ECONSTOR
Make Your Publications Visible.

A Service of
ZBW
Leibniz-Informationszentrum Wirtschaft
Leibniz Information Centre for Economics

Demeze-Jouatsa, Ghislain H.

Working Paper

A complete folk theorem for finitely repeated games

Center for Mathematical Economics Working Papers, No. 584

Provided in Cooperation with:
Center for Mathematical Economics (IMW), Bielefeld University

Suggested Citation: Demeze-Jouatsa, Ghislain H. (2018) : A complete folk theorem for finitely repeated games, Center for Mathematical Economics Working Papers, No. 584, Bielefeld University, Center for Mathematical Economics (IMW), Bielefeld, https://nbn-resolving.de/urn:nbn:de:0070-pub-29303825

This Version is available at:
https://hdl.handle.net/10419/201608

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

WWW.ECONSTOR.EU

Mitglied der Leibniz-Gemeinschaft

Center for
Mathematical Economics
Working Papers

584

July 2018

A complete folk theorem for finitely repeated games

Ghislain H. Demeze-Jouatsa

Center for Mathematical Economics (IMW)
Bielefeld University
Universitätsstraße 25
D-33615 Bielefeld · Germany

e-mail: imw@uni-bielefeld.de
http://www.imw.uni-bielefeld.de/wp/
ISSN: 0931-6558

A complete folk theorem for finitely repeated games

By Ghislain H. DEMEZE-JOUATSA, Bielefeld University[1]

email: demeze_jouatsa@uni-bielefeld.de

This version: July 31, 2018

Abstract: I analyze the set of pure strategy subgame perfect Nash equilibria of any finitely repeated game with complete information and perfect monitoring. The main result is a complete characterization of the limit set, as the time horizon increases, of the set of pure strategy subgame perfect Nash equilibrium payoff vectors of the finitely repeated game. The same method can be used to fully characterize the limit set of the set of pure strategy Nash equilibrium payoff vectors of any the finitely repeated game.

Keywords: Finitely Repeated Games, Pure Strategy, Subgame Perfect Nash Equilibrium, Limit Perfect Folk Theorem, Discount Factor.

JEL classification: C72, C73.

1 Introduction

This paper provides a full characterization of the limit set, as the time horizon increases, of the set of pure strategy subgame perfect Nash equilibrium payoff vectors of any finitely repeated game. The obtained characterization is in terms of appropriate notions of feasible and individually rational payoff vectors of the stage-game. These notions are based on Smith's (1995) notion of Nash decomposition and appropriately generalize the classic notion of feasible payoff vectors as well as the notion of effective minimax payoff defined by Wen (1994). The main theorem nests earlier results of Benoit and Krishna (1984) and Smith (1995). Using a similar method, I obtain a full characterization of the limit set, as the time horizon increases, of the set of pure strategy Nash equilibrium payoff vectors of any finitely repeated game. The obtained result nests earlier results of Benoit and Krishna (1987).

Whether non-Nash outcomes of the stage-game can be sustained by means of subgame perfect Nash equilibria of the finitely repeated game depends on whether players

[1]Ghislain H. DEMEZE-JOUATSA acknowledges DAAD for funding this research project and thanks Christoph Kuzmics, Frank Riedel, Lones Smith, Michael Greinecker, Karl Schlag and Olivier Gossner for useful comments.

can be incentivized to abandon their short term interests and to follow some collusive paths that have greater long-run average payoffs. There are two extreme cases. On the one hand, in any finite repetition of a stage-game that has a unique Nash equilibrium payoff vector such as the prisoners' dilemma, only the stage-game Nash equilibrium payoff vector is sustainable by subgame perfect Nash equilibria of finite repetitions of that stage-game. The underlying reason is that in the last round of the finitely repeated game, players can agree only on Nash equilibria of the stage-game as no future retaliation is possible. Backwardly, the same argument works at each round of the finitely repeated game since each player has a unique continuation payoff for the upcoming rounds. On the other hand, for stage-games in which all players receive different Nash equilibrium payoffs as the battle of sexes, the limit perfect folk theorem hold: Any feasible and individually rational payoff vector of the stage-game is achievable as the limit payoff vector of a sequence of subgame perfect Nash equilibria of the finitely repeated game as the time horizon goes to infinity.

Benoit and Krishna (1984) established that for the limit perfect folk theorem to hold, it is sufficient that the dimension of the set of feasible payoff vectors of the stage-game equals the number of players and that each player receives distinct payoffs at Nash equilibria of the stage-game.[2] Smith (1995) provided a weaker, necessary and sufficient condition for the limit perfect folk theorem to hold. Smith (1995) showed that it is necessary and sufficient that the Nash decomposition of the stage-game is complete; as I explain below. The distinct Nash payoffs condition and the full dimensionality of the set of feasible payoff vectors as in Benoit and Krishna (1984) or the complete Nash decomposition of Smith (1995) allow us to construct credible punishment schemes and to (recursively) leverage the behavior of any player near the end of the game. These are essential to generate a limit perfect folk theorem. In the case that the stage-game admits a unique Nash equilibrium payoff vector, Benoit and Krishna (1984) demonstrated that the set of subgame perfect Nash equilibrium payoff vectors of the finitely repeated game is reduced to the unique stage-game Nash equilibrium payoff vector.

