



HAL
open science

Finitely repeated games with semi-standard monitoring

Pauline Contou-Carrère, Tristan Tomala

► **To cite this version:**

Pauline Contou-Carrère, Tristan Tomala. Finitely repeated games with semi-standard monitoring. 2010. halshs-00524134

HAL Id: halshs-00524134

<https://shs.hal.science/halshs-00524134>

Submitted on 7 Oct 2010

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

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



Documents de Travail du Centre d'Économie de la Sorbonne

C
E
S
W
o
r
k
i
n
g
P
a
p
e
r
s



Finitely repeated games with semi-standard monitoring

Pauline CONTOU-CARRÈRE, Tristan TOMALA

2010.73



Finitely repeated games with semi-standard monitoring*

Pauline Contou-Carrère[†] & Tristan Tomala[‡]

September 6, 2010

Abstract. This paper studies finitely repeated games with semi-standard monitoring played in pure strategies. In these games, each player's action set is endowed with a partition, and the equivalence classes of the actions played are publicly observed. We characterize the limit set of equilibrium payoffs as the duration of the game increases.

Keywords: Finitely repeated games, semi-standard monitoring, folk theorem

JEL Classification: C72, C73

1 Introduction

The main result of the theory of repeated games is the Folk Theorem which states that when players perfectly observe their opponents' actions and are patient enough, every feasible and individually rational payoff vector can be sustained by an equilibrium. This was first proved by Aumann and Shapley [1] and Rubinstein [17] for infinitely repeated undiscounted games. The result was then extended to discounted games by Sorin [20] for Nash

*The authors thank Joseph Abdou and Olivier Gossner for helpful discussions and comments

[†]CES, Université Paris 1 Panthéon-Sorbonne, CNRS, UMR 8174, 106-112 Bd de l'Hôpital, Paris Cedex 13, France. p.contou@yahoo.fr

[‡]HEC Paris, Department of Economics and Decision Sciences, 78351 Jouy-en-Josas Cedex, France. tomala@hec.fr

equilibria, and by Fudenberg and Maskin [6] for subgame perfect equilibria. Benoît and Krishna [2], [3] proved Folk Theorems for finitely repeated games with pure strategies and Gossner [9] extended the results of [2] to mixed strategies.

The assumption of perfect observation is not realistic in many applied settings and a large literature is devoted to repeated games with imperfect monitoring where players observe imperfect signals depending on the actions played. Seminal results for discounted repeated games are in Fudenberg et al. [5], [7] and Fudenberg and Levine [4], who characterize subgame perfect equilibrium payoffs of repeated games with public monitoring through dynamic programming methods. While the literature on discounted repeated games with private monitoring is profuse, no general characterization is known to date, see the book of Mailath and Samuelson [14] for a survey. On another hand, undiscounted repeated games with imperfect monitoring were first studied by Lehrer [10], [11], [12], who obtained characterizations of equilibrium payoffs for two-player games and for n -player games with semi-standard monitoring. Tomala [21] obtained a characterization for repeated games with public monitoring and pure strategies, and Renault and Tomala [16] characterized the set of communication equilibrium payoffs for any repeated game with imperfect monitoring.

Surprisingly little attention has been paid to finitely repeated games with imperfect monitoring, notable exceptions being the works of Sekiguchi [18], Mailath et al. [13], and Renault et al. [15]. However, these results are quite specific and to the best of our knowledge, no Folk-Theorem-like result is known for a reasonably wide class of games with imperfect monitoring.

We combine the approaches of Benoît and Krishna [2], [3], and of Lehrer [11] and consider finitely repeated games with semi-standard monitoring. These are finitely repeated games where each player's set of actions is equipped with an equivalence relation and at each stage of the game the equivalence classes of the actions played are publicly observed. Our aim is to study the asymptotics of the equilibrium payoff sets in relation with the results of Lehrer [11]. Considering the infinitely repeated undiscounted game, Lehrer characterized the set of equilibrium payoffs as the set of individually rational payoffs which are feasible by actions profiles that offer no unilateral deviation which is both profitable and undetectable. Let us call E this set of payoffs which is determined by the stage game and the equivalence relations on the action sets.

Firstly, following Benoît and Krishna [3], we assume that for each player, there exists a strictly individually rational Nash equilibrium payoff in the one-shot game. Under this assumption, we show that the set of Nash equilibrium payoffs \mathcal{E}_T of the T -fold repeated game converges to E as T goes to infinity, thus generalizing the result of Benoît and Krishna [3] to the semi-standard setup.

Secondly, following Benoît and Krishna [2], we study subgame perfect equilibria and assume that for each player, there exist two distinct Nash equilibrium payoffs in the one-shot game and that E has a non-empty interior (conditions C1 and C2 thereafter). It turns out that the set of subgame perfect equilibrium payoffs \mathcal{E}_T^* of the T -fold repeated game *does not* converge to E . The reason is the following. To obtain a subgame perfect equilibrium, one needs to control for optimality during punishment phases, i.e. we must

provide incentives to implement punishments. Due to imperfect monitoring, only action profiles that offer no undetectable and profitable deviation can be played by a subgame perfect equilibrium. The room for punishment is thus more narrow and the individual rationality constraints are tighter. We introduce the relevant notion of minmax level that takes into account the undetectable and profitable deviations, and we let E^* be the set of payoff in E which are individually rational with respect to those minmax levels. Under conditions C1 and C2, we prove that \mathcal{E}_T^* converges to E^* as T goes to infinity, thereby providing the generalization of the result of Benoît and Krishna [2] to the semi-standard setup. It is worthwhile to note that our results apply also to discounted games (finitely or infinitely repeated): the set of discounted subgame perfect equilibrium payoffs is a subset of E^* and we show convergence of this set to E^* as the discount factor goes to one. Like Benoît and Krishna [2], our results hold for pure strategies. The extension to mixed strategies, as in Gossner [9], is left as an open problem (the issue is commented in Section 5).

To conclude the introduction, let us compare our results to relevant works on finitely repeated games. Firstly, there is a line of literature (Smith [19], Gonzales-Diaz [8]) on finitely repeated games with perfect monitoring whose aim is to dispense with conditions on the stage game (like conditions C1 and C2). Similarly, Wen [22] studies discounted games without full dimensionality assumptions, and introduces the notion of effective minmax level. The focus of our paper is different. We do not depart from the usual and simple sufficient conditions of Benoît and Krishna [2], [3] on the stage game. Rather, we study the impact of observation on the equilibrium

payoffs of the repeated game. In particular, the minmax we introduce is not comparable with Wen's effective minmax. Our notion is driven by purely informational considerations.

