A folk theorem for finitely repeated games with public monitoring

Johannes Hörner[*], Jerôme Renault[†]

April 25, 2011

Abstract

We adapt the methods from Abreu, Pearce and Stacchetti (1990) to finitely repeated games with imperfect public monitoring. Under a combination of (a slight strengthening of) the assumptions of Benoît and Krishna (1985) and those of Fudenberg, Levine and Maskin (1994), a folk theorem follows.

Keywords: Repeated games.

JEL codes: C72, C73

1 Introduction

The literature on finitely and infinitely repeated games have proceeded somewhat independently in the last twenty years. Following Abreu, Peace and Stacchetti (1990), tremendous progress has been accomplished in the analysis of infinitely repeated games with imperfect monitoring. Results in this literature have build on the fixed-point characterization that they give of the set of public perfect equilibrium payoffs, which allowed for a largely non-constructive analysis of the equilibrium payoff set, in particular, as discounting vanishes. See, in particular, Fudenberg, Levine and Maskin (1994, hereafter FLM).

Clearly, no such fixed-point characterization exists in the case of finitely repeated games, as the (public perfect) equilibrium payoff set is not independent of the number of periods left.

[*]Yale University, 30 Hillhouse Ave., New Haven, CT 06520, USA, johannes.horner@yale.edu.
[†]TSE (GREMAQ, Université Toulouse 1), 21 allée de Brienne, 31000 Toulouse, France, jerome.renault@tse-fr.eu.

As a result, folk theorems for this case have required the explicit specification of equilibrium strategies. This has forced these authors to restrict attention to perfect monitoring. See Benoît and Krishna (1985), Gossner (1995) and Smith (1995), as well as Benoît and Krishna (1996) for a survey.[1]

Yet the (Bellman-Shapley) operator involved in the definition of the fixed point applies just as well to the case of a finite horizon, giving us an immediate link between the equilibrium payoff sets that obtain as the horizon length varies. Similarly, the main idea behind the proof of FLM applies as well. Suppose that the (average) equilibrium payoff set converges to a strict subset of the feasible, and individually rational payoff set V. Then there exists a direction in which this limit set has a boundary that is locally smooth (in an appropriately defined sense) yet bounded away from the extreme point of V in that direction. Under the assumptions of FLM, then, a contradiction can be derived.

Some care must be taken in this argument, however. First, this limit set must have non-empty interior. Clearly, this requires an assumption involving the payoffs of the stage game. Indeed, if the stage game admits a unique Nash equilibrium payoff, then the finitely repeated game admits a unique (perfect public) equilibrium payoff. Therefore, we must follow Benoît and Krishna (1985) and assume that the stage game admits distinct Nash equilibrium payoffs for each player, and we strengthen this assumption by assuming that the convex hull of this set of vectors has non-empty interior. Furthermore, in the absence of discounting, the relative weights of current vs. continuation average payoffs in the definition of the operator are related to the number of periods left. Therefore, it is not exactly the same operator that is being applied repeatedly. As longer and longer horizons are considered, and as flow payoffs are assigned a vanishing weight, one must make sure that these weights do not decrease too fast in order for a contradiction to obtain.

While there has been few systematic analyses of finitely repeated games since Benoît and Krishna (1985) and Gossner (1995), very interesting examples have been produced by Mailath, Matthews and Sekiguchi (2002) to show how non-trivial equilibria can be constructed in finitely repeated with imperfect monitoring even as the stage game admits a unique Nash equilibrium. Such equilibria involve private strategies, and so are not covered by our analysis. However, all that matters for our argument is that it is known that the finitely repeated game admits a set of equilibrium payoffs whose convex hull has non-empty interior for some horizon length. This

[1] It might be argued that Gossner's construction is not entirely constructive, though, as his main innovation relies in applying approachability. Still, the overall structure of the argument follows Benoît and Krishna.

length can then be as "end-game" involving private strategies, and treated as a "blackbox" when defining public strategies. Other related recent contributions involve Contou-Carrère and Tomala (2010) for the case of semi-standard monitoring, and Gonzàlez-Dìaz (2006) for the case of Nash equilibria. Standard results or definitions will follow Mailath and Samuelson (2006, hereafter M&S).

