{aligned} \quad (4)$$

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 to the other 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 independent uniform 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.

¹⁵Koren (1992) already established a version of this result in the case of only *two* players and without assuming uniform punishments. The latter assumption greatly facilitates the formulation of the equilibrium conditions and the extension to n players.

According to the nondeviation condition (4), player i , of type θ_i , takes for granted that players $j \neq i$ follow the equilibrium strategies. He can report a type θ'_i , which is 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, the other players notice the deviation and player i is punished at the level $v_i(\theta_i)$.

Condition (4) 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 $\theta'_i = \theta_i$, (4) is equivalent to

$$\text{For every } i \text{ and } \theta = (\theta_i, \theta_{-i}) \in \Theta : \quad u_i(\theta_i, \pi(\theta)) \geq v_i(\theta_i) \quad (5)$$

which implies that

$$\text{For every } i \text{ and } \theta_i \in \Theta_i : \quad x_i(\theta_i) \geq v_i(\theta_i) \quad (6)$$

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 .

Proposition 1 is established in the appendix. Intuitively, three properties play a key role to show that an arbitrary equilibrium is payoff equivalent to a completely revealing one. First, the posterior probability distributions of every player on the other players' types converge, so that the players' behavior on the equilibrium path becomes nonrevealing. Let us make things extreme and assume that at equilibrium, no information is revealed from some stage T on; then, from T on, the players' behavior on the equilibrium path is described by a probability distribution π_T over A (which depends on the history of moves up to stage T). Second, once the players have settled on π_T , every player i may as well learn the types of the other players, since, thanks to private values, his payoff only depends on θ_i and π_T . Finally, because there is no discounting, the information that is transmitted in T stages can as well be transmitted at once, at the first one. In the appendix, we also show the role of uniform punishment strategies.

3.2 Existence of Nash equilibrium

Thanks to Proposition 1, the set $\mathcal{N}[\Gamma_\infty(q)]$ of Nash equilibrium payoffs of the undiscounted repeated game has a tractable representation so that it is rela-

tively easy to check whether it is empty or not. Let us denote as $PG_\infty(p, \omega, z)$ the infinitely repeated version of the public good game $PG(p, \omega, z)$ introduced in Section 2.2. We will show that $PG_\infty(p, \omega, z)$ has no Nash equilibrium as soon as z , p_1 and p_2 are sufficiently large.

Koren (1992) already illustrated the same phenomenon in the two-player game of Example 1 below. However, in the latter, as soon as one of the players is greedy, the set of feasible individually rational payoffs has an empty interior. The same happens in $PG_\infty(p, \omega, z)$, when both players are greedy, but only in that case. As shown in the appendix, our example can be modified slightly so as to completely get rid of this peculiarity.

Proposition 2 *Consider the undiscounted infinitely repeated public good game $PG_\infty(p, \omega, z)$. If the priors p_1, p_2 and the payoff z are sufficiently large, $PG_\infty(p, \omega, z)$ has no Nash equilibrium, i.e., $\mathcal{N}[PG_\infty(p, \omega, z)] = \emptyset$. The same holds in a variant of $PG_\infty(p, \omega, z)$ in which, for every pair of types, the interior of the set of feasible individually rational payoffs is nonempty.*

Proof: Let us consider an interim expected equilibrium payoff x of the game $PG_\infty(p, \omega, z)$. According to Proposition 1, x is described by probability distributions $\pi(\theta)$ over $A = \{c, d\} \times \{c, d\}$ for every pair of types $\theta \in \{n, g\} \times \{n, g\}$, which must satisfy the equilibrium conditions (4) in Proposition 1, in particular the ex post individual rationality conditions (6). From a geometric representation of the feasible, individually rational payoffs when one of the players is normal and the other one is greedy, it appears that if z is sufficiently large, $z \approx z + 1$. The ex post individual rationality conditions thus imply that

$$\pi((c, d) \mid n, g) \geq 1 - \varepsilon, \quad \pi((d, c) \mid g, n) \geq 1 - \varepsilon \quad (7)$$

where we write $\pi(a \mid \theta)$ for $\pi(\theta)(a)$, for every $a \in A$. A formal proof of these inequalities is given in the appendix.

Condition (4) for player 1 of type n implies that he should not benefit from the following deviation: pretend to be of type g at stage 0; thereafter, mimic type g when the other player turns out to be normal and defend his individual rationality level (namely play c) when the other player turns out to be greedy:

$$p_2 u_1(n, \pi(n, n)) + (1 - p_2) u_1(n, \pi(n, g)) \geq p_2 u_1(n, \pi(g, n)) + (1 - p_2) \cdot 1 \quad (8)$$

The latter condition amounts to ex post individual rationality when $p_2 = 0$.

Let us set $\rho_1 = \pi((c, d) \mid n, g)$ and $\rho_2 = \pi((d, c) \mid g, n)$. Feasibility implies that

$$u_1(n, \pi(n, g)) \leq \rho_1 \cdot 1 + (1 - \rho_1)(1 + \omega) \quad (9)$$

$$u_1(n, \pi(g, n)) \geq \rho_2(1 + \omega) \quad (10)$$

Hence, (8) implies that

$$p_2 u_1(n, \pi(n, n)) + (1 - p_2)(1 - \rho_1)\omega \geq p_2 \rho_2(1 + \omega)$$

so that

$$u_1(n, \pi(n, n)) \geq \rho_2(1 + \omega) - \frac{1 - p_2}{p_2}(1 - \rho_1)\omega \quad \text{if } p_2 > 0$$

and, if z is sufficiently large, from (7),

$$u_1(n, \pi(n, n)) \geq (1 + \omega) - \varepsilon(1 + \frac{\omega}{p_2}).$$

If in addition p_2 is sufficiently large, say $p_2 \geq \omega$, we must have

$$u_1(n, \pi(n, n)) \geq (1 + \omega) - 2\varepsilon$$

By proceeding similarly for player 2,

$$u_2(n, \pi(n, n)) \geq (1 + \omega) - 2\varepsilon$$

which contradicts feasibility for ε sufficiently small (since $u_1(n, \pi(n, n)) + u_2(n, \pi(n, n)) \leq 2 + \omega$).

The variant of $PG_\infty(p, \omega, z)$ is described in the appendix. ■

The previous finding should be contrasted with what happens in one-sided reputation models. If only one player is uncertain of the type of the other, there exist equilibria accounting for a reputation effect, even if the game is not discounted (see Shalev (1994), Cripps and Thomas (1995) and Israeli (1999)). However, as soon as both players believe with arbitrarily small probability that the other player could be (very) greedy, the game has no equilibrium at all. Two features of the previous example are important: lack of information on both sides and no discounting. The analysis of the discounted game, when the players become infinitely patient, leads to different conclusions, as we will see in Section 5.

4 Repetition and cooperation

Under complete information, the set of feasible and individually rational payoffs of a one-shot game coincides not only with the set of Nash equilibrium payoffs of the infinitely repeated game but also with a well founded set of cooperative solutions of the one-shot game (see *e.g.* Kalai *et al.* (2010)). The inequalities (4) in Proposition 1 can be interpreted as describing cooperative solutions for the one-shot Bayesian game $\Gamma(q)$, so that Proposition 1 establishes a relationship between “repetition” and “cooperation”. However, under incomplete information, a more natural candidate for the set of cooperative solutions 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.¹⁶ Myerson (1991)’s definitions take a simpler form in our framework of independent private values. We then establish a partial analog of the 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

$$x_i(\theta_i) = \sum_{\theta_{-i}} q_{-i}(\theta_{-i}) u_i(\theta_i, \pi(\theta_i, \theta_{-i})) \quad i = 1, \dots, n, \theta_i \in \Theta_i \quad (11)$$

A feasible payoff x achieved through π (as in (11)) is *incentive compatible* if

$$x_i(\theta_i) \geq \sum_{\theta_{-i}} q_{-i}(\theta_{-i}) u_i(\theta_i, \pi(\theta'_i, \theta_{-i})) \quad \text{for every } i, \theta_i, \theta'_i \in \Theta_i \quad (12)$$

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 that¹⁷

$$x_i(\theta_i) \geq \max_{a_i \in A_i} u_i(\theta_i, a_i, \tau_{-i}) \quad \text{for every } \theta_i \in \Theta_i \quad (13)$$

¹⁶Under an assumption that is similar to our uniform punishment strategies, Peters and Szentes (2012) also find that $\mathcal{F}[\Gamma(q)]$ corresponds to a set of reasonable cooperative solutions of $\Gamma(q)$, which are achieved through interim commitment.

¹⁷Literally, Myerson (1991)’s *interim* individual rationality condition requires that there exists a type dependent correlated strategy of players $j \neq i$, $\tau_{-i}(\theta_{-i}) \in \Delta(A_{-i})$, $\theta_{-i} \in \Theta_{-i}$, such that $x_i(\theta_i) \geq \max_{a_i \in A_i} \sum_{\theta_{-i}} q_{-i}(\theta_{-i}) u_i(\theta_i, a_i, \tau_{-i}(\theta_{-i}))$ for every $\theta_i \in \Theta_i$. But, with independent private values, (13) is an equivalent formulation, since $u_i(\theta_i, \cdot)$ is linear.

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 (11), (12) and (13). $\mathcal{F}[\Gamma(q)]$ 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.

Lemma 0 *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, (13)), x is ex post individually rational (namely, (6)): $INTIR_i[\Gamma(q)] \subseteq EXPIR_i[\Gamma(q)]$ 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 0 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 (13).