Secondly, as mentioned above, the literature on finitely repeated games with imperfect monitoring is scarce. Most relevant to the present paper is the work of Sekiguchi [18] who studies the finite repetition of a stage game which has a single equilibrium. The signals have a product structure akin to semi-standard monitoring: for each player i , a stochastic signal depending on player i 's action only, is publicly announced. Distributions of signals are assumed to have common support. Under these assumptions, Sekiguchi [18] shows that the finitely repeated game has a unique equilibrium payoff. We differ from Sekiguchi's model in two aspects. First, following Benoît and Krishna [2], we assume that the stage game has multiple equilibria. Second, we assume pure strategies, and consequently the signals do not have common support (otherwise, they are trivial). Another related paper is Mailath et al. [13] which shows on a specific finitely repeated game with public monitoring, that equilibria in public strategies do not exhaust all equilibria of the game. This result is specific to mixed strategies, as any pure strategy is fully equivalent to a pure public strategy. Finally, Renault et al. [15] prove a Folk Theorem for finitely repeated minority games with public signals. Note that the monitoring structure of this game is not semi-standard and that specific properties of the minority games are used extensively in Renault et al. [15].

The paper is organized as follows. Section 2 presents the model. In Section 3, we give necessary conditions satisfied by equilibrium payoffs. The

main theorem is in Section 4. Section 5 offers extensions and concluding remarks. Section 6 is devoted to the proof of the main theorem.

2 The model

Let $G = (N, (X_i)_{i \in N}, (g_i)_{i \in N})$ be a normal form game where $N = \{1, \dots, n\}$ is a finite set of players, for each $i \in N$, X_i denotes the set of actions of player i and $g_i : \times_{i \in N} X_i \rightarrow \mathbb{R}$ is the payoff function of player i . The set of action profiles is denoted $X = \times_{i \in N} X_i$ and $g : X \rightarrow \mathbb{R}^n$ denotes the vector payoff function $(g_i)_{i \in N}$. All action sets are assumed to be topological compact spaces and g is assumed to be continuous. We denote $M = \max_{i \in N, x \in X} |g_i(x)|$. Throughout the paper, we use the conventional notation $-i$ to denote the set of players $j \neq i$ and use it for indexing $x_{-i} = (x_j)_{j \neq i}$ and for products of sets $X_{-i} = \times_{j \neq i} X_j$.

Each action set X_i is endowed with a partition \bar{X}_i . For x_i in X_i , we denote \bar{x}_i the cell of the partition containing x_i . The partition induces an equivalence relation \sim_i on X_i : $x_i \sim_i y_i$ if $\bar{x}_i = \bar{y}_i$. We denote $\bar{X} = \times_{i \in N} \bar{X}_i$.

As particular case, games with finite action sets enter this setup. Another class of games of interest are *finite games in mixed strategies with observable distributions of signals*. For each player i , there is an underlying finite set of actions A_i endowed with a partition. We let X_i be the set of probability distributions over A_i . A distribution x_i over actions defines a distribution \bar{x}_i over equivalence classes of actions: $\bar{x}_i(\bar{a}_i) = \sum_{a_i \in \bar{a}_i} x_i(a_i)$. The partition of A_i induces an equivalence relation over X_i defined by: $x_i \sim_i y_i \Leftrightarrow \bar{x}_i = \bar{y}_i$. In this model, players choose mixed actions and if player i chooses x_i , the

distribution \bar{x}_i is publicly observed.

We turn now to the description of the repeated game. For each integer $T \geq 1$, the T -fold repeated game G_T unfolds as follows. At each stage $t \leq T$, players select actions simultaneously and if $x = (x_i)_i \in X$ is the action profile chosen, the profile of equivalence classes $\bar{x} = (\bar{x}_i)_i$ is publicly announced. The game ends after stage T and if $(x_t)_{t=1}^T$ is the sequence of actions profiles, player i gets the payoff $\frac{1}{T} \sum_{t=1}^T g_i(x_t)$. This describes a *repeated game with semi-standard monitoring*.

Now, we describe the strategies in the repeated game. For all player i and $t \leq T$, let $H_{i,t} = (X_i \times \bar{X})^{t-1}$ be the set of histories at stage t of player i ($H_{i,1}$ is a singleton) and $\mathcal{H}_{i,T} = \bigcup_{1 \leq t \leq T} H_{i,t}$ be the set of all histories of player i in G_T . The set of *public histories* at stage t is $H_{p,t} = \bar{X}^{t-1}$ and $\mathcal{H}_{p,T} = \bigcup_{1 \leq t \leq T} H_{p,t}$ is the set of all public histories in G_T . A *pure strategy* for player i is a mapping $\sigma_i : \mathcal{H}_{i,T} \rightarrow X_i$. Let $\Sigma_{i,T}$ be the set of pure strategies of player i in G_T and $\Sigma_T = \times_{i \in N} \Sigma_{i,T}$ be the set of profiles of pure strategies. A profile $\sigma = (\sigma_i)_{i \in N}$ induces a unique sequence of action profiles denoted $(x_t(\sigma))_{1 \leq t \leq T}$ where $x_t(\sigma) = (x_{i,t}(\sigma))_{i \in N}$ and $x_{i,t}(\sigma)$ is the action played by player i at stage t .

Given a strategy profile σ , the average payoff of player i is $\gamma_{i,T}(\sigma) = \frac{1}{T} \sum_{t=1}^T g_i(x_t(\sigma))$. The repeated game G_T is now given by the normal form $(N, (\Sigma_{i,T})_{i \in N}, (\gamma_{i,T})_{i \in N})$.

Definition 2.1 *A strategy profile $\sigma \in \Sigma_T$ is a Nash equilibrium of G_T if for every player i and for all $\tau_i \in \Sigma_{i,T}$, $\gamma_{i,T}(\tau_i, \sigma_{-i}) \leq \gamma_{i,T}(\sigma)$.*

We let $\mathcal{E}_T \subseteq \mathbb{R}^n$ be the set of Nash equilibrium payoffs of G_T . That is,

$u \in \mathcal{E}_T$ if there exists a Nash equilibrium σ of G_T such that $u_i = \gamma_{i,T}(\sigma)$ for each player i .

The main focus of this paper is on subgame perfect equilibria. To define this concept properly, we first introduce public strategies: a strategy of player i is *public* if it only depends on public signals, that is on profiles of equivalence classes of actions. In other words, player i 's strategy depends on her own past actions only through the equivalence classes. This concept is widely used in repeated games (see e.g. Fudenberg and Levine [4] and Fudenberg et al. [5]) and may impose restrictions on the resulting equilibria (see Mailath et al. [13]). However, restricting to public strategies is without loss of generality in games with pure strategies. The argument is simple and well known (see e.g. Tomala [21]). Consider a pure public strategy of player i . The action played at the first stage is encoded in the strategy. Given a public history, player i can use the strategy to compute the actions she played in the past. The recall of her own past actions is thus implicit.