2 Notation and Assumptions

We consider finitely repeated games between I players. Notations mostly follow M&S. Actions sets are A_i, finite, with generic element a_i, and given action profile a, there is a public signal y from a finite set Y that is publicly observed. The distribution of signals given a is denoted $\pi(\cdot|a)$. Rewards in the stage game are given by $u_i(a)$ for player i, given action profile. Action profiles are not observed, nor are realized payoffs. Let F denote the set of feasible payoffs, V the set of feasible and individually rational payoffs (where individual rationality is, as usual, defined with respect to the mixed minmax payoff, in which players $-i$ randomize independently). We maintain throughout:

Assumption 1 (A1) *The set V has non-empty interior.*

We maintain the simplest rank assumptions of Fudenberg, Levine and Maskin (1994). The profile α has individual full rank for i the $|A_i| \times |Y|$-matrix

$$\Pi_i(\alpha) = \left(\pi\left(y|a_i, \alpha_{-i}\right)\right)_{a_i,y}$$

has full row rank, i.e. the probability distributions $\{\pi(\cdot|a_i, \alpha_{-i}) : a_i \in A_i\}$ are linearly independent. It has individual full rank if it has individual full rank for all i. It has pairwise full rank for i and j if the $(|A_i| + |A_j|) \times |Y|$-matrix

$$\Pi_{ij}(\alpha) = \begin{pmatrix} \Pi_i(\alpha) \\ \Pi_j(\alpha) \end{pmatrix}$$

has rank $|A_i| + |A_j| - 1$ (i.e., maximal rank).

Assumption 2 (A2) *All the pure action profiles yielding extreme points of F have pairwise full rank for all pairs of players.*

Player i's minmax payoff is denoted u^i, and the corresponding (possibly mixed, but uncorrelated) action profile m^i.

Assumption 3 (A3) *All the minmax profiles m^i have individual full rank.*

We let $E_n \subset F$ denote the set of average equilibrium payoffs in the game that is repeated n times (without discounting). Throughout, equilibrium refers to public perfect equilibrium (see, for instance, FLM). It is known that E_n converges in the Hausdorff sense (See Renault and Tomala (2011)), though the speed of convergence is unknown. We denote by E the limit of this set, which is convex (as opposed to E_n, which typically is not). Clearly, $E_1 \subset E$. Hence, a sufficient condition that guarantees that E has non-empty interior is to assume that the convex hull of E_1 has non-empty interior, i.e., there exists distinct Nash equilibrium payoff vectors whose convex hull has non-empty interior. This is a strengthening of the assumption of Benoît and Krishna (1985).

3 A Folk Theorem with a Randomization Device

In this section, we assume a public randomization device. This drastically simplifies arguments, as it guarantees that E_n is convex, for all n.

Theorem 4 *Assume (A1)–(A3). If E has non-empty interior, then $E = V$, the set of feasible and individually rational payoffs.*

Proof. Recall that E is convex. Therefore, by a theorem due to Alexandrov, almost all its boundary points are normal, i.e. the representing function is differentiable at these points. Therefore, if $E \neq V$, there exists a vector $\lambda \in S^1 := \left\{x \in \mathbb{R}^I : \|x\| = 1\right\}$ and a point $v \in bd(E)$ with (unique) normal vector λ, and such that $\max_{v' \in V} \lambda \cdot (v' - v) = 2\kappa > 0$.

Assume first that λ is not a coordinate direction, and let l be the line through v with direction λ. In what follows, $a \in A$ refers to an action profile such that $\lambda \cdot g(a) = \max_{v' \in V} \lambda \cdot v'$ (fix one of them if multiple exist). Let $\bar{v} \in \mathbb{R}^I$ denote the unique point of l such that $\lambda \cdot \bar{v} = \lambda \cdot g(a)$. Clearly, $\bar{v} \notin E$. For all $k \in \mathbb{R}$, let $H_\lambda(k) := \left\{x \in \mathbb{R}^I : \lambda \cdot x = k\right\}$, and $H_\lambda^+(k) := \left\{x \in \mathbb{R}^I : \lambda \cdot x \geq k\right\}$.

The pairwise full rank assumption at the action profile a ensure that there exists $x : Y \to \mathbb{R}^I$ such that a is a Nash equilibrium of the game with payoff function

$$g(\cdot) + \sum \pi(y|\cdot)\, x(y),$$

and $\lambda \cdot x(y) = 0$ for all $y \in Y$ (See M&S, Lemma 8.1.1(4) and 9.2.2). Without loss of generality, we may assume that the payoff from this equilibrium is equal to $\bar{v}$ (redefine each $x(y)$ as $x(y) + \bar{v} - g(a) - \sum_{y \in Y} \pi(y|a)x(y)$). Let $M = \max_y \|x(y)\|$, and set $\kappa_0 := \frac{\kappa}{2}/\sqrt{\kappa^2 + M^2}$.