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 $\mathcal{F}[\Gamma(q)]$ ((11), (12) and (13) above): the equality in (4) is (11), the inequality in (4) implies (12) and (6), which in turn implies (13) by Lemma 0. ■

As stated in Proposition 2, unlike $\mathcal{F}[\Gamma(q)]$, $\mathcal{N}[\Gamma_\infty(q)]$ can be empty. The characterization in Proposition 1 indeed makes clear that repetition is just equivalent to a very demanding form of cooperation in the one-shot game, in which every player can decide to participate (and get his individually rational level if he does not) after having learnt the types of the other players. This participation constraint makes incentive compatibility harder in $\mathcal{N}[\Gamma_\infty(q)]$

than in $\mathcal{F}[\Gamma(q)]$. Example 1 below, taken from Koren (1992), further illustrates that $\mathcal{N}[\Gamma_\infty(q)]$ can be strictly included in $\mathcal{F}[\Gamma(q)]$.¹⁸

Example 1: 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 $\pi(\theta)$, $\theta \in \Theta$, 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

(14)

It is easily checked that $\pi(\theta)$ satisfies (12) and (13), 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

¹⁸For appropriate values of the prior p , it also happens in the public good game that $\mathcal{N}[PG_\infty(p, \omega, z)]$ is not empty and is strictly included in $\mathcal{F}[PG(p, \omega, z)]$. However, a full characterization of $\mathcal{N}[PG_\infty(p, \omega, z)]$ seems much harder in this game than in Koren (1992)'s example.

a payoff in $\mathcal{N}[\Gamma_\infty(q)]$, $\pi(\theta)$ must satisfy (4); 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 (4), 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 (14) 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 (14). ■

5 Bayesian discounted repeated game

In this section, we focus on two players and we allow every player to evaluate his payoff sequence with a discount factor $\delta \in (0, 1)$. The δ -discounted version $\Gamma_\delta(q)$ of the infinitely repeated game is played as in Section 3, but a sequence of actions $\mathbf{a} = (a^t)_{t \geq 0} \in A^\mathbb{N}$ leads now to the payoff

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

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.

The equilibrium payoffs of the undiscounted repeated game $\Gamma_\infty(q)$ are natural candidates that could belong to $\lim_{\delta \rightarrow 1} \mathcal{N}[\Gamma_\delta(q)]$. However, as stated in Proposition 2, the undiscounted repeated public good game $PG_\infty(p, \omega, z)$ may have no equilibrium. Under the assumption that there exist belief-free equilibria (in the sense of Hörner and Lovo (2009)), Peski (2013) characterizes the set $\lim_{\delta \rightarrow 1} \mathcal{N}[\Gamma_\delta(q)]$ as the set of *finitely* revealing equilibrium payoffs, in which beliefs are updated finitely many times. This is again of little use in our public good game. Indeed, Hörner and Lovo (2009)'s characterization of belief-free equilibrium payoffs and Proposition 1 imply that there cannot be any belief-free equilibrium when the undiscounted repeated game has no equilibrium.¹⁹

Let $PG_\delta(p, \omega, z)$ be the δ -discounted infinitely repeated game associated with the public good game $PG(p, \omega, z)$ introduced in Section 2.2. Even if there are no equilibria that would easily be shown to belong to $\lim_{\delta \rightarrow 1} \mathcal{N}[PG_\delta(p, \omega, z)]$ when p_1, p_2 and z are large, we establish below that (the interior of) this set is always nonempty.²⁰

The proof of this result will be fully constructive: we will show that if both players are sufficiently patient, $PG_\delta(p, \omega, z)$ has a family of perfect Bayesian equilibria of the form of a *war of attrition*, in which the normal players do not contribute until one of them gives up and does contribute. Then, the one who has first given up goes on by contributing at most stages. The contributing player thus reveals that his type is normal, whereas the other player only contributes occasionally and keeps his type unknown. We will first define precisely the war of attrition equilibria in auxiliary reputation games, in which the payoffs of the normal players have a general structure and the greedy players reduce to automata. We consider a rather large class of automata to reflect the strategic possibilities of the greedy players.

¹⁹Hörner and Lovo (2009) show that all belief-free equilibrium payoffs must belong to the set V^* of ex post incentive compatible and ex post individually rational payoffs. By Proposition 1, $V^* \subseteq \cap_q \mathcal{N}[\Gamma_\infty(q)]$.

²⁰More precisely, $\liminf_{\delta \rightarrow 1} \mathcal{N}[PG_\delta(p, \omega, z)] \neq \emptyset$, with 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 $x_n \rightarrow x$.

5.1 A family of reputation games

We introduce a family of reputation games $\mathcal{G}_\delta(p, i, \gamma)$, where $\delta = (\delta_1, \delta_2)$, $p = (p_1, p_2)$, $i = 1$ or 2 , $\gamma = (\gamma_1, \gamma_2)$, $p_k, \gamma_k \in [0, 1]$, $\delta_k \in [0, 1)$, $k = 1, 2$, and $\gamma_1 + \gamma_2 < 1$. There are two agents. At stage -1 , according to independent moves of nature, agent k ($k = 1, 2$) is a player, with probability p_k , or an automaton, with probability $1 - p_k$. At every stage $t = 0, 1, \dots$, every agent chooses c or d . The moves are made simultaneously and are observed after every stage. As usual, automata do not have payoffs. The stage payoffs of the players are described by

$$\begin{array}{cc} & c & d \\ c & m_1, m_2 & 0, g_2 \\ d & g_1, 0 & -l_1, -l_2 \end{array}$$

where $0 \leq m_k < g_k$ and $0 < l_k$, $k = 1, 2$.²¹ The stage game between the players is a game *with strictly competitive interests* in the sense of Cripps et al. (2005) and Atakan and Ekmekci (2013). Player k discounts streams of payoffs using a personal discount factor δ_k .²² We also assume that at every stage $t = 0, 1, \dots$, independently of past events, a random variable R_t is drawn according to a uniform distribution over $[0, 1]$ and is publicly revealed.²³ Finally, the behavior of the automata depends on the parameters i and γ . Let $j \neq i$, i.e., $j = i + 1 \pmod{2}$. At every stage $t = 0, 1, \dots$, automaton i plays d unless the next three conditions all hold²⁴

- agent j played c for the first time at an odd stage $s < t$
- at any stage $u \in \{s + 1, \dots, t - 1\}$, agent j played d if and only if $R_u < \gamma_j$,
- $R_t < \gamma_j$.

²¹The reputation game $\mathcal{G}_\delta(p, i, \gamma)$ is described by the following parameters: $\delta_k, p_k, m_k, g_k, l_k, \gamma_k, k = 1, 2$ and i . The notation just keeps track of the main ones. There is a slight abuse of notation in using a single matrix: the payoffs of a player are the same, whether he faces another player or an automaton.

²²In this section, we follow the tradition of the reputation literature and allow the players to use different discount factors.

²³As usual (see Peski (2013) for a recent reference), the public random device is introduced to simplify the exposition but is not necessary.

²⁴If one of the conditions does not hold, the automaton plays c .

In other words, automaton i may consent to an agreement that gives the average payoff $\gamma_j g_j$ to player j . If agent j did not initiate the agreement by playing c at an odd stage, or if the agreement was broken, automaton i always plays d .

The behavior of automaton j is similar, and is defined by inverting i and j and by replacing “odd” by “even” in the former description.

We go on by identifying specific strategies in $\mathcal{G}_\delta(p, i, \gamma)$, in which the players behave as in a *war of attrition*. We first define a *war of attrition protocol*, which depends on i and γ and also on a positive integer parameter T determining the horizon of the war.

War of attrition protocol: Before stage 0, every player chooses the first stage at which he will play c if the other agent did not play c yet. To this aim, player i (resp., player $j \neq i$) chooses $\tau_i \in \{0, 2, \dots, 2T\}$ (resp., $\tau_j \in \{1, 3, \dots, 2T + 1\}$) according to some probability distribution. If player i (resp., j) plays c for the first time at an even (resp., odd) stage $s \leq 2T + 1$, the players initiate an agreement in which player i (resp., j) chooses c most of the time. More precisely, if player k ($k = i, j$) is the one who gives in, then, at every stage $u \geq s + 1$, if $R_u \geq \gamma_k$, player k plays c and the other player plays d and if $R_u < \gamma_k$, player k plays d and the other player plays c .²⁵

The previous description tells what a player should do at every stage as long as he follows the protocol and faces either an automaton or a player who also follows the protocol. A complete description of the players’ *strategies*, which take account of all histories that are not consistent with the protocol, will be given in the appendix. In particular, if a player adopts a war of attrition strategy, he follows the war of attrition protocol as long as the other agent’s moves are consistent with it and behaves as his associated automaton otherwise.

The next three propositions give precise, easy to check, conditions that guarantee that $\mathcal{G}_\delta(p, i, \gamma)$ has a perfect Bayesian equilibrium in war of attrition strategies. The proofs are given in the appendix, together with the explicit expression of the equilibrium payoffs. The condition in the first proposition is particularly simple: it just says that the automata cooperate with

²⁵This latter part is similar to the standard construction to prove the folk theorem under complete information.

an arbitrary small positive probability after a player has first cooperated.

Proposition 4 *Let p and γ be fixed, with $\gamma_1, \gamma_2 > 0$. There exists $\underline{\delta} \in (0, 1)$ such that, for any $\delta_1 > \underline{\delta}$ and any $\delta_2 > \underline{\delta}$, $\mathcal{G}_\delta(p, i, \gamma)$ has a perfect Bayesian equilibrium in war of attrition strategies for some $i \in \{1, 2\}$.*

As many other results of the reputation literature (see Mailath and Samuelson (2006)), the following proposition makes use of the fact that a player may be significantly more patient than the other.