More precisely, let $\sigma_i \in \Sigma_{i,T}$ be a pure strategy of player i and $h_p \in \mathcal{H}_{p,T}$ be a public history. Then, σ_i induces a unique sequence of actions compatible with h_p , and therefore there is a unique private history $h_i := \tilde{h}_i(\sigma_i, h_p)$ compatible with h_p and σ_i . Thus, the next action $\sigma_i(\tilde{h}_i(\sigma_i, h_p))$ depends on the public history only. This construction defines a public strategy σ'_i which is *equivalent* to σ_i , i.e. for all σ_{-i} and all stages t , $x_{i,t}(\sigma_i, \sigma_{-i}) = x_{i,t}(\sigma'_i, \sigma_{-i})$. In the remainder of the paper, we identify a pure strategy with the equivalent public strategy. This allows to give a sound definition of subgame perfect equilibria. In repeated games with imperfect monitoring, there is no proper subgame since actions are not publicly observed. However,

when all players rely on public information only, we may safely define the subgame that follows a public history. This leads to the concept of perfect public equilibrium (Fudenberg and Levine [4] and Fudenberg et al. [5]) which we simply call subgame perfect equilibrium in the present context. Given a strategy σ_i of player i and a public history h , we denote by $\sigma_i[h]$ the continuation strategy following h . That is, $\sigma_i[h](h') = \sigma_i(hh')$ for each public history h' (hh' denotes the concatenated history h followed by h'). We denote $\sigma[h]$ the profile $(\sigma_i[h])_i$.

Definition 2.2 *A strategy profile $\sigma \in \Sigma_T$ is a subgame perfect equilibrium of G_T if for all $t < T$ and $h \in H_{p,t}$, $\sigma[h]$ is a Nash equilibrium of G_{T-t} .*

We let $\mathcal{E}_T^* \subseteq \mathbb{R}^n$ be the set of subgame perfect equilibrium payoffs of G_T . That is, $u \in \mathcal{E}_T^*$ if there exists a subgame perfect equilibrium σ of G_T such that $u_i = \gamma_{i,T}(\sigma)$ for each player i . In the sequel, we characterize the limit (in the Hausdorff topology) of \mathcal{E}_T^* as T goes to infinity. Note that, in our pure strategy setup, \mathcal{E}_T^* is non empty for some T if and only if the game G admits a Nash equilibrium in pure strategies. We maintain this assumption throughout the paper.

3 Necessary conditions

The usual necessary conditions for equilibrium payoffs of repeated games are feasibility and individual rationality. We give the counterpart of each condition in our model. First, we revisit feasibility.

3.1 Feasibility

Following Lehrer [11], we argue that the actions profiles played at equilibrium must not offer undetectable and profitable deviations to the players. The intuition is as follows. Assume that player i has two equivalent actions $x_i \sim_i y_i$, and that player i receives a higher payoff by playing x_i than by playing y_i , given the actions y_{-i} of the other players. It is then impossible to have player i choosing y_i against y_{-i} at an equilibrium. If so, player i would deviate to x_i to increase her stage payoff. Since $\bar{x}_i = \bar{y}_i$, the deviation goes unnoticed and therefore has no impact on future payoffs. The deviation is thus profitable in the repeated game.

For each player $i \in N$, let D_i be the set actions profiles $x = (x_i)_{i \in N}$ such that if player i deviates to an action in \bar{x}_i , her payoff does not increase.

Definition 3.1 For each player $i \in N$,

$$D_i = \left\{ (x_i, x_{-i}) \in X_i \times X_{-i} : g_i(x_i, x_{-i}) = \max_{y_i \sim_i x_i} g_i(y_i, x_{-i}) \right\}$$

We let $D := \bigcap_{i \in N} D_i$. An action profile in D is such that there exists no unilateral deviation which is both profitable and undetectable. Remark that a Nash equilibrium of G belongs to D . Indeed, if $x \in X$ is a Nash equilibrium of G , then for each i , x_i is a best response of player i to x_{-i} . Then x_i maximizes player i 's payoff also on the set of actions which are equivalent to x_i . As a consequence, D is a non-empty compact subset of X (recall that g is continuous).

Examples 3.2

1. Monitoring is *perfect* when $x_i \sim_i y_i \Rightarrow x_i = y_i$. It follows directly that $D_i = D = X$.
2. Monitoring is *trivial* when $x_i \sim_i y_i$ for all $x_i, y_i \in X_i$. In this case, D_i is the set of action profiles (x_i, x_{-i}) such that x_i is a best response against x_{-i} . It follows that D is the set of Nash equilibria of G .
3. D needs not be a product set of actions. Consider the following numerical two-player example. The payoff matrix is

	y_1	y_2	y_3	y_4
x_1	3, 4	1, 3	3, 2	1, 1
x_2	-1, 1	0, 0	4, 2	3, 0
x_3	2, 1	-1, 2	5, 0	1, -1
x_4	-1, 1	4, 0	-1, 1	0, 0

The partitions are $\bar{X}_1 = \{\{x_1, x_2\}, \{x_3, x_4\}\}$ and $\bar{X}_2 = \{\{y_1, y_2\}, \{y_3, y_4\}\}$.

One can easily check that $D_1 = \{(x_1, y_1), (x_1, y_2), (x_2, y_3), (x_2, y_4), (x_3, y_1), (x_3, y_3), (x_3, y_4), (x_4, y_2)\}$ and $D_2 = \{(x_1, y_1), (x_1, y_3), (x_2, y_1), (x_2, y_3), (x_3, y_2), (x_3, y_3), (x_4, y_1), (x_4, y_3)\}$. It follows that $D = \{(x_1, y_1), (x_2, y_3), (x_3, y_3)\}$.

The next lemma shows that a Nash equilibrium of the repeated game prescribes actions profiles in D on the equilibrium path. Further, a subgame perfect equilibrium prescribes actions profiles in D after all histories.

Lemma 3.3

- If σ is a Nash equilibrium of G_T , then for all $t \leq T$, $x_t(\sigma) \in D$.

- If σ is a subgame perfect equilibrium of G_T , then for each public history h , $\sigma(h) \in D$.

Proof. Let us prove the second point. Let σ be a subgame perfect equilibrium of G_T and assume by contradiction that there is a public history h of length t such that $\sigma(h) \notin D$. Without loss of generality, assume $\sigma(h) \notin D_1$. Let τ_1 be a strategy of player 1 such that for all $h' \in \mathcal{H}_{p,T} \setminus \{h\}$, $\tau_1(h') = \sigma_1(h')$ and $\tau_1(h)$ is an action such that $\bar{\tau}_1(h) = \bar{\sigma}_1(h)$ and $g_1(\tau_1(h), \sigma_{-1}(h)) > g_1(\sigma(h))$. The public signal after h is the same under (τ_1, σ_{-1}) and under (σ_1, σ_{-1}) . Thus, the sequence of action profiles induced by (τ_1, σ_{-1}) after stage $t + 1$ coincides with the one induced by (σ_1, σ_{-1}) . It follows that $\gamma_{1,T-t}(\tau_1[h], \sigma_{-1}[h]) > \gamma_{1,T-t}(\sigma[h])$, contradicting the fact that $\sigma[h]$ is a Nash equilibrium of G_{T-t} .

The first point follows by the same argument applied on the equilibrium path. □

As a consequence, an equilibrium payoff is a convex combination of points in $g(D)$.