A part of the puzzle remains unresolved. Namely, for a stage-game that does not admit a complete Nash decomposition, what is the exact range of payoff vectors that are achievable as the limit payoff vector of a sequence of subgame perfect Nash equilibria of finite repetitions of that stage-game?

[2] Fudenberg and Maskin (1986) introduced the notion of full dimensionality of the set of feasible payoff vectors and used it to provide a sufficient condition for the perfect folk theorem for infinitely repeated games.

The Nash decomposition of a normal form game is a strictly increasing sequence of non-empty groups of players. Players of the first group are those who receive at least two distinct Nash equilibrium payoffs in the stage-game. The second group of players of the Nash decomposition, if any, contains each player of the first group as well as some new players. New players are those who receive at least two distinct Nash equilibrium payoffs in the new game that is obtained from the stage-game by setting the utility function of each player of the first group equal to a constant. This idea can be iterated. After a finite number of iterations, the player set no longer changes. The Nash decomposition is complete if its last element equals the whole set of players.

If the stage-game has an incomplete Nash decomposition, then the set of players naturally breaks up into tow blocks where the first block contains all the players whose behavior can recursively be leveraged near the end of the finitely repeated game. In contrast, it is not possible to control short run incentives of players of the second block of the latter partition. Therefore, each player of the second block has to play a stage-game pure best response at any profile that occurs on a pure strategy subgame perfect Nash equilibrium play path. Stage-game action profiles eligible for pure strategy subgame perfect Nash equilibrium play paths of the finitely repeated game are therefore exactly the stage-game pure Nash equilibria of what one could call the effective one shot game, the game obtained from the initial stage-game by setting the utility function of each player of the first block equal to a constant.

This restriction of the set of eligible actions for pure strategy subgame perfect Nash equilibrium play paths has two main implications. Firstly, for a feasible payoff vector to be approachable by pure strategy subgame perfect Nash equilibria of the finitely repeated game, it has to be in the convex hull of the set of Nash equilibrium payoff vectors of the effective one shot game. I introduce the concept of a recursively feasible payoff vector. I call a payoff vector recursively feasible if it belongs to the convex hull of the set of payoff vectors to profile of actions that are Nash equilibria of the effective one shot game. Secondly, as subgame perfect Nash equilibria are protected against unilateral deviations even off equilibrium paths, any player of the second block has to be at her best response at any action profile occurring on a credible punishment path. Therefore, only pure Nash equilibria of the effective one shot game are eligible for credible punishment paths in any finite repetition of the original stage-game. Consequently, a player of the first block can guarantee herself a payoff that is strictly greater than her effective minimax payoff. I call this payoff the recursive effective minimax payoff.

The main finding of this paper says that, as the time horizon increases, the set of

payoff vectors of pure strategy subgame perfect Nash equilibria of the finitely repeated game converges to the set of recursively feasible payoff vectors that dominate the recursive effective minimax payoff vector.

The paper proceeds as follows. In Section 2 I introduce the model and the definitions. Section 3 states the main finding of the paper and sketches the proof. In Section 4, I discuss some extensions and Section 5 concludes the paper. Proofs are provided in the Appendices.

2 Model and definitions

2.1 The Stage-game

Let $G = (N, A = \times_{i \in N} A_i, u = (u_i)_{i \in N})$ be a stage-game where the set of players $N = \{1, ..., n\}$ is finite and where for all player $i \in N$ the set A_i of actions of player i is compact. Given player $i \in N$ and an action profile $a = (a_1, ..., a_n) \in A$, let $u_i(a)$ denote the stage-game utility of player i given the action profile a. Given an action profile $a \in A$, $i \in N$ a player, and $a'_i \in A_i$ an action of player i, let (a'_i, a_{-i}) denote the action profile in which all players except player i choose the same action as in a, while player i chooses a'_i. A stage-game pure best response of player i to the action profile a is an action $b_i(a) \in A_i$ that maximizes the stage-game payoff of player i given that the choice of other players is given by a_{-i}. An action profile $a \in A$ is a **pure Nash equilibrium of the stage-game** G (denoted by $a \in \text{Nash}(G)$) if $u_i(a'_i, a_{-i}) \leq u_i(a)$ for all player $i \in N$ and all action $a'_i \in A_i$.

Let γ be a real number that is strictly greater than any payoff a player might receive in the stage-game G.[3] A player is said to have to have distinct pure Nash payoffs in the stage-game if there exist two pure Nash equilibria of the stage-game in which this player receives different payoffs. Let $\tau(G) = (N, A, (u'_i)_{i \in N})$ be the normal form game where the utility function of player i is defined by

$$u'_i = \begin{cases} \gamma & \text{if } i \text{ has distinct Nash payoffs in } G \\ u_i & \text{otherwise} \end{cases}.$$

Let $G^0 := G$ and $G^{l+1} := \tau(G^l)$ for all $l \geq 0$. For all $l \geq 0$, let N_l be the set of players with a utility function that is constant to γ in the game G^l. As N is finite, there is an $h \in [0, +\infty)$ such that $N_{l+1} = N_l$ for all $l \geq h$. Let $\widetilde{A} = \text{Nash}(G^h)$ be the set of pure Nash equilibria of the game G^h.