Proposition 5 *Fix p , γ and i . There exists $\underline{\delta} \in (0, 1)$ and $c > 0$ such that, for every $\delta_1, \delta_2 \in (0, 1)$ satisfying $\delta_i > \underline{\delta}$ and $1 - \delta_j \leq c(1 - \delta_i)$, $\mathcal{G}_\delta(p, i, \gamma)$ has a perfect Bayesian equilibrium in war of attrition strategies.*

The next proposition assumes that the players have the same discount factor and identifies a sufficient condition, condition (15) below, which is satisfied as soon as

- the players have the same payoffs ($g_1 = g_2, l_1 = l_2$)²⁶,
- automata always play d (i.e., $\gamma = 0$) and
- $p_i > p_j$ (which of course always holds, for $i =$ either 1 or 2, unless $p_1 = p_2$).

Proposition 6 *Consider the game with equal discount factors $\delta = \delta_1 = \delta_2$ and assume that p , γ and i satisfy*

$$\log(1 - p_i) \frac{l_i + \gamma_i g_i}{g_i} < \log(1 - p_j) \frac{l_j + \gamma_j g_j}{g_j}. \quad (15)$$

Then there exists $\underline{\delta} \in (0, 1)$ such that, for any $\delta > \underline{\delta}$, $\mathcal{G}_\delta(p, i, \gamma)$ has a perfect Bayesian equilibrium in war of attrition strategies.

²⁶As observed by Atakan and Ekmekci (2013), $g_1 = g_2$ can be achieved through a normalization and is thus without loss of generality.

Proposition 6 is consistent with Atakan and Ekmekci (2013)'s theorem 2. They specially study a reputation game in which the players of the basic game, with strictly competitive interests, have an arbitrary number of actions (as opposed to two, c and d , here) but they focus on very simple automata, which would always play d in our framework. In other words, they set $\gamma_1 = \gamma_2 = 0$, while we allow for $\gamma_1, \gamma_2 > 0$. Assuming that both players have the same discount factor δ , they show that if condition (15) holds with $\gamma_1 = \gamma_2 = 0$, then, when δ is sufficiently large, at *any* perfect Bayesian equilibrium of the reputation game, one of the players (player i)²⁷ gets approximately his minmax payoff 0 while the other player (player j) benefits from a reputation effect, namely gets a substantial part of g_j . As already suggested in the introduction, the main difference between our results and those of Atakan and Ekmekci (2013) is that, being interested in a reputation result, they establish the *asymptotic uniqueness* of a very special kind of equilibria, while we are interested in characterizing a larger class of equilibria, which *exist* for *every sufficiently large* discount factor. We will pursue the comparison in the appendix, when the expression of the equilibrium payoffs is available.

5.2 War of attrition equilibria in the public good game

Let us go back to the public good game $PG_\delta(p, \omega, z)$, in which there are two players and no automaton. When both players are normal, they play a game with strictly competitive interests in which (after subtracting 1 to all payoffs) $m_k = 0$, $g_k = \omega$, $l_k = 1 - \omega$, $k = 1, 2$. Building on the previous section, we consider specific strategies for the players of $PG_\delta(p, \omega, z)$, which we call *war of attrition strategies*, as in the previous reputation games. Let $i = 1$ or 2 , $j \neq i$ and $\gamma_k \in [0, \frac{1}{z}]$, $k = 1, 2$. In a *war of attrition strategy profile* parametrized by i and $\gamma = (\gamma_1, \gamma_2)$, player k behaves as automaton k in $\mathcal{G}_\delta(p, i, \gamma)$ if his type is greedy and as player k in $\mathcal{G}_\delta(p, i, \gamma)$ if his type is normal.²⁸

The next proposition states that equilibria in war of attrition strategies yield explicit elements in $\lim_{\delta \rightarrow 1} \mathcal{N} [PG_\delta(p, \omega, z)]$.

²⁷In the particular case $g_1 = g_2$, $l_1 = l_2$, agent i is the one who is the more likely to be a player rather than an automaton, namely, $p_i > p_j$.

²⁸Observe that given the description of automata, the behavior of greedy players is defined on every possible history. The behavior of normal players is fully described in the appendix.

Proposition 7 Consider the δ -discounted infinitely repeated public good game $PG_\delta(p, \omega, z)$. For every $\gamma_1, \gamma_2 \in (0, \frac{1}{z})$, there exists $\underline{\delta} \in (0, 1)$ such that, for every $\delta > \underline{\delta}$, $PG_\delta(p, \omega, z)$ has a perfect Bayesian equilibrium in war of attrition strategies parametrized by $i = 1$ or 2 and $\gamma = (\gamma_1, \gamma_2)$. Furthermore, if the priors are such that $p_1 \neq p_2$, the same holds for $\gamma_1 = \gamma_2 = 0$. The corresponding equilibrium payoffs converge, as $\delta \rightarrow 1$, to

$$\begin{aligned} x_i(n; \gamma) &= 1 + \gamma_i \omega, \\ x_j(n; \gamma) &= 1 + \pi_0(1 - \gamma_i)\omega + \pi_0 \gamma_j \omega, \\ x_i(g; \gamma) &= z + (1 - \gamma_j + \gamma_j(1 - z)) a_i, \\ x_j(g; \gamma) &= z + (1 - \gamma_i + \gamma_i(1 - z)) (\pi_0 + (1 - \pi_0)a_j), \end{aligned}$$

where $i = 1$ or 2 , $j \neq i$, $\pi_0 = 1 - \frac{1-p_i}{(1-p_j) \frac{1-\omega+\gamma_i\omega}{1-\omega+\gamma_j\omega}}$, $a_i = \frac{1-\omega+\gamma_i\omega}{1-\omega+(1-\gamma_j)\omega} \left(1 - (1-p_j) \frac{1-\omega+(1-\gamma_j)\omega}{1-\omega+\gamma_i\omega} \right)$, and $a_j = \frac{1-\omega+\gamma_j\omega}{1-\omega+(1-\gamma_i)\omega} \left(1 - (1-p_j) \frac{1-\omega+(1-\gamma_i)\omega}{1-\omega+\gamma_i\omega} \right)$. In particular, $\liminf_{\delta \rightarrow 1} \mathcal{N} [PG_\delta(p, \omega, z)]$ has a nonempty interior.

It may be useful to compare the latter result with our findings in the undiscounted case, namely, Propositions 1 and 2. Note that in a war of attrition equilibrium, once one of the players has revealed that his type is normal, the players follow a cooperative agreement, 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 agreement is not modified, of course).

However, the previous equilibria cannot be reduced to completely revealing ones (from the beginning of the game). 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 war of attrition equilibria. In the undiscounted game, a normal player is always better off waiting for the other player revealing his type, because waiting is free.

Note also that the previous equilibria are not finitely revealing in the sense of Peski (2013), namely, they involve a number of changes of the players' posteriors that increases with δ . However, we can refine our construction of equilibria by introducing, in the war of attrition, an arbitrary number of non-revealing stages between each stage of information transmission. It is then possible to keep the number of information transmission stages bounded as

$\delta \rightarrow 1$ and to obtain similar limit payoffs, so that $\liminf_{\delta \rightarrow 1} \mathcal{N} [PG_\delta(p, \omega, z)]$ contains a set (with a non-empty interior) of finitely revealing equilibrium payoffs (see the remark after the proof of Proposition 7 in the Appendix for details).

6 Appendix: proofs and complements

6.1 Undiscounted game (Proposition 1)

6.1.1 Strategies and payoff functions

A strategy of player i in $\Gamma_\infty(q)$ is a sequence of mappings $\sigma_i = (\sigma_i^t)_{t \geq 0}$, $\sigma_i^t : \Theta_i \times A^{t-1} \rightarrow \Delta(A_i)$. 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.

6.1.2 Sufficient conditions for an equilibrium

Let us assume that the conditions (4) 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 $v_j(\theta_j)$.

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 $\sigma_i(\theta_i)$ 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 (4) 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 $t_0 + 1, t_0 + 2, \dots$).²⁹

More precisely, the equilibrium strategies $\sigma_i(\theta_i)$ generate probability distributions $\mu_\sigma(\cdot | \theta_1, \dots, \theta_n)$ 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

$$x_i(\theta_i) = U_i(\theta_i, \sigma) = E_{q, \mu_\sigma}(u_i(\theta_i, \tilde{\pi}) | \theta_i) \quad \text{for every } i, \theta_i \quad (16)$$

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.³⁰

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

²⁹Note that player i may reveal further information on his type by playing so as to guarantee himself $v_i(\theta_i)$. This typically happens out of equilibrium.

³⁰If 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.

obtain that

$$x_i(\theta_i) \geq E_{q, \mu_\sigma} (\max \{u_i(\theta_i, \tilde{\pi}), v_i(\theta_i)\} \mid \theta'_i) \quad \text{for every } i, \theta_i, \theta'_i \quad (17)$$

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 $\theta' \in \Theta$, chooses a frequency of moves $\pi \in \Delta(A)$ according to $\mu_\sigma(\cdot \mid \theta')$ and recommends π to all players³¹. 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 . (16) says that by telling the truth and following the recommendation of the mediator, the players get the same *interim* payoff as by playing σ . (17) 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 (16) and (17) differ from (4) in two respects. (16) and (17) involve (type dependent) probability distributions over $\Delta(A)$, while (4) is formulated in terms of deterministic distributions $\pi(\theta)$, $\theta \in \Theta$. Moreover, in (17), the probability distribution μ_σ is not necessarily completely revealing³².