Corollary 3.4 For each T , $\mathcal{E}_T^* \subseteq \mathcal{E}_T \subseteq \text{co } g(D)$

3.2 Individual rationality

In this section, we revisit individual rationality and we start by recalling the usual notion. We denote the *minmax level* of player i by

$$v_i = \min_{x_{-i} \in X_{-i}} \max_{x_i \in X_i} g_i(x_i, x_{-i})$$

and we let IR be the set of payoff vectors $u \in \mathbb{R}^n$ such that for all $i \in N$, $u_i \geq v_i$. As usual in repeated games, any equilibrium payoff is individually rational: $\mathcal{E}_T \subseteq IR$ for each T . The proof is standard, for any strategy σ_{-i} of players $-i$, playing a best-reply to $\sigma_{-i}(h)$ after each history h , yields a payoff no less than v_i to player i .

However, with semi-standard monitoring, we get tighter constraints for subgame perfect equilibria. The intuition is the following. From Lemma 3.3, a subgame perfect equilibrium prescribes action profiles in D after each history. It follows that if the minmax above cannot be achieved by a joint action in D , then it is not possible to punish player i to v_i within a subgame perfect equilibrium. We introduce the relevant notion of minmax level. For each $i \in N$, let D_{-i}^* be the projection of D on X_{-i} , i.e.,

$$D_{-i}^* = \{x_{-i} \in X_{-i} : \exists x_i \in X_i, (x_i, x_{-i}) \in D\}.$$

Definition 3.5 *For each player i in N , the semi-standard minmax level of player i is,*

$$v_i^* = \min_{x_{-i} \in D_{-i}^*} \max_{x_i \in X_i} g_i(x_i, x_{-i})$$

Let $x_{-i}^*(i) \in D_{-i}^*$ be an action profile for players $-i$ which achieves the minimum above. Against this profile, player i 's payoff is at most v_i^* , no matter which action she plays. Let $x_i^*(i) \in X_i$ be such that $x^*(i) := (x_{-i}^*(i), x_i^*(i)) \in D$. If $x^*(i)$ is played, then there is no player who can profitably deviate without being detected, and player i 's payoff is at most v_i^* . We denote IR^* the set of payoffs which are individually rational with respect to the semi-standard minmax levels.

Lemma 3.6 For each T , $\mathcal{E}_T^* \subseteq IR^*$.

Proof. Let σ be a subgame perfect equilibrium of G_T and τ_i be the strategy of player i which plays a stage-best reply after each history. That is, $\forall h \in \mathcal{H}_{p,T}$, $\tau_i(h) \in \operatorname{argmax}_{x_i \in X_i} g_i(x_i, \sigma_{-i}(h))$. By Lemma 3.3, $\sigma_{-i}(h) \in D_{-i}^*$, for each $h \in \mathcal{H}_{p,T}$. It follows:

$$g_i(\tau_i(h), \sigma_{-i}(h)) \geq \min_{x_{-i} \in D_{-i}^*} \max_{x_i \in X_i} g_i(x_i, x_{-i}) = v_i^*$$

and thus $\gamma_{i,T}(\tau_i, \sigma_{-i}) \geq v_i^*$. Since σ is an equilibrium of G_T , $\gamma_{i,T}(\sigma_i, \sigma_{-i}) \geq \gamma_{i,T}(\tau_i, \sigma_{-i}) \geq v_i^*$. \square

We provide now an example where the semi-standard minmax level differs from the usual one.

Example 3.7 Consider a two-player game where the sets of actions are $X_1 = \{a_1, b_1, c_1\}$ and $X_2 = \{a_2, b_2, c_2\}$ for players 1 and 2 respectively. The partitions are $\bar{X}_1 = \{\{a_1, b_1\}, \{c_1\}\}$ and $\bar{X}_2 = \{\{a_1, b_2\}, \{c_2\}\}$. The payoff matrix is:

	a_2	b_2	c_2
a_1	3, 4	6, 0	2, -1
b_1	-1, 2	5, 1	0, 1
c_1	2, 4	3, 0	3, 3

The action b_1 of player 1 is strictly dominated by a_1 and $\bar{a}_1 = \bar{b}_1$. It follows that $D_1 = \{(x_1, x_2) : x_1 \neq b_1\}$. Similarly, $D_2 = \{(x_1, x_2) : x_2 \neq b_2\}$.

Therefore, $D = \{a_1, c_1\} \times \{a_2, c_2\}$. It follows that

$$v_2^* = \min_{x_1 \neq b_1} \max_{x_2 \in X_2} g_2(x_1, x_2) = 4$$

while $v_2 = 2$.

For this example, it is worthwhile to note that the minmax levels (v_2 and v_2^*) are the same if the game is played in mixed strategies with observable distributions of signals. Indeed, in this model, D is the set of mixed action profiles such that the dominated actions b_1 of player 1 and b_2 of player 2 are played with probability 0. •

4 The main results

We examine now the convergence of the set of equilibrium payoffs to the set of individually rational payoffs that are feasible and robust to undetectable unilateral deviations. Denote $E = \text{co } g(D) \cap IR$ and $E^* = \text{co } g(D) \cap IR^*$. Lehrer [11] proved that E is the set of equilibrium payoffs of the undiscounted infinitely repeated game. On another hand, for perfect monitoring, Benoît and Krishna [3] proved that, under some condition on the stage game, \mathcal{E}_T converges to E as T goes to infinity. This result easily extends to semi-standard monitoring.

Theorem 4.1 *Assume that for each player $i \in N$, there exists a Nash equilibrium e^i of G_1 such that $g_i(e^i) > v_i$. Then \mathcal{E}_T converges to E in the Hausdorff topology as T goes to infinity.*

Proof. The proof is a straightforward extension of the construction of Benoît and Krishna [3] and we just provide an outline. Given a payoff vector $u \in E$, there exists a finite sequence of actions that induce an average payoff vector ε -close to u . The strategy recommends the players to follow the sequence cyclically for finitely many cycles. At the end of this phase, if no deviation is detected, then each static equilibrium e^i is played during R_i consecutive stages at the end of play. Therefore, in case of no detectable deviation, the play ends up in a sequence of one-shot Nash equilibria during $R = \sum_{i=1}^n R_i$ consecutive stages. If at some stage t , some player i deviates from the main path in a detectable way, then her opponents play a minmax strategy against player i from stage t to the end of the game. Following Benoît and Krishna [3], we can adjust the lengths R_i and the number of cycles in such a way that this strategy profile forms a Nash equilibrium of G_T provided T is large enough. \square

Now, we extend the result of Benoît and Krishna [2] for subgame perfect equilibria, to repeated games with semi-standard monitoring. Benoît and Krishna [2] introduced the following conditions:

Condition C 1 *For each player $i \in N$, there exist two Nash equilibria of the static game e^i and f^i such that $g_i(e^i) > g_i(f^i)$.*

Condition C 2 $\dim(\text{co } g(D)) = n$

Benoît and Krishna [2] proved that in games with perfect monitoring, \mathcal{E}_T^* converges to E as T goes to infinity, under these conditions. It follows that, under conditions C1 and C2, \mathcal{E}_T and \mathcal{E}_T^* have the same limit for games with

perfect monitoring. Our next result shows that this is not the case for games with semi-standard monitoring.