Given $x \in \mathbb{R}^I$, let

$$C_x := \left\{x' \in \mathbb{R}^I : \lambda \cdot (x - x') > \kappa_0 \|x - x'\|\right\}.$$

See Figure 1. Because E is smooth at v, there exists $k < \lambda \cdot v$ and a compact set $D \subset \mathbb{R}^I$ such that

$$H_\lambda(k) \cap C_v \subset D \subset int(E).$$

Because $E_n \to E$, we can assume without loss that $D \subset E_n$ for all n. Let $v_n = \arg\max_{x \in l \cap E_n} \lambda \cdot x$ be the highest point of E_n on the line l, and set $k_n = \lambda \cdot v_n$. We restrict attention to $n \geq n_0$ such that $k_n - \cdot\kappa/n > k$. This implies that $C_{v_n} \cap H_\lambda^+(k) \in E_n$, because $C_{v_n} \subset C_v$, v_n and D are in E_n, and E_n is convex.

Given v_n, let

$$w_n(y) := v_n + \frac{1}{n}x(y) - \frac{\kappa}{n}\lambda,$$

for all $y \in Y$. Note that

$$\frac{\lambda \cdot (v_n - w_n(y))}{\|v_n - w_n(y)\|} \geq \frac{\kappa}{\sqrt{\kappa^2 + M^2}} > \kappa_0,$$

so that $w_n(y) \in C_{v_n}$ and hence $w_n(y) \in E_n$ (note that $\lambda \cdot w_n(y) = k_n - \kappa/n > k$). Note that the action profile a is a Nash equilibrium of the game with payoff vector

$$\begin{aligned} & \frac{1}{n+1}g(\cdot) + \frac{n}{n+1}\sum \pi(y|\cdot)\, w_n(y) \\ = \; & \frac{1}{n+1}\left(g(\cdot) + \sum \pi(y|\cdot)\, x(y)\right) + \frac{n}{n+1}v_n - \frac{\kappa\lambda}{n+1}. \end{aligned}$$

Because $w_n \in E_n$, this implies that the resulting equilibrium payoff vector h_{n+1} is in E_{n+1}. Observe that

$$\lambda \cdot (h_{n+1} - v_n) = \frac{1}{n+1}\lambda \cdot (g(a) - v_n) - \frac{\kappa}{n+1} \geq \frac{\kappa}{n+1}.$$

Furthermore, observe that, by construction, h_{n+1} is on the line l (recall that $g(a) + \sum \pi(y|a)\, x(y) = \bar{v} \in l$). Hence,

$$k_{n+1} \geq \lambda \cdot h_{n+1} \geq \lambda \cdot v_n + \frac{\kappa}{n+1} = k_n + \frac{\kappa}{n+1}.$$

This is not possible, as it implies that $k_n \to \infty$. Consider the case in which λ is a coordinate

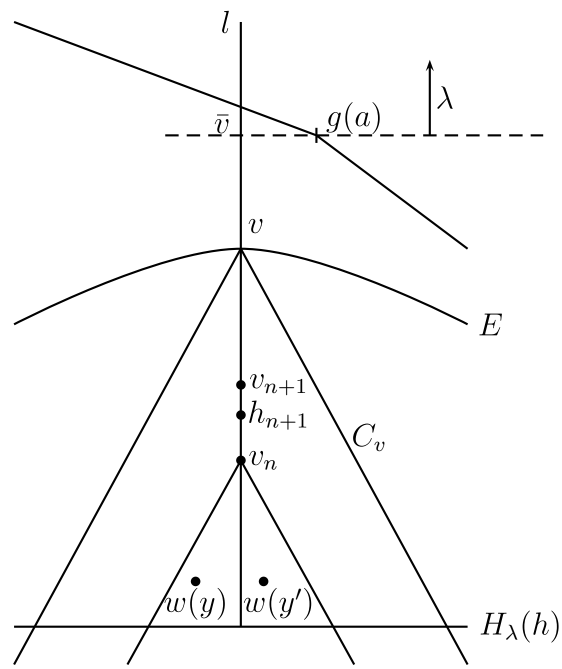

Figure 1: Proof of Lemma 4

direction, i.e. $\lambda = \pm e^i$, for some basis vector e^i. Then it is also the case, by Assumption 2 and 3 (see M&S, Lemma 9.2.1(1) and (3)) that there exists $x(y)$ such that $\lambda \cdot x(y) = 0$ for all y, and such that the action profile a (resp. m^i) that maximizes $g_i(a)$ (resp. minmaxes player i) is an equilibrium of the game with payoffs $g(\cdot) + \sum_{y \in Y} \pi(y|\cdot)x(y)$. In both cases, the remainder of the proof is identical to the previous one. ■

4 Dispensing with the Public Randomization Device

The proof can be extended to the case in which no public randomization device is assumed. There are two difficulties.