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

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

Hence, for every i, θ_i , (16) 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

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

³¹As 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}).

³²The above reliable mediator selects π as a random function of the players' reported types but does not reveal these reported types.

which is the first part of (4) if we set $\pi(\theta) = E_{\mu_\sigma}(\tilde{\pi} \mid \theta)$.

By proceeding similarly and using in addition that “max” is convex, for every i, θ_i, θ'_i , (17) 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 (4). ■

6.2 Generalization of the public good game, complements on Proposition 2

The main goal of this section is to formally establish inequalities (7). We actually consider a variant of the public good game $PG(p, \omega, z)$ of Section 2.2, in which every player has an extra action e , with the following payoffs:

		$\theta_2 = n$			$\theta_2 = g$			
		c	d	e	c	d	e	
$\theta_1 = n$	c	(1, 1)	(1, 1 + ω)	(1, 0)	c	(1, 1)	(1, 1 + z)	(1, 0)
	d	(1 + ω , 1)	(ω , ω)	(0, 0)	d	(1 + ω , 1)	(ω , z)	(0, 0)
	e	(0, 1)	(0, 0)	(0, 0)	e	(0, 0)	(0, z)	(0, 1 + z)
		c	d	e	c	d	e	
$\theta_1 = g$	c	(1, 1)	(1, 1 + ω)	(0, 0)	c	(1, 1)	(1, 1 + z)	(0, 0)
	d	(1 + z , 1)	(z , ω)	(z , 0)	d	(1 + z , 1)	(z , z)	(z , 0)
	e	(0, 1)	(0, 0)	(1 + z , 0)	e	(0, 0)	(0, z)	(1 + z , 1 + z)

In the new game, for every pair of types (θ_1, θ_2) , the set of feasible, individually rational payoffs of the public good game with complete information, in which θ_1 and θ_2 are common knowledge, has a nonempty interior, a property that is not satisfied when $\theta_1 = \theta_2 = g$ in the initial game.

In the new game,

- the ex post individual rationality levels are still $v_i(n) = 1$, $v_i(g) = z$, $i = 1, 2$
- uniform punishments still exist (play d)
- for every $(\theta_1, \theta_2) \neq (g, g)$, the set of feasible, individually rational payoffs is the same as in $PG(p, \omega, z)$. In particular, there exist feasible, strictly individually rational payoffs for every $(\theta_1, \theta_2) \neq (g, g)$. But now, there also exist strictly individually rational payoffs for $(\theta_1, \theta_2) = (g, g)$.

Let us show that, without discounting, if z , p_1 and p_2 are sufficiently large, there are no equilibria. As in Section 3.2, let $\pi(\cdot | \theta)$ be the equilibrium distributions over $A = \{c, d, e\} \times \{c, d, e\}$ when θ is reported and consider the equilibrium conditions in Proposition 1 (namely, (12)). We first prove that:

Claim: if z is sufficiently large, the feasibility and ex post individual rationality conditions imply that $\pi((d, c) | g, n) \geq 1 - \varepsilon$ and $\pi((c, d) | n, g) \geq 1 - \varepsilon$.

Proof of the claim: we establish the second inequality, the first one can be established in the same way.

$$\text{Let } \pi(\cdot | g, n) = \begin{pmatrix} \beta_c & \alpha & \gamma_{13} \\ \rho & \beta_d & \gamma_{23} \\ \gamma_{31} & \gamma_{32} & \gamma_{33} \end{pmatrix} \text{ and set } \gamma = \gamma_{23} + \gamma_{33}.$$

The ex post individual rationality condition for player 1, of greedy type, implies that

$$\rho z + \beta_d(z - 1) + \gamma z \geq z - 1 \quad (18)$$

(write the ex post individual rationality condition for the payoffs -1 and forget about some negative weights). Similarly, the ex post individual rationality condition for player 2, of normal type, implies that

$$\rho \omega + \beta_d + \gamma(\omega + 1) \leq \omega \quad (19)$$

(proceed as for player 1, but also replace α by $1 -$ all the other probabilities in $\pi(\cdot | g, n)$).

The feasibility constraints are $\rho \geq 0$, $\beta_d \geq 0$, $\gamma \geq 0$ and $\rho + \beta_d + \gamma \leq 1$.

If the latter constraint is binding, namely if $\rho + \beta_d + \gamma = 1$, then $\rho \omega + \beta_d + \gamma(\omega + 1) \geq \omega$, so that, by (19), $\rho \omega + \beta_d + \gamma(\omega + 1) = \omega$ and $\rho = 1$, which establishes the claim ((18) is also satisfied).

The constraint $\rho \geq 0$ cannot be binding as soon as $z \geq 1 + \frac{1}{\varepsilon}$ ($\Leftrightarrow \frac{z}{z-1} \leq \omega + 1$). Indeed, if $\rho = 0$, (18) and (19) imply then that $1 - \omega \leq \gamma \left(\frac{z}{z-1} - (\omega + 1) \right) \leq 0$ contradicting $\omega \leq 1$.

Hence the relevant extreme points are determined by the following binding constraints:

- (18), (19), $\beta_d = 0$: $\gamma = \frac{\omega}{z}$, $\rho = 1 - \frac{1+\omega}{z}$.
- (18), (19), $\gamma = 0$: $\rho = 1 - \frac{1}{z(1-\omega)+\omega}$.
- (18), $\beta_d = \gamma = 0$: $\rho = 1 - \frac{1}{z}$.
- (19), $\beta_d = \gamma = 0$: $\rho = 1$.

In all cases, the claim is satisfied. ■

If all γ_{kl} 's are set to 0, the proof of the claim establishes the inequalities (7) in the original public good game.

Furthermore, in the new game, the nondeviation condition (8) must be satisfied and feasibility still implies (9) and (10). Hence, by proceeding exactly as in Section 3.2, one shows that the set of Nash equilibrium payoffs is also empty in the new game.

6.3 Role of uniform punishments (Propositions 1 and 3)

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 (4). Examples 2 and 3 below illustrate how the absence of uniform punishments modifies the results.

In example 2, the conditions (4) 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 3, 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 (4). However, Proposition 3 still fails.

In both examples 2 and 3, 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,

$$\begin{aligned} x_1(\theta_1) &= u_1(\theta_1, \pi(\theta_1)) \\ &\geq u_1(\theta_1, \pi(\theta'_1)) \quad \forall \theta_1, \theta'_1 \in \Theta_1 \quad \text{i.e., incentive compatibility} \end{aligned} \quad (20)$$

$$\geq v_1(\theta_1) \quad \forall \theta_1 \in \Theta_1 \quad \text{i.e., ex post individual rationality} \quad (21)$$

and, for player 2,

$$x_2 = u_2(\pi(\theta_1)) \geq v_2 \quad \forall \theta_1 \in \Theta_1 \quad \text{i.e., ex post individual rationality} \quad (22)$$

As shown by Hart (1985), in order to characterize the equilibrium payoffs of $\Gamma_\infty(q)$, *ex post* individual rationality (namely, (6) or (21) 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.³³ 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

$$\forall p_1 \in \Delta(\Theta_1), \quad \sum_{\theta_1} p_1(\theta_1) x_1(\theta_1) \geq val_1 \left[\sum_{\theta_1} p_1(\theta_1) u_1(\theta_1, \cdot) \right] \quad (23)$$

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 (23) 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, $EXPIR_1[\Gamma(q)]$

³³The same condition holds as well in two-person games with independent private values (see Koren (1992)).

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, (13) implies (23). Furthermore, (23) holds in particular at the extreme points of $\Delta(\Theta_1)$, so that it implies *ex post* individual rationality (*i.e.*, (21)). 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 0, under the assumption of uniform punishments, the three sets coincide. In examples 2 and 3 below, this assumption does not hold. In example 2, the two inclusions are strict. In example 3, the first inclusion is strict but $INTIR_1[\Gamma_\infty(q)] = EXPIR_1[\Gamma(q)]$.

Example 2

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 (23)). 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 (23), 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 (13): 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 (20), (21) and (22) 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 (23). Hence π does not define an equilibrium of $\Gamma_\infty(q)$.

Example 2 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 3

The framework is the same as in example 2 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 (23) if and only if it is *ex post* individually rational (namely, (21): $x_1(h) \geq 1$ 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 (20), (21) and (22).³⁴

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 (13): 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 2 and 3, *interim* individual rationality takes a different form in the one-shot game and in the infinitely repeated game. In example 2, in order to defend himself, player 1 must play in a non-revealing way in the repeated game. In example 3, 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)).

6.4 Reputation games (Propositions 4, 5 and 6)

Propositions 4, 5 and 6 state that, under appropriate assumptions, the reputation game $\mathcal{G}_\delta(p, i, \gamma)$ has a perfect Bayesian equilibrium in war of attrition strategies, which are based on a war of attrition protocol. We start by giving a full description of war of attrition strategies. Recall that, in a war of attrition protocol, player i (resp., player $j \neq i$) chooses the first stage $\tau_i \in \{0, 2, \dots, 2T\}$ (resp., $\tau_j \in \{1, 3, \dots, 2T + 1\}$) at which he will play c if the other agent did not play c yet. If player i (resp., j) plays c for the first time at an even (resp., odd) stage $s \leq 2T + 1$, the players initiate an agreement in which player i (resp., j) chooses c with frequency $1 - \gamma_i$ (resp.,

³⁴The 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).