Theorem 4.2 *Under conditions C1 and C2, \mathcal{E}_T^* converges in the Hausdorff topology to E^* as T goes to infinity. I.e., for all $\varepsilon > 0$, there exists $T_0 \in \mathbb{N}$ such that for all $T \geq T_0$ and $u \in E^*$, there exists $v \in \mathcal{E}_T^*$, $\|u - v\| \leq \varepsilon$.*

Examples 4.3 1. In the case of perfect monitoring, we have for each player i , $D_i = X$, thus $D = X$ and $D_{-i}^* = X_{-i}$. It follows that $E^* = \text{co } g(X) \cap IR$. Theorem 4.2 is then the Folk Theorem of Benoît and Krishna [2] when specialized to perfect monitoring.

2. In the case of trivial information, D is the set of Nash equilibrium payoffs of G and E^* is the convex hull of the associated set of payoff vectors. Clearly \mathcal{E}_T^* converges to E^* , it is enough to approximate a convex combination of Nash payoffs by playing static equilibria with the right proportions.

The proof of Theorem 4.2 follows Gossner [9]. The idea is the following. We define a normal path as a finite sequence of actions profiles that yield an average payoff vector close to the target payoff. When a unilateral deviation is detected, the play enters a punishment phase. Punishers are rewarded at late stages of the game. Not applying punishments entails losing the terminal reward. This deters detectable deviations, both from the normal path and from the punishment phase. The strategy only uses actions profiles in D and thus undetectable deviations are not profitable. The detailed construction is in the Appendix.

5 Extensions and open issues

5.1 Discounted games

Our main results extend easily to discounted games. Let $\delta \in [0, 1)$ be a discount factor. The T -fold discounted game $G_{T,\delta}$ is described as the game G_T except for the payoff function which is given by:

$$\gamma_{T,\delta}(\sigma) = \sum_{t=1}^T \frac{1-\delta}{1-\delta^T} \delta^{t-1} g(x_t(\sigma)).$$

The discounted game $G_{\infty,\delta}$ is described similarly but is repeated infinitely many times. The payoff function is:

$$\gamma_{\infty,\delta}(\sigma) = \sum_{t=1}^{\infty} (1-\delta) \delta^{t-1} g(x_t(\sigma)).$$

We let $\mathcal{E}_{T,\delta}^*$ (resp. $\mathcal{E}_{\infty,\delta}^*$) be the set of subgame perfect equilibrium payoffs of the game $G_{T,\delta}$ (resp. $G_{\infty,\delta}$). The counterpart of Theorem 4.2 for discounted game is the following.

Theorem 5.1 • For each T, δ , $\mathcal{E}_{T,\delta}^* \subseteq \mathcal{E}_{\infty,\delta}^* \subseteq E^*$.

- Under conditions C1 and C2, $\mathcal{E}_{T,\delta}^*$ (resp. $\mathcal{E}_{\infty,\delta}^*$) converges in the Hausdorff topology to E^* as $T \rightarrow \infty$ and $\delta \rightarrow 1$ (resp. $\delta \rightarrow 1$).

Proof. To prove the first point, first note that Lemma 3.3 extends to discounted games. That is, if σ is a subgame perfect equilibrium of $G_{T,\delta}$ (resp. $G_{\infty,\delta}$), then for each history h , $\sigma(h) \in D$. The proof is exactly the same. The inclusion in E^* follows as in Lemma 3.6. To see the inclusion

$\mathcal{E}_{T,\delta}^* \subseteq \mathcal{E}_{\infty,\delta}^*$, consider a subgame perfect equilibrium σ of $G_{T,\delta}$ and repeat it periodically every T stages (past history is forgotten at the beginning of each T -period). The resulting strategy is clearly a subgame perfect equilibrium of $G_{\infty,\delta}$ and the overall discounted payoff is $\gamma_{T,\delta}(\sigma)$.

To prove the second point, it is thus enough to prove that $\mathcal{E}_{T,\delta}^*$ converges to E^* as $T \rightarrow \infty$ and $\delta \rightarrow 1$. One can see from the proof of Theorem 4.2 that any detectable deviation yields a payoff strictly less than the equilibrium payoff for a suitable choice of the parameters. These equilibrium constraints are thus satisfied in $G_{T,\delta}$ as well, provided that δ is close enough to 1. \square

Regarding the convergence of $\mathcal{E}_{\infty,\delta}^*$ to E^* as $\delta \rightarrow 1$, a stronger result can be obtained. Namely, Condition C1 can be dispensed with. To show this, one may adapt the construction of Fudenberg and Maskin [6] to our setup. We use only action profiles in D (on and off the equilibrium path), so that deviations can be assumed to be detectable. The construction then adapts easily.

Note that Fudenberg et al. [7] provide a characterization of $\lim_{\delta \rightarrow 1} \mathcal{E}_{\infty,\delta}^*$ for games with public monitoring, using linear programming methods. Our result provides an alternative characterization which allows a constructive approach to the equilibrium strategies.

5.2 Other signalling structures

Let $f : X \rightarrow S$ be a function mapping the set of actions profiles to a set of signals. The associated *repeated game with public monitoring* is such that $f(x)$ is publicly observed when x is played. The semi-standard monitoring

case is such that $f(x) = (\bar{x}_i)_i$. A natural generalization is the following.

Definition 5.2 *The mapping $f : X \rightarrow S$ is rectangular if for every $s \in S$, the inverse image $f^{-1}(s) = \{x \in X : f(x) = s\}$ is a product set $\times_i X_i(s)$ where $X_i(s)$ is a subset of X_i .*

Our results extend naturally to this case. Let D_i be the set of action profiles x such that, $g_i(x_i, x_{-i}) = \max_{y_i \in X_i(f(x))} g_i(y_i, x_{-i})$, and $D = \cap_i D_i$. The definition of the modified minmax v_i^* then follows. The main arguments for adapting the proofs are the following. Clearly, if σ is a subgame perfect equilibrium, $\sigma(h) \in D$ for each h . The inclusion of equilibrium payoffs in E^* follows. To adapt the constructions of the equilibrium strategies, it is enough to remark that when a deviation is detected, the deviating player is identified. To see this, assume by contradiction that there exists $x \in X$, $y_1 \in X_1$ and $y_2 \in X_2$ such that,

$$t := f(y_1, x_{-1}) = f(y_2, x_{-2}) \neq f(x)$$

This implies that $(y_1, x_2, x_{-12}) \in X_1(t) \times X_2(t) \times X_{-12}(t)$ and $(x_1, y_2, x_{-12}) \in X_1(t) \times X_2(t) \times X_{-12}(t)$. But then, $(x_1, x_2, x_{-12}) \in X_1(t) \times X_2(t) \times X_{-12}(t)$ which contradicts $f(x) \neq t$.

More generally, the results extend whenever the signalling function allows to ascribe any detectable unilateral deviation to a single player. Without this latter assumption on signals, characterizing the limit of \mathcal{E}_T^* is an open problem (see Fudenberg et al. [7] and Tomala [21] for characterizations of equilibrium payoffs in infinitely repeated games with public monitoring).