First, continuation payoffs cannot be fine-tuned to be precisely those used in Theorem 1. This is not a major difficulty as long as directions $\lambda \in \mathbb{R}^I$ are considered for which score is maximized by a pure action profile. In that case, the same action profile can be enforced by giving strict incentives to the players, and so what matters is whether the set E_n is "rich" enough that we can approximate the desired continuations payoffs by actual available elements of E_n, while preserving these strict incentives. This can be done constructively, assuming full support of the monitoring.

Assumption 5 (B1) *There exists a partition of the set of signals Y into subsets Y_1, Y_2, such that, for all $a \in A$, $k = 1, 2$, $\sum_{y \in Y_k} \pi(y|a) > 0$.*

Of course, the result is trivial (or rather, follows from Benoît and Krishna (1985)) in the case of perfect monitoring. We suspect that Assumption (B1) can be dropped by a direct constructive argument whenever there is not full support, but have not verified the details.

On the other hand, a difficulty arises while considering negative coordinate directions, which correspond to the minimization of a player's payoff. If player i's payoff is to be minimized, players $-i$ might have to use a mixed strategy to minmax i. The question then arises, whether incentives to (approximately) minmax i can be given to player $-i$ for some continuation payoffs that can only be chosen approximately (i.e., within some neighborhood). This can easily be done with two players, or of course when the minmaxing action profile is pure. More generally, consider the minmaxing action profile m^i_{-i}. It is known that there is a best-reply for player i for which the resulting action profile m^i is a Nash equilibrium of the ("team zero-sum") game in which player i's payoff function is u_i, while player $-i$'s payoff function is $-u_i$. See von Stengel and Koller (1997) for details. It suffices that games with payoff functions that are sufficiently close to $(u_i, -u_i)$ admit an equilibrium whose payoff is close to the minmax payoff u^i. This is ensured, for instance, if m is an essential equilibrium. See Wu and Jiang (1962).

Assumption 6 (B2) *Assume that for all i, either (i) m^i_{-i} involves at most one player j $(\neq i)$ not using a pure strategy, or (ii) m^i is an essential equilibrium.*

Note that (i) covers the case of two players. In fact, it is essentially the two-player case, as with more players, players $k \neq i, j$ can be given strict incentives to play the pure action that m specifies. We suspect that this assumption can be dispensed with altogether, but have been unable to do so so far. Alternatively, this assumption can be dropped if the minmax payoff is defined with respect to pure strategies.

We can then show:

Theorem 7 *Under Assumptions (A1)–(A3) and (B1)–(B2), if E has non-empty interior, then $E = V$.*

Proof. To be completed. ■

References

[1] Abreu, D., D. Pearce, and E. Stacchetti (1990). "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica*, **58**, 1041–1063.

[2] Benoît, J.P. and V. Krishna (1985). "Finitely repeated games," *Econometrica*, **53**, 905–922.

[3] Benoît, J.P. and V. Krishna (1996). "The Folk Theorems for Repeated Games: A synthesis," Econ-WPA, C-72.

[4] Contou-Carrère, P. and T. Tomala (2010). "Finitely repeated games with semi-standard monitoring," mimeo.

[5] Fudenberg, D. and D.K. Levine (1994). "Efficiency and observability with long-run and short-run players," *Journal of Economic Theory*, **62**, 103–135.

[6] Fudenberg, D., D.K. Levine and E. Maskin (1994). "The Folk theorem with imperfect public information," *Econometrica*, **62**(5), 997–1039.

[7] González-Díaz, J. (2006). "Finitely repeated games: a generalized Nash folk theorem," *Games and Economic Behavior*, **55**, 100–111.

[8] Gossner, O. (1995). "The Folk Theorem for Finitely Repeated Games with Mixed Strategies," *International Journal of Game Theory*, **24**, 95–107.

[9] 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*, **2**(1).

[10] Mailath, G.J. and L. Samuelson (2006). *Repeated Games and Reputations: Long-Run Relationships*, Oxford University Press.

[11] Renault, J. and T. Tomala (2011). "General properties of long-run supergames," mimeo.

[12] Sekiguchi, T. (2001). "A negative result in finitely repeated games with product monitoring," *Economic Letters*, **74**, 67–70.

[13] Smith, L. (1995). "Necessary and Sufficient Conditions for the Perfect Finite Horizon Folk Theorem," *Econometrica*, **63**, 425-430.

[14] von Stengel B. and D. Koller (1997). "Team-Maxmin Equilibria," *Games and Economic Behavior*, **21**, 309–321.

[15] Wu Wen-Tsun and Jiang Jia-He (1962). "Essential equilibrium points of n-person non-cooperative games," *Scientia Sinica*, **11**, 1307–1322.