$1 - \gamma_j$). A player can observe three kinds of histories that are inconsistent with the protocol

I c has been played for the first time at a “wrong” stage, namely,

I' by agent i at an odd stage t or by agent j at an even stage t ($t < 2T + 2$ or $t \geq 2T + 2$),

I'' by player i only at an even stage t or by player j only at an odd stage t , with $t \geq 2T + 2$;

II c is has been played for the first time at a “right” stage but the agreement described by γ_i or γ_j is not followed by some player at some stage;

III the stage is $\geq 2T + 2$ and c has never been played.

War of attrition strategies: A war of attrition strategy for player k ($k = 1, 2$) in $\mathcal{G}_\delta(p, i, \gamma)$ is defined from a war of attrition protocol (which depends on the parameters of the game, T and appropriate probability distributions) as follows: at every stage $t = 0, 1, \dots$,

- if the history up to stage $t - 1$ is consistent with the protocol, play according to it.

Otherwise,

- on type **I'** histories, if c has been played for the first time by player i (resp., j) at an odd (resp., even) stage, player i (resp., j) chooses c and player j (resp., i) chooses d .
- on type **I''** histories, players initiate an agreement, exactly as in the protocol before $2T + 1$.
- on type **II** histories, if player i (resp., only player j) did not follow the agreement, player i (resp., j) chooses c and player j (resp., i) chooses d from then on.
- on type **III** histories, player i (resp., j) plays a best response to automaton j (resp., i) (and the history becomes of **type I** after one or two stages).

Propositions 4, 5 and 6 will result from Lemma 8 below. To prepare for it, let us set ε_1 and v_2 such that (ε_1, v_2) is the mean stage payoff in the sequence of actions that begins when player 1 concedes first, i.e. $\varepsilon_1 = \delta_1 \gamma_1 g_1$ and $v_2 = (1 - \delta_2)g_2 + \delta_2(1 - \gamma_1)g_2$. We define ε_1 and v_2 symmetrically. Note that $\frac{v_1 - \varepsilon_1}{g_1} = 1 - \delta_1(\gamma_1 + \gamma_2)$ and $\frac{v_2 - \varepsilon_2}{g_2} = 1 - \delta_2(\gamma_1 + \gamma_2)$, so that the fact that $\gamma_1 + \gamma_2 \leq 1$ implies that $\varepsilon_1 < v_1$ and $\varepsilon_2 < v_2$. We then set

$$\alpha_1 = \frac{(1 - \delta_1^2)(l_1 + \varepsilon_1)}{\delta_1(l_1 + v_1) - \delta_1^2(l_1 + \varepsilon_1)},$$

and

$$T_1 = \left\lfloor \frac{\log(1 - p_2)}{\log(1 - \alpha_1)} \right\rfloor,$$

with the convention that $T_1 = 0$ if $\alpha_1 \geq 1$.

We also define α_2 and T_2 , by inverting the roles of player 1 and player 2.

Lemma 8 *Assume that the following conditions hold:*

$$T_i \leq T_j, \tag{24}$$

$$(1 - \delta_i)(\alpha_i \delta_i m_i - l_i) \leq \varepsilon_i, \tag{25}$$

$$(1 - \delta_j)(\alpha_j \delta_j m_j - l_j) \leq \varepsilon_j, \tag{26}$$

$$\frac{1 - p_i}{(1 - \alpha_j)^{T_i}} \leq \frac{v_j - (1 - \delta_j)m_j}{v_j + l_j - \delta_j(l_j + \varepsilon_j) - (1 - \delta_j)m_j}. \tag{27}$$

Then the reputation game $\mathcal{G}_\delta(p, i, \gamma)$ has a perfect Bayesian equilibrium in war of attrition strategies with payoffs

$$\begin{aligned} x_i &= \varepsilon_i, \\ x_j &= \left(1 - \frac{1 - p_i}{(1 - \alpha_j)^{T_i}}\right) v_j + \frac{1 - p_i}{(1 - \alpha_j)^{T_i}} (-l_j + \delta_j(l_j + \varepsilon_j)). \end{aligned}$$

Before establishing the lemma, it may be useful to understand how it applies in a particular case where the automata always play d ($\gamma = 0$), while the players have the same discount factor $\delta_k = \delta$ and the same payoffs $g_k = g$, $l_k = l$, $m_k = 0$, $k = 1, 2$. In this case, we can set $\alpha = \alpha_k$, $k = 1, 2$; $\alpha \leq 1$ as soon as $\delta \geq \frac{l}{l+g}$. Condition (24), which fixes the index i , is satisfied if agent i is more likely to be a player than agent j (namely, i is such that $p_i \geq p_j$).

Conditions (25) and (26) are satisfied, for every δ , since $m_k = 0$, $k = 1, 2$. The whole point is to guarantee (27), which reduces to

$$\frac{1 - p_i}{(1 - \alpha)^{T_i}} \leq \frac{g}{g + (1 - \delta)l}$$

The right hand side is < 1 for every $\delta < 1$ and can be made arbitrarily close to 1 by choosing δ sufficiently close to 1. The left hand side is $\leq \frac{1 - p_i}{1 - p_j}$ by definition of T_i . Hence, if $p_i > p_j$ (i.e., if $p_1 \neq p_2$ and index i is chosen to satisfy (24)), (27) is satisfied for δ sufficiently large.

Proof:

In order to show that the strategies described above define an equilibrium, we first construct the unique distributions for τ_i and τ_j that make player i (resp. j) indifferent between playing c for the first time at stage $2t$ (resp. $2t + 1$) for any $t \in \{0, 1, \dots, T\}$. Then we check that the derived strategic profile is indeed an equilibrium. We also show that T necessarily equals T_i . We adopt the convention that $\tau_i = +\infty$ if agent i is an automaton. 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).

Note also that $\mathbb{P}(\tau_j = t) = 0$ if t is an even number or if $t > 2T + 1$, and $\mathbb{P}(\tau_j = +\infty) = 1 - p_j$. Similarly, $\mathbb{P}(\tau_i = t) = 0$ if t is an odd number or if $t > 2T$, and $\mathbb{P}(\tau_i = +\infty) = 1 - p_i$.

Probability distribution of τ_j

The fact that player i has to be indifferent between starting contributing at any stage $t = 0, 2, \dots, 2T$ will enable us to find the entire distribution of τ_j . We denote by $G_i(t)$ the overall payoff of player i when he starts to contribute at stage t , the law of τ_j being fixed with the former restrictions. Note that $G_i(2t)$ refers to the expected payoff of the strategy that dictates to start contributing at stage $2t$ if player j did not contribute before and that initiate an agreement. $G_i(2t + 1)$ is the expected payoff of an arbitrary strategy which is only required to start contributing at stage $2t + 1$ if agent j did not contribute before.

We have, for any $t \geq 0$:

$$G_i(2t) = \sum_{s=0}^{t-1} \mathbb{P}(\tau_j = 2s+1)(-l_i + \delta_i^{2s+1}(l_i + v_i)) + \mathbb{P}(\tau_j > 2t-1)(-l_i + \delta_i^{2t}(l_i + \varepsilon_i)).$$

In particular, $G_i(0) = \varepsilon_i$.

So, the law of τ_j has to be such that $\varepsilon_i = G_i(2) = \dots = G_i(2T)$. From equality $G_i(2) = \varepsilon_i$, we get $\mathbb{P}(\tau_j = 1) = \alpha_i$.

Note that if $\alpha_i > p_j$ (which is equivalent to $T_i = 0$), the latter equality can not hold. In this case, $T = T_i = 0$ and $\mathbb{P}(\tau_j = 1) = p_j$.

Now, let us show by induction that:

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

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 $s = t$. From the definition of G_i , one has:

$$\begin{aligned} & G_i(2t+2) - G_i(2t) \\ = & \mathbb{P}(\tau_j = 2t+1) (-l_i + \delta_i^{2t+1}(l_i + v_i)) \\ & + \mathbb{P}(\tau_j > 2t+1) (-l_i + \delta_i^{2t+2}(l_i + \varepsilon_i)) \\ & - \mathbb{P}(\tau_j > 2t-1) (-l_i + \delta_i^{2t}(l_i + \varepsilon_i)) \\ = & \mathbb{P}(\tau_j = 2t+1) (-l_i + \delta_i^{2t+1}(l_i + v_i)) \\ & + (\mathbb{P}(\tau_j > 2t-1) - \mathbb{P}(\tau_j = 2t+1)) (-l_i + \delta_i^{2t+2}(l_i + \varepsilon_i)) \\ & - \mathbb{P}(\tau_j > 2t-1) (-l_i + \delta_i^{2t}(l_i + \varepsilon_i)) \\ = & \mathbb{P}(\tau_j = 2t+1) (l_i + v_i - \delta_i(l_i + \varepsilon_i)) \delta_i^{2t+1} \\ & - \mathbb{P}(\tau_j > 2t-1) (l_i + \varepsilon_i) (1 - \delta_i^2) \delta_i^{2t}. \end{aligned}$$

As

$$\begin{aligned} \mathbb{P}(\tau_j > 2t-1) &= 1 - \sum_{s=0}^{t-1} \mathbb{P}(\tau_j = 2s+1) \\ &= 1 - \sum_{s=0}^{t-1} (1 - \alpha_i)^s \mathbb{P}(\tau_j = 1) \\ &= 1 - \alpha_i \frac{1 - (1 - \alpha_i)^t}{1 - (1 - \alpha_i)} \\ &= (1 - \alpha_i)^t, \end{aligned} \tag{28}$$

one finally has

$$\begin{aligned} & G_i(2t+2) - G_i(2t) \\ &= (l_i + v_i - \delta_i(l_i + \varepsilon_i)) \delta_i^{2t+1} [\mathbb{P}(\tau_j = 2t+1) - \alpha_i(1 - \alpha_i)^t]. \end{aligned} \tag{29}$$

The fact that $G_i(2t+2) = G_i(2t)$ then implies that

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

which ends the induction.