5.3 Mixed strategies

Following Benoît and Krishna [2], we have restricted the analysis to pure strategies. Consider a finitely repeated game with semi-standard monitoring, finite action spaces and mixed strategies. Assume first that distributions of signals are observable and denote E_m^* the corresponding limit equilibrium payoffs set (i.e. $\text{co } g(D) \cap IR^*$). Then \mathcal{E}_T^* converges to E_m^* (under conditions C1 and C2) from Theorem 4.2. Assume now that only realized signals are observable and let $\mathcal{E}_{m,T}^*$ be the corresponding set of (mixed strategies) equilibrium payoffs. It is easy to show that $\mathcal{E}_{m,T}^* \subseteq E_m^*$ and a reasonable guess is that convergence still holds (under conditions C1 and C2).

Gossner [9] extends the Folk Theorem of Benoît and Krishna [2] to finitely repeated games with perfect monitoring and mixed strategies where only realized actions are observed. The idea is the following. The normal path defines pure actions that approximate the target payoff. Punishment phases require the use of mixed actions. Statistical tests are performed to check the empirical frequency of actions for each player. Players are rewarded at late stages of the game if they pass the test. The key idea of Gossner's proof is to avoid the explicit construction of the strategies. The statistical test and the rewarding schemes are designed so that each player has a strong incentive to pass the test with high probability and such that passing the test implies that the average payoff during the punishment phase is close to the minmax level. Gossner then argues that a subgame perfect equilibrium of the game defined by the rewarding schemes approximates the target payoff.

This approach does not work in our setup. We can use Gossner's construction to force the players to induce a prescribed empirical frequencies of signals. However, it is not possible to control for the actions played, within a given equivalence class. In particular, two mixed action profiles may yield the same distribution over signals and different payoff vectors. As a consequence, a subgame perfect equilibrium of the game defined by the statistical tests and the rewarding schemes, may pass the test with high probability and yield a payoff far away from the target payoff. An open issue is thus to prove convergence of $\mathcal{E}_{m,T}^*$ to E_m^* when only realized signals are observable.

6 Appendix: Proof of the main theorem

The proof adapts the method of Gossner [9] to semi-standard observation. We first prove an analog of Theorem 4.2 for repeated games with terminal payoffs.

A *repeated game with terminal payoffs* is a T -fold repeated game such that, at the end of the play, each player i receives a history-dependent reward. Given a mapping $\text{term} = (\text{term}_i)_i : \mathcal{H}_{p,T} \rightarrow \mathbb{R}^n$, $G_T(\text{term})$ is the game with strategy sets $(\Sigma_{i,T})_i$ and payoff functions:

$$\theta_{i,T}(\sigma) = \frac{1}{T} \sum_{t=1}^T g_i(x_t(\sigma)) + \frac{1}{T} \text{term}_i((x_t(\sigma))_{1 \leq t \leq T})$$

That is, players choose actions at stages $t \leq T$ and receive the average payoff plus the terminal payoff. We let $\mathcal{E}_T^*(\text{term})$ be the set of subgame perfect equilibrium payoffs of $G_T(\text{term})$.

In the sequel, we focus on terminal payoffs that take the three values $-W, 0, W$ with $W > 0$. We let $\text{Term}(W)$ be the set of mappings from $\mathcal{H}_{p,T}$ to $\{-W, 0, W\}^n$ and $\mathcal{E}_T^*(W) = \cup_{\text{term} \in \text{Term}(W)} \mathcal{E}_T^*(\text{term})$ be the union of the corresponding sets of subgame perfect equilibrium payoffs.

The next proposition is the analog of Theorem 4.2 for repeated games with terminal payoffs and is the main step of the proof. We assume that E^* has a non-empty interior. Note that this holds under conditions C1 and C2. Indeed, any Nash equilibrium of G is in E^* and it follows from C1 and C2 that the convex combination $\frac{1}{2N} \sum_{i \in N} (g(e^i) + g(f^i))$ is in the interior of E^* .

Proposition 6.1 *For all $\varepsilon > 0$ and $u \in \text{int}(E^*)$, there exist $T_0 \in \mathbb{N}$ and $W_0 \in \mathbb{R}_+$ such that for all $T \geq T_0$ and $W \geq W_0$, there exists $v \in \mathcal{E}_T^*(W)$ with $\|u - v\| \leq \varepsilon$ ($\text{int}(\cdot)$ denotes the topological interior).*

Proof. Let us choose $\varepsilon > 0$ and $u \in \text{int}(E^*)$. Without loss of generality, we assume that u is a convex combination of points in $g(D)$ with rational coefficients (otherwise, u can be arbitrarily approximated by such convex combinations). We write $u = \sum_{m=1}^M \frac{\alpha_m}{\alpha} \cdot g(z_m)$ with $z_m \in D$, α_m positive integers and $\alpha = \sum_m \alpha_m$. Note that u is the average payoff vector along a sequence of α action profiles where z_m is played α_m times.

Let us fix an integer $K > 1$, a positive real number W , and $P = k_P \alpha$ where k_P is a positive integer such that $1 < k_P < K$. We let $T = K\alpha$ be the length of the game. We call *normal path* the sequence of actions profiles $(y_t)_{t=1, \dots, T}$ such that $y_t = z_m$ if $t = m \bmod(\alpha)$. That is, along the normal path, z_1 is played α_1 times, \dots , z_M is played α_M times, and this is repeated

K times cyclically. We define a strategy profile σ^* as follows.

- The play starts in the normal path: y_1 is played at stage 1.
- If the play is on the normal path at stage t , then y_t is played. If the public signal is \bar{y}_t , the play remains on the normal path at stage $t + 1$. Otherwise a deviation is detected, i.e. there exists a player i who plays at stage t an action $x_{i,t}$ such that $\bar{x}_{i,t} \neq \bar{y}_{i,t}$.
- If a deviation of player i from the normal path is detected at stage $t < T - P$, the play switches to the *punishment phase of player i* (if there are several such players, we choose one arbitrarily).
- During the punishment phase of player i , the action profile $(x_{-i}^*(i), x_i^*(i))$ is played for P stages. Then, the play switches back to the normal path.
- If a deviation from the normal path is detected at stage $t \geq T - P$, the play remains on the normal path at stage $t + 1$.
- After any other possible history, a fixed Nash equilibrium of G is played.

The terminal payoff function is defined as follows.

- If the history follows the main path until stage T , then each player gets 0 as terminal payoff. Otherwise, a deviation is detected and we consider the last deviation detected along the history.
- If the last deviation is detected at some *early* stage $t < T - P$, consider the punishment phase that follows and let i be the player punished in

this phase. An *effective punisher* is a player $j \neq i$ who played an action x_j such that $\bar{x}_j = \bar{x}_j^*(i)$ during the punishment phase of player i . If player $j \neq i$ is an effective punisher then her terminal payoff is W . Otherwise, her terminal payoff is 0.

- If a deviation is detected at some *late* stage $t \geq T - P$, then all players get $-W$ as terminal payoff.