[3] As the set A of action profiles is compact and the utility function u is continuous on A, the set $u(A) = \{u(a) \mid a \in A\}$ is compact and therefore bounded. This guarantee the existence of γ.

Definition 1 *The set of* ***recursively feasible payoff vectors*** *of the game G is defined as the convex hull $Conv[u(\widetilde{A})]$ of the set $u(\widetilde{A}) = \{u(a) \mid a \in \widetilde{A}\}$.*

Let $\sim$ be the equivalence relation defined on the set of players as follows: Player i is equivalent to j (denoted by $i \sim j$) if there exists $\alpha_{ij} > 0$ and $\beta_{ij} \in \mathbb{R}$ such that for all $a \in \widetilde{A}$, we have $u_i(a) = \alpha_{ij} \cdot u_j(a) + \beta_{ij}$. For all $i \in N$, let $\mathcal{J}(i)$ be the equivalence class of player i and let

$$\widetilde{\mu}_i = \min_{a \in \widetilde{A}} \max_{j \in \mathcal{J}(i)} \max_{a'_j \in A_j} \left[\alpha_{ij} \cdot u_j(a'_j, a_{-j}) + \beta_{ij}\right]$$

and $\widetilde{\mu} = (\widetilde{\mu}_1, \cdots, \widetilde{\mu}_n)$.

If the stage-game G does not have any pure Nash equilibrium, then the set of pure strategy subgame perfect Nash equilibrium payoff vectors of the finitely repeated game is empty. If the stage-game G admits at least one pure Nash equilibrium, then $\widetilde{A}$ is non-empty and $\widetilde{\mu}$ is well defined.

Definition 2 *The payoff $\widetilde{\mu}_i$ is the* ***recursive effective minimax*** *of player i in the stage-game G.*

Call a payoff vector recursively individually rational if it dominates the recursive effective minimax payoff vector $\widetilde{\mu}$. Let $\widetilde{I} = \{x = (x_1, \ldots, x_n) \in \mathbb{R} \mid x_i \geq \widetilde{\mu}_i \text{ for all } i \in N\}$ be the set of recursively individually rational payoff vectors.

2.2 The Finitely Repeated Game

Let G be the stage-game. Given $T > 0$, let $G(T)$ denote the $T-$repeated game obtained by repeating the stage-game T times. A pure strategy of player i in the repeated game $G(T)$ is a contingent plan that provides for each history the action chosen by player i given this history. That is, a strategy is a map $\sigma_i : \bigcup_{t=1}^{T} A^{t-1} \to A_i$ where A^0 contains only the empty history. The strategy profile $\sigma = (\sigma_1, ..., \sigma_n)$ of $G(T)$ generates a **play path** $\pi(\sigma) = [\pi_1(\sigma), ..., \pi_T(\sigma)] \in A^T$ and player $i \in N$ receives a sequence $(u_i(\pi_t(\sigma))_{1 \leq t \leq T}$ of payoffs. The preferences of player $i \in N$ among strategy profiles are represented by the average utility $u_i^T(\sigma) = \frac{1}{T} \sum_{t=1}^{T} u_i[\pi_t(\sigma)]$.

A strategy profile $\sigma = (\sigma_1, ..., \sigma_n)$ is a **pure strategy Nash equilibrium of** $G(T)$ if $u_i^T(\sigma'_i, \sigma_{-i}) \leq u_i^T(\sigma)$ for all $i \in N$ and for all pure strategies σ'_i of player i.

A strategy profile $\sigma = (\sigma_1, ..., \sigma_n)$ is a **pure strategy subgame perfect Nash equilibrium** of $G(T)$ if given any $t \in \{1, ..., T\}$ and any history $h^t \in A^{t-1}$, the restriction

$\sigma_{|h^t}$ of σ to the history h^t is a Nash equilibrium of the finitely repeated game $G(T-t+1)$.

Let d be the Euclidean distance of $\mathbb{R}^n$, A and B be two closed and bounded non-empty subsets of the metric space $(\mathbb{R}^n, d)$.[4] The Hausdorff distance (based on d) between A and B is given by

$$d_H(A, B) = \max\left\{\sup_{x\in A} d(x, B), \sup_{y\in B} d(y, A)\right\},$$

where $d(x, Y) = \inf_{y\in Y} d(x, y)$.

For any $T > 0$, let $E(T)$ be the set of subgame perfect Nash equilibrium payoff vectors of $G(T)$. Let E be such that the Hausdorff distance between $E(T)$ and E goes to 0 as T goes to infinity. The set E is the Hausdorff limit of the set of subgame perfect Nash equilibrium payoff vectors of the finitely repeated game. As I show later in the Appendix 1, the limit set E exists and is unique.

3 Main result

Theorem 1 *Let G be a normal form stage-game with a finite number of players and a compact set of action profiles. As the time horizon increases, the set of pure strategy subgame perfect Nash equilibrium payoff vectors of the finitely repeated game converges (in the Hausdorff sense) to the set of recursively feasible and recursively individually rational payoff vectors.*

The proof of Theorem 1 is provided in the Appendix 1. It consists of four steps that I describe below.