Our construction must ensure that player i does not prefer conceding at stage $2T+2$, i.e. $G_i(2T+2) \leq G_i(2T)$. In fact, we construct an equilibrium such that this preference is strict, i.e. $G_i(2T+2) < G_i(2T)$. Because equality (29) still holds for $t = T$ this is equivalent to

$$\mathbb{P}(\tau_j = 2T+1) < \alpha_i(1 - \alpha_i)^T. \tag{30}$$

Moreover $\mathbb{P}(\tau_j \leq 2T+1) = p_j$, and we have

$$\sum_{t=0}^{T-1} \alpha_i(1 - \alpha_i)^t = 1 - (1 - \alpha_i)^T \leq p_j < \sum_{t=0}^T \alpha_i(1 - \alpha_i)^t = 1 - (1 - \alpha_i)^{T+1}.$$

This implies that $T = T_i$.

Probability distribution of τ_i

Here, the point is that player j has to be indifferent between conceding at stages $t = 1, 3, \dots, 2T+1$. Let us denote $G_j(t)$ the overall payoff of player j when he is willing to concede at stage t , the law of τ_i being fixed. Again, " $2t+1$ " refers to a regular war of attrition strategy, whereas " $2t$ " refers to an arbitrary strategy that dictates to contribute at stage $2t$ if i did not contribute before. We have, for any $t \geq 0$:

$$G_j(2t+1) = \sum_{s=0}^t \mathbb{P}(\tau_i = 2s)(-l_j + \delta_j^{2s}(l_j + v_j)) + \mathbb{P}(\tau_i > 2t)(-l_j + \delta_j^{2t+1}(l_j + \varepsilon_j)).$$

The law of τ_i has to be such that $G_j(1) = \dots = G_j(2T_i+1)$. By the same means as in the former case, the fact that $G_j(1) = G_j(3)$ implies that

$$\mathbb{P}(\tau_i = 2) = \alpha_j(1 - \mathbb{P}(\tau_i = 0)).$$

Again, this equality can not hold if $\alpha_j(1 - \mathbb{P}(\tau_i = 0)) > p_i$, but this can not happen. Indeed, if $\alpha_j > p_i$ then $T_j = 0$, and $T_i = 0$ because of (24). As mentioned before, we then have $T = 0$. Thus $\mathbb{P}(\tau_i = 0) = p_i$ (which is consistent with equality (31) below) and $\mathbb{P}(\tau_i = 2t) = 0$ for any $t \geq 1$.

One can prove by induction that:

$$\forall t \in \{1, \dots, T\}, \mathbb{P}(\tau_i = 2t) = (1 - \alpha_j)^{t-1} \mathbb{P}(\tau_i = 2)$$

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

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

and to

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

so that

$$\mathbb{P}(\tau_i = 0) = 1 - \frac{1 - p_i}{(1 - \alpha_j)^{T_i}}. \quad (31)$$

Thus, we define a proper distribution of τ_i , because then $\mathbb{P}(\tau_i = 0) < 1$ and

$$\mathbb{P}(\tau_i = 0) \geq 1 - \frac{1 - p_i}{(1 - \alpha_j)^{T_j}} \geq 1 - \frac{1 - p_i}{(1 - \alpha_j)^{\frac{\log(1-p_i)}{\log(1-\alpha_j)}}} = 0.$$

Type dependent probabilities defining the equilibrium strategies

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

$$\beta_0^i = \frac{\mathbb{P}(\tau_i = 0)}{p_i} = \frac{1}{p_i} \left(1 - \frac{1 - p_i}{(1 - \alpha_j)^{T_i}} \right)$$

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

$$\begin{aligned} \beta_{2t}^i &= \frac{\mathbb{P}(\tau_i = 2t)}{p_i} = \frac{1}{p_i} (1 - \alpha_j)^{t-1} \mathbb{P}(\tau_i = 2) \\ &= \frac{1}{p_i} (1 - \alpha_j)^{t-1} \alpha_j (1 - \mathbb{P}(\tau_i = 0)) \\ &= \frac{1 - p_i}{p_i} (1 - \alpha_j)^{t-1-T_i} \alpha_j, \end{aligned}$$

Player j chooses stages $2t + 1$ ($0 \leq t \leq T_i - 1$) with probability

$$\beta_{2t+1}^j = \frac{\mathbb{P}(\tau_j = 2t + 1)}{p_j} = \frac{1}{p_j} (1 - \alpha_i)^t \alpha_i,$$

and stage $2T_i + 1$ with probability

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

Equilibrium conditions

Let us check that this strategic profile is indeed an equilibrium.

First, any player $k \in \{1, 2\}$ is right to follow an agreement that was previously initiated by himself or by the other player. Indeed, if he plays c instead of d , he will get m_k instead of g_k and if he plays d instead of c he will get $-l_k$ instead of 0. For the rest of the stages, he will be punished and win at most 0, which is less than what the agreement would have given him. Note that this covers the case of \mathbf{I}'' histories.

Now, let us show that player i 's strategy is a best response to j 's one. We have $G_i(2T + 2) \leq G_i(2T)$ by construction, meaning that player i has no interest in starting contributing at stage $2T + 2$. He obviously does not prefer

to start contributing at stage t , $t > 2T + 2$, because if j did not contribute before stage $2T + 1$, j is necessarily an automaton and then there is no more reason for i to wait for starting contributing.

Let us now show that he had not better start contributing at odd stages $2t + 1$, $t \in \{0, 1, \dots, T\}$. We compare the payoff of an arbitrary strategy that starts contributing $2t + 1$ to the regular pure strategy that starts contributing at stage $2t$. Note that starting contributing at stage $2t + 1$ is punished by agent j . We have:

$$\begin{aligned} G_i(2t + 1) - G_i(2t) &\leq \mathbb{P}(\tau_j \leq 2t) \times 0 \\ &\quad + \mathbb{P}(\tau_j \geq 2t + 1) [-l_i + \delta_i^{2t+1}l_i - (-l_i + \delta_i^{2t}(l_i + \varepsilon_i))] \\ &\quad + \mathbb{P}(\tau_j = 2t + 1)(1 - \delta_i)\delta_i^{2t+1}m_i \\ &\leq \mathbb{P}(\tau_j \geq 2t + 1)\delta_i^{2t} [\delta_i l_i - (l_i + \varepsilon_i)] \\ &\quad + \mathbb{P}(\tau_j = 2t + 1)(1 - \delta_i)\delta_i^{2t+1}m_i \end{aligned}$$

For any $t \in \{0, \dots, T - 1\}$, $\mathbb{P}(\tau_j = 2t + 1) = \alpha_i (1 - \alpha_i)^t$, and $\mathbb{P}(\tau_j = 2T + 1) \leq \alpha_i (1 - \alpha_i)^T$ (cf. (30)), so that:

$$\begin{aligned} G_i(2t + 1) - G_i(2t) &\leq (1 - \alpha_i)^t \delta_i^{2t} [\delta_i l_i - (l_i + \varepsilon_i)] \\ &\quad + \alpha_i (1 - \alpha_i)^t (1 - \delta_i) \delta_i^{2t+1} m_i \\ &\leq (1 - \alpha_i)^t \delta_i^{2t} [\delta_i l_i - (l_i + \varepsilon_i) + \alpha_i \delta_i (1 - \delta_i) m_i], \end{aligned}$$

for any $t \in \{0, 1, \dots, T\}$. By means of inequality (25), we then have $G_i(2t + 1) \leq G_i(2t)$.

Finally, player i is indifferent between $0, 2, \dots, 2T$ by construction, so his strategy is a best response.

Similar arguments show that player j 's strategy is a best response to i 's one. In particular, one can show that

$$G_i(2t+2) - G_i(2t+1) \leq (1 - \mathbb{P}(\tau_i = 0))(1 - \alpha_j)^t \delta_j^{2t+1} [\delta_j l_j - (l_j + \varepsilon_j) + \alpha_j \delta_j (1 - \delta_j) m_j],$$

for any $t \in \{0, 1, \dots, T\}$. So $G_i(2t + 2) \leq G_i(2t + 1)$ by inequality (26).

Yet, it remains to show that player j should not concede at stage 0, i.e.:

$$G_j(1) \geq \mathbb{P}(\tau_i = 0)(1 - \delta_j)m_j + (1 - \mathbb{P}(\tau_i = 0)) \times 0.$$

One can check that this is equivalent to inequality (27).

Equilibrium payoffs

One has $x_i(n; \gamma) = G_i(0)$ and $x_j(n; \gamma) = G_j(1)$, the expressions in the lemma follow easily.

Bayesian perfectness

We now check that the players' strategies are best responses to each other, for appropriate beliefs, on histories that are not consistent with the war of attrition protocol. The case of **I'** histories has already been considered above.

On **I'** and **II** histories, the player, say k , who is supposed to play c has revealed earlier that he is not an automaton, which determines the belief of the other player, $\ell \neq k$. Given this belief and the fact that player k plays c , it is a best reply for player ℓ to play d . For player k , whatever his beliefs on agent ℓ (which may be a player or an automaton), playing c is a best response to d .