Now, we show that σ^* is a subgame perfect equilibrium of $G_T(\text{term})$ for suitable choice of the parameters K, W, P . The one-shot deviation principle applies: it is enough to check that there is no profitable unilateral deviation that deviates only once from the equilibrium strategies. Firstly, if all players abide by this strategy profile, then the average payoff is u . Secondly, by construction, $\sigma^*(h) \in D$ for each history h . It follows that a deviation that does not change the public signals is not profitable. Thus, we only need to check that there is no profitable and detectable deviation.

Early deviation from the normal path. Assume that player i deviates from σ_i^* at stage $t < T - P$ by playing an action $x_{i,t}$ such that $\bar{x}_{i,t} \neq \bar{y}_{i,t}$. Player i may increase her stage payoff at stage t . However, she is punished for P stages instead of getting the target payoff u_i . Overall the deviation is not profitable if the punishment phase is long enough. Observe that the terminal payoff of player i does not depend on her deviation. Let U_i be the total payoff of player i from stage $t + P + 1$ to stage T , and let term_i be player i 's terminal payoff. The deviation is not profitable if:

$$M + Pv_i^* + U_i + \text{term}_i \leq -M + Pu_i + U_i + \text{term}_i$$

This condition is satisfied if $P \geq P_0 := \max_{i \in N} \frac{2M}{u_i - v_i^*}$.

Late deviation from the normal path. Assume that player i deviates from σ_i^* at stage $t \geq T - P$. By construction of the strategies, player i may increase her stage payoff until the end of the game. However, she induces the bad terminal payoff $-W$. The deviation is not profitable if:

$$PM - W \leq -PM$$

This condition is satisfied if $W \geq W_0 := 2PM$.

Deviation from a punishment phase. Consider a unilateral detectable deviation by player j from player i 's punishment phase ($j \neq i$). Player j may increase her stage payoffs during the punishment phase, but she loses the reward W . Denote U_j the total payoff of player j from the end of player i 's punishment phase to stage T . The deviation is not profitable if:

$$PM + U_j + 0 \leq -PM + U_j + W$$

This condition is satisfied if $W \geq W_0$.

After any other history, only Nash equilibria of the game G are played, and therefore no deviation is profitable. The strategy profile σ^* is thus a subgame perfect equilibrium of $G_T(\text{term})$ if $T = K\alpha$, $P \geq P_0$ and $W \geq W_0$.

If T is not a multiple of α , we complete the definition of σ^* by playing a fixed Nash equilibrium of G at the last $T \bmod(\alpha)$ stages of the game, irrespective of the history. For each T , we have thus constructed a subgame perfect equilibrium of $G_T(\text{term})$, and a corresponding payoff $\theta_T(\sigma^*) \in \mathcal{E}_T^*(W)$.

Provided that $T \bmod(\alpha)$ is small with respect to T (i.e. K is large), we have $\|\theta_T(\sigma^*) - u\| \leq \varepsilon$. \square

Next, we show that the parameters of the above strategy can be chosen “uniformly” with respect to the target payoff.

Corollary 6.2 *For all $\varepsilon > 0$, there exist $T_0 \in \mathbb{N}$ and $W_0 \in \mathbb{R}_+$ such that for all $T \geq T_0$, $W \geq W_0$ and $u \in E^*$, there exists $v \in \mathcal{E}_T^*(W)$, $\|u - v\| \leq \varepsilon$.*

Proof. Let $\varepsilon > 0$ and $u \in E^*$. From Proposition 6.1, there exist T_0 and W_0 such that for all $T \geq T_0$ and $W \geq W_0$, there exists $u^* \in \mathcal{E}_T^*(W)$ such that $\|u - u^*\| \leq \varepsilon$. Since E^* is compact, it can be covered by finitely many balls: $E^* \subseteq \bigcup_{l=1}^L B(r_l, \frac{\varepsilon}{2})$ where $r_l \in E^*$ and $B(r_l, \frac{\varepsilon}{2})$ denotes the ball with center r_l and radius $\frac{\varepsilon}{2}$. For each $l \in \{1, \dots, L\}$, there exist T_0^l and W_0^l satisfying Proposition 6.1 for the vector r_l . It is enough to consider $T \geq T_0 := \max_{l \in \{1, \dots, L\}} T_0^l$ and $W \geq W_0 := \max_{l \in \{1, \dots, L\}} W_0^l$. \square

The next lemmas 6.3 and 6.4 adapt respectively lemmas 4.1 and 4.2 in Gossner [9] and establish the link between $\mathcal{E}_T^*(W)$ and \mathcal{E}_T^* . The idea consists in using the last stages of the game to simulate the terminal payoffs.

Lemma 6.3 *Under conditions C1 and C2, there exists $T_1 \in \mathbb{N}$ such that $\dim(\text{co } \mathcal{E}_{T_1}^*) = n$.*

Proof. The proof is almost the same as in Gossner [9]. Under Condition C2, there exist action profiles $x_0, \dots, x_n \in D$ such that $\dim(\text{co } \{g(x_1) - g(x_0), \dots, g(x_n) - g(x_0)\}) = n$.

For each $m = 1, \dots, n$, we define a strategy profile σ_m as follows: at stage $t = 1$, play x_m . If \bar{x}_m is publicly observed, then the Nash equilibrium e^1 is played for Q consecutive stages, e^2 for Q consecutive stages, \dots , and e^n for Q consecutive stages. If player i deviates in a detectable way at stage $t = 1$, then f^i is played instead of e^i (if there are several such players i , one is chosen arbitrarily). For Q large enough, σ_m is a subgame perfect equilibrium of G_{nQ+1} . An undetectable deviation of player i at stage $t = 1$ is not profitable. If player i deviates in a detectable way at stage $t = 1$, she gets at most $M + Qg_i(f^i) \leq -M + Qg_i(e^i)$ for Q large enough, since $g_i(e^i) > g_i(f^i)$.

The payoff induced by σ_m is $(g(x_m) + Qg(e^1) + \dots + Qg(e^n))/(nQ + 1)$ and is in \mathcal{E}_{nQ+1}^* for Q large enough. It follows that $\dim(\text{co } \mathcal{E}_{nQ+1}^*) = n$. \square

As a consequence, for $T_1 \geq nQ + 1$, the set of subgame perfect equilibrium payoffs $\mathcal{E}_{T_1}^*$ is rich enough to generate any terminal payoff, up to a fixed translation in the payoff space. Precisely,

Lemma 6.4 *Under conditions C1 and C2, there exists $\rho_0 > 0$ such that for all $W > 0$, there exist $T_2 \in \mathbb{N}$ and $U \in \mathbb{R}^n$,*

$$\forall r \in \{-W, 0, W\}^n, \exists v \in \mathcal{E}_{T_2}^*, \|r + U - T_2 v\| \leq \rho_0$$

The formal proof is exactly the same as in Gossner [9] and is therefore omitted. The idea is the following. For each k , repeating k times an equilibrium of G_{T_1} , defines an equilibrium of G_{kT_1} . More precisely, letting $T\mathcal{E}_T^*$ be the set of subgame perfect equilibrium *total* payoffs of G_T , one has

$k \cdot T_1 \mathcal{E}_{T_1}^* \subseteq kT_1 \mathcal{E}_{kT_1}^*$. Since $T_1 \mathcal{E}_{T_1}^*$ contains an open ball, for each $W > 0$, there exists an integer k and a vector $U \in \mathbb{R}^n$ such that $\{-U\} + kT_1 \mathcal{E}_{kT_1}^*$ contains $\{-W, 0, W\}^n$. Note that ρ_0 only depends on the game G and on T_1 . We fix T_1 and ρ_0 from now on.