First step. Using the Hausdorff distance, I show that the limiting set E is well defined. This means that, as the time horizon increases, the set of subgame perfect Nash equilibrium payoff vectors of the finitely repeated game converges. The main ingredient of this proof is the conjunction lemma borrowed from Benoit and Krishna (1984); see Lemma 2. The conjunction lemma says that, if π and $\overline{\pi}$ are, respectively, subgame perfect Nash equilibrium play paths of $G(T)$ and $G(\overline{T})$, then the conjunction $(\pi, \overline{\pi})$ is a subgame perfect Nash equilibrium play path of $G(T + \overline{T})$.

Second step. I prove by induction on the time horizon that on every pure strategy subgame perfect Nash equilibrium play path of a finite repetition of the stage-game G, only action profiles in $\widetilde{A}$ are played. It follows that the set of pure strategy subgame

[4] The choice of the euclidean distance is without loss of generality as all distances derived from norms are equivalent in finite dimension.

perfect Nash equilibrium payoff vectors of the finitely repeated game is included in the set of recursively feasible payoff vectors, see Lemma 6 and Corollary 1.

Third step. I show that for all $T > 0$, any pure strategy subgame perfect Nash equilibrium payoff vector of the finitely repeated game $G(T)$ dominates the recursive effective minimax payoff vector. This means that in any pure strategy subgame perfect Nash equilibrium of the finitely repeated game $G(T)$, each player receives at least her recursive effective minimax payoff, see Lemma 7.

Fourth step. Given $t > 0$ and a recursively feasible payoff vector y that dominates the recursive effective minimax payoff vector, I construct a subgame perfect Nash equilibrium payoff vector y^t of the finitely repeated game $G(t)$ such that the sequence $(y^t)_{t\geq 1}$ converges to y. The family of equilibrium strategies that I use to sustain a target play path is similar to those used by Smith (1995), Fudenberg and Maskin (1986), Abreu et al. (1994) and Gossner (1995). The challenge here is to independently motivate each player of the block N_h to be an effective punisher during a punishment phase. Indeed, as some players of the block N_h might have equivalent utility functions, the payoff asymmetry lemma of Abreu et al. (1994) does not generate a suitable reward payoff family. To overcome this difficulty, I make use of a more powerful lemma, Lemma 9, which guarantees the existence of a multi-level reward path function. The following five phases briefly describe the above later family of strategy profiles.

The first phase (Phase $\mathbf{P}_0$) of the considered strategy consists to repeatedly follow a target play path π^y that has an average payoff equal to y. The second phase [Phase $\mathbf{P}(i)$] is a punishment phase and prescribes a way to punish a player, say i, if she belongs to the block N_h and is the only one who deviated from the first phase. During this phase, each player of the block $N_h \backslash \mathcal{J}(i)$ can play whatever pure action she wants while players of the block $\mathcal{J}(i) \cup (N \backslash N_h)$ are required to play according to a profile $\widetilde{m}^i$.[5] The third phase serves as a compensation for players of the equivalence class $\mathcal{J}(i)$. Indeed, those players might receive strictly less than their recursive effective minimax payoff in each period of the phase $\mathbf{P}(i)$. The fourth phase is a transition. During the fifth phase, players of the block N_h are rewarded. The reward level of each player depends on whether she was effective punisher during the last punishment phase or not. It turns out that an utility maximizing player will find it strictly dominant to be an effective punisher during the phase $\mathbf{P}(i)$.

[5] At the profile of actions $\widetilde{m}^i$, player i does not have to be at a pure best response. If she plays a pure best response to $\widetilde{m}^i$, she receives at least her stage-game pure minimax payoff but no more than her stage-game recursive effective minimax payoff.

4 Discussion and extension

4.1 Case of the Nash solution

Theorem 1 provides a complete characterization of the limit set of the set of pure strategy subgame perfect Nash equilibrium payoff vectors of the finitely repeated game. In this section, I provide similar result for the set of pure strategy Nash equilibrium payoff vectors of the finitely repeated game.

I find convenient to introduce few notations.

Let $G = (N, A = \times_{i \in N} A_i, u = (u_i)_{i \in N})$ be a compact normal form game. For all player i, let $\mu_i = min_{a \in A} max_{a_i \in A_i} u_i(a_i, a_{-i})$ be the minimax payoff of player i and $\mu = (\mu_1, ..., \mu_n)$ be the minimax payoff vector of the game G.

Let $\tau^*(G) = (N, A, (u_i^*)_{i \in N})$ be the normal form game where the utility function u_i^* of player $i \in N$ is the same as in the original game G, unless the original game G has a pure Nash equilibrium in which player i has a payoff that is strictly greater than her minimax payoff μ_i. In that case, her utility function u_i^* equals the constant γ.

Let $G^{*0} := G$ and $G^{*l+1} := \tau^*(G^{*l})$ for all $l \geq 0$. For all $l \geq 0$, let N_l^* be the set of players with a utility function that is constant to γ in the game G^{*l}. As N is finite, there is an $h \in [0, +\infty)$ such that $N_{l+1}^* = N_l^*$ for all $l \geq h$. Let $A^* = \text{Nash}(G^{*h})$ be the set of pure Nash equilibria of the game G^{*h}.

Definition 3 *The set of* ***Nash-feasible payoff vectors*** *of the game G is defined as the convex hull $Conv[u(A^*)]$ of the set $u(A^*) = \{u(a) \mid a \in A^*\}$.*

Recall that a payoff vector is called individually rational if it dominates the minimax payoff vector of the stage-game.