Finally, if a normal player gets to see a **III** history, then the other agent is an automaton with probability 1. It is then optimal for this normal player to react accordingly. ■

Remark: The previous proof is written in terms of the players ex ante, overall payoffs. Starting with Kreps and Wilson (1982), the reputation literature rather considers conditional expected payoffs at every stage. We can as well construct our war of attrition equilibria inductively. By doing so, we find that player i is indifferent between playing c first at a stage $2t$ and waiting two more stages (hoping that agent j will concede at stage $2t + 1$) if and only if the probability that agent j plays c at stage $2t + 1$ given past history is α_i . That determines strategy of player j at stages $2t + 1$, $t \in \{0, \dots, T - 1\}$: play c with probability $\frac{\alpha_i}{p_j^{2t+1}}$ and d otherwise, where p_j^{2t+1} is the probability that player i assigns to j being a player (as opposed to an automaton) given that c has not been played before stage $2t + 1$. The same holds for player i 's strategy at stages $2t$, $t \in \{1, \dots, T\}$: play c with probability $\frac{\alpha_j}{p_i^{2t}}$ and d otherwise, where p_i^{2t} is the probability that player j assigns to i being a player given that c has not been played before stage $2t$. For the construction to be complete, it remains to see that:

- if c has never been played before stage $2T + 1$, player j plays c with probability 1 at stage $2T + 1$. Player i must then prefer c to d at stage

$2T$ if c has not been played before, and we impose the preference to be strict, as in the former proof. This is equivalent to inequality $\frac{\alpha_i}{p_j^{2T+1}} > 1$.

- if c has never been played before stage $2T$, player i plays c with probability 1 at stage $2T$. Player j must then be indifferent between c to d at stage $2T$ if c has not been played before. This is equivalent to inequality $\frac{\alpha_j}{p_i^{2T}} = 1$.

These two constraints enable us to determine T and $\mathbb{P}(\tau_i = 0)$.

We will now prove Propositions 4, 5 and 6. They all consists in giving simple sufficient condition for inequalities (24), (25), (26), (27), so that Lemma 8 holds.

Proof of Proposition 4 As $\gamma_k > 0$, one has $\delta_k(l_k + \varepsilon_k) > l_k$ as long as δ_k is close enough to 1. Thus we can choose $\underline{\delta}$ close enough to 1 so that, for any $\delta_1, \delta_2 > \underline{\delta}$, one has $\delta_1(l_1 + \varepsilon_1) > l_1$ and $\delta_2(l_2 + \varepsilon_2) > l_2$, and so that Inequalities (25) and (26) hold. Then, for any of these δ_1, δ_2 , one can choose $i \in \{1, 2\}$ such that $T_i \leq T_j$. The left hand side of inequality (27) is lower than 1:

$$\frac{1 - p_i}{(1 - \alpha_j)^{T_i}} \leq \frac{1 - p_i}{(1 - \alpha_j)^{T_j}} \leq 1,$$

whereas the inequality $\delta_j(l_j + \varepsilon_j) > l_j$ ensures that the right hand side is greater than 1. Therefore, inequalities (24), (25), (26), (27) of Lemma 8 all hold. \blacksquare

Proof of Proposition 5 One can check that $\alpha_k \underset{\delta_k \rightarrow 1}{\sim} (1 - \delta_k) \frac{2(l_k + \gamma_k g_k)}{g_k(1 - \gamma_1 - \gamma_2)}$ for any $k \in \{1, 2\}$, so that

$$T_k = \frac{\beta_k}{1 - \delta_k} + \frac{\eta_k(1 - \delta_k)}{1 - \delta_k}, \quad (32)$$

where $\beta_k = -\log(1 - p_l) \frac{g_k(1 - \gamma_1 - \gamma_2)}{2(l_k + \gamma_k g_k)}$ (with $l \in \{1, 2\}$, $l \neq k$) and η_k is a function such that $\eta_k(x) \xrightarrow{x \rightarrow 0} 0$.

Fix $c = \frac{1}{2} \frac{\beta_j}{\beta_i}$. If $1 - \delta_j \leq c(1 - \delta_i)$, one has:

$$\begin{aligned}
T_i &= \frac{\beta_i}{1 - \delta_i} + \frac{\eta_i(1 - \delta_i)}{1 - \delta_i} \leq \frac{c\beta_i}{1 - \delta_j} + \frac{c\eta_i(1 - \delta_i)}{1 - \delta_j} \\
&= \frac{1}{2} \frac{\beta_j}{1 - \delta_j} + \frac{c\eta_i(1 - \delta_i)}{1 - \delta_j} \\
&\leq T_j - \frac{1}{2} \frac{\beta_j}{1 - \delta_j} - \frac{\eta_j(1 - \delta_j)}{1 - \delta_j} + \frac{c\eta_i(1 - \delta_i)}{1 - \delta_j} \\
&= T_j - \frac{1}{1 - \delta_j} \left(\frac{\beta_j}{2} + \eta_j(1 - \delta_j) - c\eta_i(1 - \delta_i) \right).
\end{aligned}$$

Moreover, if $1 - \delta_j \leq c(1 - \delta_i)$, δ_j goes to 1 when δ_i goes to 1, so that $\eta_j(1 - \delta_j)$ and $\eta_i(1 - \delta_i)$ are arbitrary small if δ_i is close enough to 1. This ensures that $T_i < T_j$ if δ_i is close enough to 1.

Then we can find a simple upper bound for the right hand side of inequality (27):

$$\frac{1 - p_i}{(1 - \alpha_j)^{T_i}} \leq \frac{1 - p_i}{(1 - \alpha_j)^{T_j - 1}} = (1 - \alpha_j) \frac{1 - p_i}{(1 - \alpha_j)^{T_j}} \leq 1 - \alpha_j.$$

If $\gamma_j > 0$, the left hand side of inequality (27) tends to $\frac{1 - \gamma_i}{1 - \gamma_i - \gamma_j} > 1$ as δ_j goes to 1, so that (27) holds as long as δ_j is close enough to 1.

If $\gamma_j = 0$, one can check that the left hand side of inequality (27) equals $1 - \frac{l_j}{g_j(1 - \gamma_i)}(1 - \delta_j) + o(1 - \delta_j)$, whereas $1 - \alpha_j = 1 - 2 \frac{l_j}{g_j(1 - \gamma_i)}(1 - \delta_j) + o(1 - \delta_j)$. Thus, (27) is always valid as long as δ_j is close enough to 1.

Finally, inequalities (25) and (26) are always valid when δ_i and δ_j are close enough to 1. This means that if $1 - \delta_j \leq c(1 - \delta_i)$ and if δ_i is close enough to 1 (i.e. $\delta_i \geq \underline{\delta}$ for some $\underline{\delta} \in \{0, 1\}$), then δ_j is also close to 1, and the four equations of Lemma 8 all hold. \blacksquare

Proof of Proposition 6 This proof is similar to the former proof of Proposition 5. Equation (32) still holds, and the inequality $\log(1 - p_i) \frac{l_i + \gamma_i g_i}{g_i} < \log(1 - p_j) \frac{l_j + \gamma_j g_j}{g_j}$ is equivalent to $\beta_i < \beta_j$. We have

$$T_i = T_j - \frac{1}{1 - \delta} (\beta_j - \beta_i + \eta_j(1 - \delta) - \eta_i(1 - \delta)),$$

so that $T_i < T_j$ if δ is close enough to 1. The rest follows as in the proof of Proposition 5. \blacksquare

Further elements of comparison between Proposition 6 and Atakan and Ekmekci (2013)'s theorem 2

We can now make more precise the comparison that we started at the end of Section 5.1. In their theorem 2, Atakan and Ekmekci (2013) assume without loss of generality that the basic payoffs satisfy $g_1 = g_2$ and focus on the case $\gamma = 0$. Under the latter assumption, both automata always play d , but the war of attrition strategies (described at the beginning of Section 6.4) still identify one of the players, player i , as the one who can start to cooperate at even stages (in particular at stage 0), while the other (player $j \neq i$) can start to cooperate at odd stages. This construction has no counterpart in Atakan and Ekmekci (2013) who rely on an auxiliary war of attrition in continuous time. When $g_1 = g_2 = g$ and $\gamma = 0$, condition (15) in Proposition 6, which identifies player i , reduces to $l_i \log(1-p_i) < l_j \log(1-p_j)$. Furthermore, in this case, by proceeding as in the proofs of Propositions 5 and 6, one can show that, when $\delta \rightarrow 1$, the players' equilibrium payoffs, given in Lemma 8, converge to

$$\begin{aligned} x_i &= 0 \\ x_j &= \lim_{\delta \rightarrow 1} \left(1 - \frac{1-p_i}{(1-\alpha_j)^{T_i}} \right) g = \left[1 - (1-p_i) \left(\frac{1}{1-p_j} \right)^{\frac{l_j}{l_i}} \right] g. \end{aligned}$$

These payoffs are exactly the same as the ones that are implicitly obtained in Atakan and Ekmekci (2013)'s theorem 2, as the limit payoffs of any perfect Bayesian equilibrium of the reputation game with $\gamma = 0$. Proposition 6 is not only consistent with Atakan and Ekmekci (2013)'s result but also provides an explicit construction of a larger family of equilibria (namely, with $\gamma > 0$), for any sufficiently large discount factor, in the original discrete time framework.³⁵

6.5 Public good game (Proposition 7)

Proposition 7 results from the following lemma.