We may now complete the proof of Theorem 4.2.

Proof. Fix $\varepsilon > 0$. From Corollary 6.2, there exists $T_0 \in \mathbb{N}$ and $W_0 \in \mathbb{R}_+$ such that for all $T \geq T_0$, $W \geq W_0$ and $u \in E^*$, there exists $v \in \mathcal{E}_T^*(W)$ such that $\|u - v\| \leq \varepsilon$. Fix a terminal payoff function \mathbf{term} and a subgame perfect equilibrium σ^* of $G_T(\mathbf{term})$ such that $v = \theta_T(\sigma^*)$, as constructed in the proof of Proposition 6.1. Lemma 6.4 gives $T_2 \in \mathbb{N}$ and $U \in \mathbb{R}^n$ such that for all $h \in H_{p,T}$, there exists $v(h) \in \mathcal{E}_{T_2}^*$ such that $\|\mathbf{term}(h) + U - T_2 v(h)\| \leq \rho_0$. Since adding the constant vector U to the terminal payoff does not change the equilibrium constraints, σ^* is a subgame perfect equilibrium of $G_T(\mathbf{term} + U)$ and we may assume w.l.o.g. $U = 0$. Define then $\mathbf{term}'(h) = T_2 v(h)$.

It is easy to check that, if $W \geq W_0 + 2\rho_0$, then σ^* is a subgame perfect equilibrium of $G_T(\mathbf{term}')$. The terminal payoff is used to deter deviations from punishment phases and late deviations. Following a punishment phase of player i , the terminal payoff \mathbf{term}' of player j is between $W - \rho_0$ and $W + \rho_0$ if she was an effective punisher, and is between $-\rho_0$ and ρ_0 otherwise. This is enough to provide incentives to confirm with the punishment phase. The same reasoning holds for late deviations.

Finally, we define a strategy profile σ^{**} of G_{T+T_2} as follows. The profile σ^* is played until stage T . Then if h is the sequence of actions played up to stage T , a subgame perfect equilibrium of G_{T_2} with payoff $v(h)$ is played.

From the construction, σ^{**} is a subgame perfect equilibrium of G_{T+T_2} . The induced payoff is close to v and thus close to u . Indeed, let h^* be the history induced by σ^* at stage T . By construction, $\text{term}(h^*) = 0$, so $v = \gamma_T(\sigma^*)$ and $w := \gamma_{T+T_2}(\sigma^{**}) = \frac{T}{T+T_2}\gamma_T(\sigma^*) + \frac{T_2}{T+T_2}v(h^*)$. It follows,

$$\begin{aligned} \|v - w\| &\leq \left\| \gamma_T(\sigma^*) \frac{T_2}{T+T_2} \right\| + \left\| \frac{T_2}{T+T_2} v(h^*) \right\| \\ &\leq \frac{T_2}{T+T_2} M + \frac{1}{T+T_2} (\rho_0 + \|U\|) \end{aligned}$$

which is arbitrarily small for T large enough. □

References

- [1] Aumann, R.J., Shapley L.S. (1976), *Long-term competition - A game theoretic analysis*, re-edited in 1994, in N. Megiddo, editor, *Essays on game theory*, pages 1–15. Springer-Verlag, New York.
- [2] Benoît, J.P., Krishna, V. (1985), *Finitely repeated games*, *Econometrica*, 53, 905-922.
- [3] Benoît, J.P., Krishna, V. (1987), *Nash equilibria of finitely repeated games*, *International Journal of Game Theory*, 16(3), 197-204.
- [4] Fudenberg, D., D.K. Levine (1994), *Efficiency and observability with long-run and short-run players*, *Journal of Economic Theory*, 62, 103-135.

- [5] Fudenberg, D., D.K. Levine and E. Maskin (1994), *The Folk theorem with imperfect public information*, *Econometrica*, 62(5), 997-1039.
- [6] Fudenberg, D., E. Maskin (1986), *The folk theorem in repeated games with discounting or with incomplete information*, *Econometrica*, 54, 533-554.
- [7] Fudenberg, D., D.K. Levine and S. Takahashi (2007), *Perfect public equilibrium when players are patient*, *Games and Economic Behavior*, 61(1), 27-49.
- [8] González-Díaz, J. (2006), *Finitely repeated games: a generalized Nash folk theorem*, *Games and Economic Behavior*, 55, 100-111.
- [9] Gossner, O. (1995), *The Folk Theorem for Finitely Repeated Games with Mixed Strategies*, *International Journal of Game Theory*, 24, 95-107.
- [10] Lehrer, E. (1989), *Lower equilibrium payoffs in two-player repeated games with non-observable actions*, *International Journal of Game Theory*, 18(1), 57-89.
- [11] Lehrer, E. (1990), *Nash equilibria of n-player repeated games with semi-standard information*, *International Journal of Game Theory*, 19, 191-217.
- [12] Lehrer, E. (1992), *Two-player repeated games with non observable actions and observable payoffs*, *Mathematics of Operations Research*, 17, 200-224.

- [13] Mailath, G.J, S. A. Matthews and T. Sekiguchi (2002), *Private Strategies in Finitely Repeated Games with Imperfect Public Monitoring*, The B.E. Journal in Theoretical Economics, vol. 2, issue 1.
- [14] Mailath, G.J. and L. Samuelson (2006), *Repeated Games and Reputations: Long-Run Relationships*, Oxford University Press.
- [15] Renault, J., Scarlatti S. and M. Scarsini (2008), *Discounted and finitely repeated minority games with public signals*, Mathematical Social Sciences, 56, 44-74.
- [16] Renault, J. and T. Tomala (2004), *Communication equilibrium payoffs of repeated games with imperfect monitoring*, Games and Economic Behavior, 49, 313-344.
- [17] Rubinstein, A. (1977) *Equilibrium in supergames*, Center for Research in Mathematical Economics and Game Theory, Research Memorandum 25, 1977.
- [18] Sekiguchi, T. (2001) *A negative result in finitely repeated games with product monitoring*, Economic Letters, 74, 67-70.
- [19] Smith, L. (1995) *Necessary and sufficient conditions for the perfect finite horizon folk theorem*, Econometrica, 62, 949-954.
- [20] Sorin, S. (1986) *On repeated games with complete information*, Mathematics of Operations Research, 11, 147-160.
- [21] Tomala, T. (1998), *Pure equilibria with public observation*, International Journal of Game Theory, 27(1), 93-109.

- [22] Wen, Q. (1994), *The “folk theorem” for repeated games with complete information*, *Econometrica*, 62(4), 949-954.