Theorem 2 *Let G be a normal form stage-game with a finite number of players and a compact set of action profiles. As the time horizon increases, the set of pure strategy Nash equilibrium payoff vectors of the finitely repeated game converges (in the Hausdorff sense) to the set of Nash-feasible and individually rational payoff vectors.*

The proof of Theorem 2 is provided in Appendix 2.

4.2 Alternative statement of Theorem 1 and Theorem 2

Theorem 1 and Theorem 2 respectively provide the limit set of the set of pure strategy subgame perfect Nash equilibrium payoff vectors of any finitely repeated game and the limit set of the set of pure strategy Nash equilibrium payoff vectors of any finitely repeated game. Theorem 1 and Theorem 2 can equivalently be stated as necessary and sufficient conditions on a feasible payoff vector of any given stage-game to be approachable by equilibrium strategies of finite repetitions of that stage-game.

Recall that a payoff vector is called feasible if it belongs to the convex hull of the set of stage-game payoff vectors $u(A) = \{u(a) \mid a \in A\}$.

Definition 4 *A feasible payoff vector x is approachable by means of pure strategy subgame perfect Nash equilibria of the finitely repeated game if for all $\varepsilon > 0$ there exists an integer T_ε such that for all $T > T_\varepsilon$, the finitely repeated game $G(T)$ has a pure strategy subgame perfect Nash equilibrium whose average payoff vector is within ε of x.*

Definition 5 *A feasible payoff vector x is approachable by means of pure strategy Nash equilibria of the finitely repeated game if for all $\varepsilon > 0$ there exists an integer T_ε such that for all $T > T_\varepsilon$, the finitely repeated game $G(T)$ has a pure strategy Nash equilibrium whose average payoff vector is within ε of x.*

Theorem 3 *Let G be a normal form stage-game with a finite number of players and a compact set of action profiles. Let x be a feasible payoff vector. The following statements are equivalent.*

1 *The payoff vector x is recursively feasible and recursively individually rational.*

2 *The payoff vector x is approachable by means of pure strategy subgame perfect Nash equilibria of the finitely repeated game.*

Theorem 4 *Let G be a normal form stage-game with a finite number of players and a compact set of action profiles. Let x be a feasible payoff vector. The following statements are equivalent.*

1 *The payoff vector x is Nash-feasible and individually rational.*

2 *The payoff vector x is approachable by means of pure strategy Nash equilibria of the finitely repeated game.*

The equivalence of Theorem 1 (respectively Theorem 2) and Theorem 3 (respectively Theorem 4) follow from Lemma 5 (respectively Lemma 13).

4.3 Case with discounting

Theorem 1 and Theorem 2 assume no discounting. This assumption is without loss of generality. The underlying reason is that a payoff continuation lemma for finitely repeated game with discounting holds. This lemma allows to approach any feasible payoff vector by means of deterministic paths in the case that there exists a discount factor. I show in the Appendix 3 how to make use this payoff continuation lemma to prove the effective folk theorem for finitely repeated games with discounting.

Lemma 1 (Payoff continuation lemma for finitely repeated game) *For any $\varepsilon > 0$, there exists $k > 0$ and $\underline{\delta} < 1$ such that for any feasible payoff vector x, there exists a deterministic sequence of profile of stage-game actions $\{a^\tau\}_{\tau=1}^{k}$ whose discounted average payoff is within ε of x for all discount factor $\delta \geq \underline{\delta}$.*

This lemma establishes that for any positive ε, there exists an uniform $k > 0$ and $\underline{\delta}$ such that any feasible payoff is within ε of the discounted average of a deterministic path of length k for any discount factor greater than or equal to $\underline{\delta}$.

4.4 Relation with the literature

Finitely repeated games with complete information and perfect monitoring has extensively been studied. This paper provides a generalization of earlier results by Benoit and Krishna (1984), Benoit and Krishna (1987), Smith (1995) and González-Díaz (2006).

The sequence of subset $(N_l)_{l\geq 0}$ defined in Section 2.1 induces a Nash decomposition $0 \subsetneq N_1 \subsetneq \cdots \subsetneq N_h$. The Nash decomposition is called complete if $N_h = N$. Smith (1995) proved that having a complete Nash decomposition is a necessary and sufficient condition for the limit perfect folk theorem to hold. Under a complete Nash decomposition, the set of recursively feasible payoff vectors equals the classic set of feasible payoff vectors and the recursive effective minimax payoff vector equals the classic effective minimax payoff vector. In that case, Theorem 3 says that any feasible payoff vector that dominates the effective minimax payoff vector is approachable by means of pure strategy subgame perfect Nash equilibria of the finitely repeated game. That is the message of the limit perfect folk theorem.

Benoit and Krishna (1984) showed that, if the dimension of the set of feasible payoff vectors of the stage-game equals the number of players and each player receives at least two distinct payoffs at pure Nash equilibria of the stage-game, then the limit perfect folk

theorem holds. This result is a particular case of Theorem 3. Indeed, under the distinct stage-game Nash equilibrium payoffs condition of Benoit and Krishna (1984), the Nash decomposition of the stage-game equals $\varnothing \subsetneq N_h = N$ which is complete and therefore the set of the recursively feasible payoff vectors equals the classic set of the feasible payoff vectors and the recursive effective minimax payoff vector equals the classic effective minimax payoff vector. Furthermore, under the full dimensionality condition, the effective minimax payoff vector equals the minimax payoff vector.