³⁵There is an additional subtle point. As observed above, inequality (15) in Proposition 6 must be strict, which excludes for instance the case $p_1 = p_2$ and $l_1 = l_2$. No such restriction is needed in Atakan and Ekmekci (2013)'s continuous time framework. We can recover their result in the border case by deriving the equilibrium payoffs for $\gamma_k > 0$ and letting $\gamma_k \rightarrow 0$, $k = 1, 2$.

Lemma 9 *Let us assume that the δ -discounted infinitely repeated public good game $PG_\delta(p, \omega, z)$ and $\gamma = (\gamma_1, \gamma_2)$ satisfy the following conditions:*

$$\gamma_k \in \left[0, \frac{1}{z}\right), \quad \delta \geq \frac{z-1}{z-\gamma_k z} \quad k = 1, 2 \quad (33)$$

and inequalities (24) and (27) with $m_k = 0$, $g_k = \omega$, $l_k = 1 - \omega$, $\delta_1 = \delta_2 = \delta$, for $i = 1$ or 2 and $j \neq i$. Then $PG_\delta(p, \omega, z)$ has a perfect Bayesian equilibrium in war of attrition strategies parametrized by i and γ . The corresponding equilibrium payoffs converge, as $\delta \rightarrow 1$, to

$$\begin{aligned} x_i(n; \gamma) &= 1 + \gamma_i \omega, \\ x_j(n; \gamma) &= 1 + \pi_0(1 - \gamma_i)\omega + \pi_0 \gamma_j \omega, \\ x_i(g; \gamma) &= z + (1 - \gamma_j + \gamma_j(1 - z)) a_i, \\ x_j(g; \gamma) &= z + (1 - \gamma_i + \gamma_i(1 - z)) (\pi_0 + (1 - \pi_0)a_j), \end{aligned}$$

where $\pi_0 = 1 - \frac{1-p_i}{(1-p_j) \frac{1-\omega+\gamma_i\omega}{1-\omega+\gamma_i\omega}}$, $a_i = \frac{1-\omega+\gamma_i\omega}{1-\omega+(1-\gamma_j)\omega} \left(1 - (1-p_j) \frac{1-\omega+(1-\gamma_j)\omega}{1-\omega+\gamma_i\omega}\right)$, and $a_j = \frac{1-\omega+\gamma_j\omega}{1-\omega+(1-\gamma_i)\omega} \left(1 - (1-p_j) \frac{1-\omega+(1-\gamma_i)\omega}{1-\omega+\gamma_i\omega}\right)$.

Proof:

The situation of normal players i and j in $PG_\delta(p, \omega, z)$ is the same as in the reputation game $\mathcal{G}_\delta(p, i, \gamma)$ with $m_k = 0$, $g_k = \omega$, $l_k = 1 - \omega$ because a greedy player behaves as an automaton. Thanks to Lemma 8, we mostly need to check that playing as an automaton is indeed a best response for a greedy player. A greedy player k has basically four ways to deviate from the strategy of an automaton.

- Playing c first, when not supposed to (namely, at an odd (resp., even) stage if $k = i$ (resp., $k = j$)): the player will gain 1 and will then be punished and gain at most z for the remaining stages, whereas he could have won at least z all along by simply always playing d .
- Playing c first, when supposed to: this will be interpreted as the initiation of an agreement by the other player. Depending on whether and when he will break this agreement, the player will first gain 1 and then his average gain will be between $(1 - \gamma_k)1 + \gamma_k(z + 1)$ and z . As $(1 - \gamma_k)1 + \gamma_k(z + 1) \leq 2$, this is again beaten by always playing d .

- Breaking an agreement initiated by the other player by playing c instead of d : he will get 1 and then will be punished and win at most z for all the remaining stages, which is always less than he would have got by always playing d .
- Breaking an agreement initiated by the other player by playing d instead of c : he will win z and then be punished and win at most z for all remaining stages; by following the agreement, he first get 1 and then get averagely $\gamma_l + (1 - \gamma_l)(z + 1)$, where $l \in \{1, 2\}$, $l \neq k$. This means that k is right to follow the agreement if:

$$z \leq (1 - \delta)1 + \delta(\gamma_l \cdot 1 + (1 - \gamma_l)(z + 1)),$$

and this is equivalent to the second inequality in (33).

The expression of limit payoffs is a straightforward application of the formulas that can be found in the proof of Lemma 8.

For a fixed δ , payoffs are given by:

$$\begin{aligned} x_i^\delta(n; \gamma) &= 1 + \gamma_i \omega, \\ x_j^\delta(n; \gamma) &= 1 + \mathbb{P}(\tau_i = 0)(1 - \gamma_i)\omega + \mathbb{P}(\tau_i = 0)\gamma_j \omega, \\ x_i^\delta(g; \gamma) &= z + \sum_{t=0}^T \mathbb{P}(\tau_j = 2t + 1)v_i \delta^{2t+1}, \\ x_j^\delta(g; \gamma) &= z + \sum_{t=0}^T \mathbb{P}(\tau_j = 2t)v_j \delta^{2t}. \end{aligned}$$

and the result comes from letting δ go to 1. ■

Elements of comparison with Peski (2013)

Following Peski (2013)'s ideas, we could also exhibit limit payoffs of finitely revealing equilibria. To this aim, one needs to refine the construction of the war of attrition equilibria by adding an arbitrary number of non-revealing stages between the information transmission stages. This means, e.g., that player i (resp., j) possibly concedes at stages $2tk$ (resp., $(2t + 1)k$), for some $k \geq 1$, instead of stages $2t$ (resp., $2t + 1$), the additional stages being devoted to playing (d, d) with probability 1 if no one conceded before. With an appropriate choice of k as an increasing function of δ , one can construct

a class of equilibria such that the number of information transmission stages is bounded as δ goes to 1, and such that equilibrium payoffs still converge³⁶.

To be more precise, such a construction would basically change δ into δ^k in the inequalities that drive the existence of war of attrition equilibria (Lemma 8). To ensure this existence, it is then sufficient to choose k such that δ^k tends to a limit that is arbitrarily close to 1. This can be achieved by choosing values of k of order $\frac{-c}{\log \delta}$, with $c > 0$ and small enough. With such a choice, one can check that the corresponding value of T_i is bounded, and so is the number of information transmission stages. Note that, in this construction, even if the number of information transmission stages is bounded, the horizon of the war of attrition (i.e., the last stage where some information may be revealed) still goes to infinity as $\delta \rightarrow 1$.

References

- [1] Abreu, D. and F. Gul (2000), “Bargaining and reputation”, *Econometrica* 68, 85-117.
- [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.

³⁶Moreover, these payoffs can be chosen arbitrarily close to the payoffs of Proposition 7.

- [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 (2003), “Some asymptotic results in discounted repeated games of one-sided incomplete information”, *Mathematics of Operations Research* 28, 433-462.
- [10] 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.
- [11] Forges, F. (1988), “Communication equilibria in repeated games with incomplete information”, *Mathematics of Operations Research* 13, 191-231.
- [12] Forges, F. (1992), “Repeated games of incomplete information : non-zero-sum”, in: R. Aumann and S. Hart, *Handbook of Game Theory with Economic Applications*, Elsevier Science Publishers, North-Holland, chapter 6, 155-177.
- [13] Forges, F. (2013), “A folk theorem for Bayesian games with commitment”, *Games and Economic Behavior* 78, 64–71.
- [14] Fudenberg, D. and J. Tirole (1991), *Game theory*, M.I.T. Press.
- [15] Hart S. (1985), “Nonzero-sum two-person repeated games with incomplete information”, *Mathematics of Operations Research* 10, 117–153.
- [16] Hörner, J. and S. Lovo (2009), “Belief-free equilibria in games with incomplete information”, *Econometrica* 77, 453–487.
- [17] 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.
- [18] Israeli, E. (1999), “Sowing doubt optimally in two-person repeated games,” *Games and Economic Behavior* 28, 203-216.

- [19] Koren, G. (1992), “Two-person repeated games where players know their own payoffs”, mimeo, Courant Institute of Mathematical Sciences, New York University.
- [20] Kreps, D. and R. Wilson (1982), “Reputation and imperfect information”, *Journal of Economic Theory* 27, 253-279.
- [21] Kalai, A., E. Kalai, E. Lehrer and D. Samet (2010), “A commitment folk theorem”, *Games and Economic Behavior* 69, 127-137.
- [22] Mailath, G. and L. Samuelson (2006), *Repeated Games and Reputations*, *Oxford University Press*.
- [23] 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.
- [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. (2013), “Repeated games with incomplete information and discounting”, *Theoretical Economics*, forthcoming.
- [28] Peters, M. and B. Szentes (2012), “Definable and contractible contracts”, *Econometrica* 80, 363-411.
- [29] Schmidt, K. (1993), “Reputation and equilibrium characterization in repeated games with conflicting interests”, *Econometrica* 61, 325-351.
- [30] Shalev, J. (1994), “Nonzero-sum two-person repeated games with incomplete information and known-own payoffs”, *Games and Economic Behavior* 7, 246-259.

- [31] Simon R., (1995), S. Spież and H. Toruńczyk, (1995), “The existence of equilibria in certain games, separation for families of convex functions and a theorem of Borsuk-Ulam type”, *Israel Journal of Mathematics* 92, 1-21.
- [32] Sorin, S. (1983), “Some results on the existence of Nash equilibria for non-zero sum games with incomplete information”, *International Journal of Game Theory* 12, 193-205.
- [33] Sorin, S. (1999), “Merging, reputation, and repeated games with incomplete information”, *Games and Economic Behavior* 29, 274-308.
- [34] Wiseman, T. (2012), “A partial folk theorem for games with private learning”, *Theoretical Economics* 7, 217-239.