Benoit and Krishna (1987) provided a sufficient condition under which any feasible and individually rational payoff vector can be approximated by the average payoff in a Nash equilibrium of the finitely repeated game. The authors showed that it is sufficient that any player receives in at least one stage-game Nash equilibrium a payoff that is strictly greater than her minimax payoff vector. Basically, under this condition, the decomposition $\varnothing \subsetneq N_1^* = N$ is complete and the set of Nash-feasible payoff vectors equals the set of feasible payoff vector. In such a case, Theorem 4 says that any feasible and individually rational payoff vector of the stage-game can be approached by means of pure strategy Nash equilibria of the finitely repeated game.

González-Díaz (2006) studied the set of Nash equilibrium payoff vectors of a finitely repeated game. His analysis however, differs from that of Section 4.1 of this paper . Indeed, González-Díaz (2006) restricted attention to a particular set of payoff vectors –the set of payoff vectors that belong to the convex hull of the set of payoff vectors to profile of pure actions of the stage-game that dominate the pure minimax payoff vector of the stage-game–. This restriction is not without loss of generality, since the set of Nash equilibrium payoff vectors of the finitely repeated game might converge to a higher-dimension upper set. Theorem 2 and Theorem 4 of this paper provide a full characterization of the whole limit set of the set of pure strategy Nash equilibrium payoffs of the finitely repeated game.

5 Conclusion

This paper analyzed the set of pure strategy subgame perfect Nash equilibrium payoff vectors of the finitely repeated games with complete information. The main finding is an effective folk theorem. It is a complete characterization of the limit set, as the time horizon increases, of the set of pure strategy subgame perfect Nash equilibrium payoff vectors of the finitely repeated game. As the time horizon increases, the limiting set always exists, is closed, convex and can be strictly in between the convex hull of

the set of stage-game Nash equilibrium payoff vectors and the classic set of feasible and individually rational payoff vectors. Our finding exhibits the exact range of cooperative payoffs that players can achieve in finite time horizon. One might wonder if similar results holds in the case that players can employ unobservable mixed strategies or in the case that equilibrium strategies are are protected against renegotiation.

6 Appendix 1: Proof of the Complete perfect folk theorem

6.1 On the existence of the limit set of the set of pure strategy subgame perfect Nash equilibrium payoff vectors of the finitely repeated game

In this section, I show that the limit set of the set of pure strategy subgame perfect Nash equilibrium payoff vectors of any finitely repeated game is well defined. Precisely, I prove that for any stage-game, the set of feasible payoff vectors that are approachable by means of pure strategy subgame perfect Nash equilibria of the finitely repeated game equals the limit set E. As corollary, I obtain that the limit set E is a compact and convex subset of the set of feasible payoff vectors of the stage-game. The main ingredient of this proof is the conjunction lemma established by Benoit and Krishna (1984) . The conjunction lemma says that the conjunction of two subgame perfect Nash equilibrium play paths is a subgame perfect Nash equilibrium play path of the corresponding finitely repeated game. I state it below. Note that the convexity and the compactness of E considerably simplify the proof of Theorems 1 and 3.

Lemma 2 *(See Benoit and Krishna (1984)) If π and $\overline{\pi}$ are two subgame perfect Nash equilibrium play paths of $G(T)$ and $G(\overline{T})$ respectively, then the conjunction $(\pi, \overline{\pi})$ is a subgame perfect Nash equilibrium play path of $G(T + \overline{T})$.*

Let G be a compact normal form game and let ASPNE(G) be the set of all feasible payoff vectors of the stage-game G that are approachable by means of pure strategy subgame perfect Nash equilibrium payoff vectors of the finitely repeated game (see Definition 4).

Lemma 3 *The set* ASPNE(G) *is compact and convex.*

Proof of Lemma 3.
The reader can check that ASPNE(G) is a closed subset of the set of feasible payoff vectors which is compact. The set ASPNE(G) is therefore compact. Since ASPNE(G) is closed, its convexity holds if $z = \frac{1}{2}(x+y) \in$ ASPNE(G) for all $x, y \in$ ASPNE(G). Let $x, y \in$ ASPNE(G) and let $\varepsilon > 0$. Choose T_0^x and T_0^y from the Definition 4 such that for all $T > \max\{T_0^x, T_0^y\}$, the finitely repeated game $G(T)$ has two pure strategy subgame perfect Nash equilibria σ^x and σ^y such that $d(x, u^T(\sigma^x)) < \frac{\varepsilon}{5}$ and $d(y, u^T(\sigma^y)) < \frac{\varepsilon}{5}$. Let $T > \max\{T_0^x, T_0^y\}$, σ^x and σ^y be two pure strategy subgame perfect Nash equilibria of the game $G(T)$ such that $d(x, u^T(\sigma^x)) < \frac{\varepsilon}{5}$ and $d(y, u^T(\sigma^y)) < \frac{\varepsilon}{5}$. Let $\pi = (\pi(\sigma^x), \pi(\sigma^y))$ be the conjunction of the subgame perfect Nash equilibrium play paths $\pi(\sigma^x)$ and $\pi(\sigma^y)$ generated by the strategies σ^x and σ^y respectively. Let $a \in$ Nash(G) be a pure Nash equilibrium of the stage-game G and $\pi' = (a, \pi(\sigma^x), \pi(\sigma^y))$ be the conjunction of the pure Nash equilibrium a and the play path π. From Lemma 2, π and π' are respectively subgame perfect Nash equilibrium play paths of $G(2T)$ and $G(2t+1)$. In addition, $d(z, u^{2T}(\pi)) < \frac{4\varepsilon}{5}$ and

$$d(z, u^{2T+1}(\pi')) < d(z, u^{2T}(\pi)) + d(u^{2T}(\pi), u^{2T+1}(\pi')) < \frac{4\varepsilon}{5} + \frac{2\rho}{2T+1}$$

where $\rho = 2\max_{a\in A} \|u(a)\|_\infty$. Consequently, for all $T > 2\max\{T_0^x, T_0^y, \frac{10\rho}{\varepsilon}\}$, the finitely repeated game $G(T)$ has a pure strategy subgame perfect Nash equilibrium whose average payoff is within ε of z. That is $z \in$ ASPNE(G). ■

Lemma 4 *For all $T > 0$, $E(T) \subseteq$ ASPNE(G).*

Proof of Lemma 4.
Let σ be a pure strategy subgame perfect Nash equilibrium of the finitely repeated game $G(T)$ and $\pi(\sigma) = (\pi_1(\sigma), \cdots, \pi_T(\sigma))$ be the play path generated by σ. Let $x = u^T(\sigma)$. For all $s \geq 0$ and $t \in \{2, \cdots, T\}$, let

$$\pi(s,t) = (\pi_t(\sigma), \cdots, \pi_T(\sigma), \underbrace{\pi(\sigma), \cdots, \pi(\sigma)}_{s \text{ times}})$$

be a play path of $G((s+1)T - t + 1)$. From Lemma 2, $\pi(s,l)$ is a pure strategy subgame perfect Nash equilibrium play path of the finitely repeated game $G((s+1)T - t + 1)$. Moreover, the sequence of payoff vectors $\left(u^{(s+1)T-t+1}[\pi(s,l)]\right)_{s\geq 0}$ converges to x. ■

Lemma 5 *As the time horizon increases, the set of pure strategy subgame perfect Nash equilibrium payoff vectors of the finitely repeated game converges to the set* ASPNE(G).[6]

[6]The convergence in this lemma uses the Hausdorff distance. See Section 2.2.

Proof of Lemma 5. Let $\varepsilon > 0$. We search for $T_\varepsilon > 0$ such that for all $T > T_\varepsilon$, $d_H(\text{ASPNE(G)}, E(T)) < \varepsilon$. Let $\{B(x^l, \frac{\varepsilon}{2}) \mid x^l \in P, l = 1, ..., L\}$ be a finite covering of ASPNE(G).[7] For all $l = 1, ..., L$ take T_0^l given by the definition of "$x^l \in$ ASPNE(G)" with $\frac{\varepsilon}{2}$.[8] Pose $T_0 = \max_{l \leq L} T_0^l$. Let $T > T_0$ and let $x \in$ ASPNE(G). Let $x^{l_0} \in$ ASPNE(G) be such that $x \in B(x^{l_0}, \frac{\varepsilon}{2})$ and let $y \in E(T)$ be such that $d(x^{l_0}, y) < \frac{\varepsilon}{2}$. We have $d(x, y) \leq d(x, x^{l_0}) + d(x^{l_0}, y) < \varepsilon$. This implies that $d(x, E(T)) < \varepsilon$. Consequently, $\sup_{x \in \text{ASNPE}(G)} d(x, E(T)) \leq \varepsilon$. Furthermore, from Lemma 4, $d(y, \text{ASPNE}(G)) = 0$ for all $y \in E(T)$. That is $\sup_{y \in E(T)} d(y, \text{ASPNE}(G)) = 0$. It follows that $d_H(\text{ASPNE}(G), E(T))$ $= \sup_{x \in P} d(x, E(T)) \leq \varepsilon$ for all $T > T_0$. Take $T_\varepsilon = T_0$. ■

6.2 The recursive feasibility of pure strategy subgame perfect Nash equilibrium payoff vectors of the finitely repeated game

Lemma 6 *Let G be a compact normal form game, let $T > 0$, and let σ be a pure strategy subgame perfect Nash equilibrium of $G(T)$. The support* $\text{Supp}(\pi(\sigma)) = \{\pi_1(\sigma) \dots \pi_T(\sigma)\}$ *of the subgame perfect Nash equilibrium play path $\pi(\sigma) = (\pi_1(\sigma) \dots \pi_T(\sigma))$ is included in the set* $\text{Nash}(G^h)$ *of pure Nash equilibrium profiles of the effective game G^h.*

Proof of Lemma 6.
If $N_h = N$, then $\text{Nash}(G^h) = A$ and $\text{Supp}(\pi(\sigma)) \subseteq \text{Nash}(G^h)$. Now assume that $N \backslash N_h \neq \emptyset$. Let's proceed by induction on the time horizon T.
For $T = 1$, the pure strategy subgame perfect Nash equilibrium σ is a pure Nash equilibrium of the stage-game G. By construction, the sequence $(\text{Nash}(G^l))_{l \geq 0}$ is increasing and therefore $\text{Nash}(G) = \text{Nash}(G^0) \subseteq \text{Nash}(G^h)$.
Suppose that $T > 1$ and that the support of any subgame perfect Nash equilibrium play path of the finitely repeated game $G(t)$ with $t \in \{1, \dots, T-1\}$ is included in the set $\text{Nash}(G^h)$ and let's show that $\{\pi_1(\sigma), \dots, \pi_T(\sigma)\} \subseteq \text{Nash}(G^h)$. The restriction $\sigma_{|\pi_1(\sigma)}$ of σ to the history $\pi_1(\sigma)$ is a pure strategy subgame perfect Nash equilibrium of the game $G(T-1)$ and the induction hypothesis implies that the support $\{\pi_2(\sigma) \dots \pi_T(\sigma)\}$ of the play path $\pi(\sigma_{|\pi_1(\sigma)})$ generated by the strategy profile $\sigma_{|\pi_1(\sigma)}$ is included in $\text{Nash}(G^h)$. It remains to show that $\pi_1(\sigma) \in \text{Nash}(G^h)$.
At this point I proceed by contradiction. Assume that $\pi_1(\sigma) \notin \text{Nash}(G^h)$. Then, in the game G^h, there exists a player $i \in N$ who has a strict incentive to deviate from the pure action profile $\pi_1(\sigma)$. This player has to be in the block $N \backslash N_h$ since any player of the block N_h has a constant utility function in the game G^h. Let σ_i' be a pure strategy

[7] $B(x, \varepsilon) = \{y \in \mathbb{R}^n \ / \ d(x, y) < \varepsilon\}$
[8] See Definition 4.

one shot deviation of player i from σ that consists in playing a stage-game pure best response $b_i[\pi_1(\sigma)]$ to $\pi_1(\sigma)$ in the first round of the finitely repeated game $G(T)$ and conforming to σ_i from the second round on. At the pure strategy profile (σ_i', σ_{-i}), player i receives $u_i(\pi^1) + e$ (with $e > 0$) in the first round. Let $h^1 = (b_i(\pi_1(\sigma)), \pi_1(\sigma)_{-i})$ be the observed history after this first round and $\sigma_{|h^1}$ be the restriction of σ to the history h^1. We have $(\sigma_i', \sigma_{-i})_{|h^1} = \sigma_{|h^1}$ and $\sigma_{|h^1}$ is a pure strategy subgame perfect Nash equilibrium of $G(T-1)$. By induction hypothesis, the support of the play path generated by $\sigma_{|h^1}$ is included in Nash(G^h). Therefore, at the profile (σ_i', σ_{-i}) player i receives the sequence of stage-game payoffs $\{u_i(\pi^1) + e, n_i, \ldots, n_i\}$ where n_i is her unique stage-game pure Nash equilibrium payoff.[9] Since player i receives $\{u_i(\pi_1(\sigma)), n_i, ...n_i\}$ at the strategy profile σ, we have $u_i^T(\sigma_i', \sigma_{-i}) > u_i^T(\sigma)$. This contradicts the fact that σ is a pure strategy subgame perfect Nash equilibrium of $G(T)$ and concludes the proof. ■

Let $\widetilde{F}$ be the set of recursively feasible payoff vectors. We have the following corollary.

Corollary 1 *Let G be a compact normal form game, let $T > 0$, and let σ be a pure strategy subgame perfect Nash equilibrium of $G(T)$. Then the average payoff vector $u^T(\sigma)$ belongs to the set $\widetilde{F}$.*

6.3 Necessity of the recursive effective minimax payoff for the complete perfect folk theorem

Wen (1994) shows that any subgame perfect Nash equilibrium payoff vector of the infinitely repeated game weakly dominates the effective minimax payoff vector. This domination also holds for finitely repeated games. The following lemma provides a sharp upper bound. The lemma says that, any pure strategy subgame perfect Nash equilibrium payoff vector of the finitely repeated game weakly dominates the recursive effective minimax payoff vector.

Lemma 7 *Let G be a compact normal form game, let $T \geq 1$, and let σ be a pure strategy subgame perfect Nash equilibrium of the finitely repeated game $G(T)$. Then the average payoff vector $u^T(\sigma)$ dominates the recursive effective minimax payoff vector of the stage-game.*

I find convenient to recall the definition of the recursive effective minimax payoff before proceeding to the proof of Lemma 7.

[9] Recall that each player of the block $N \backslash N_h$ has a unique pure Nash equilibrium payoff in the game G^h. This payoff equals her unique pure Nash equilibrium payoff in the original game G.

Let $\sim$ be the equivalence relation defined on the set of players as follows: Player i is equivalent to j (denoted by $i \sim j$) if there exists $\alpha_{ij} > 0$ and $\beta_{ij} \in \mathbb{R}$ such that for all $a \in \widetilde{A}$, we have $u_i(a) = \alpha_{ij} \cdot u_j(a) + \beta_{ij}$. For all $i \in N$, let $\mathcal{J}(i)$ be the equivalence class of player i and let

$$\widetilde{\mu}_i = \min_{a \in \widetilde{A}} \max_{j \in \mathcal{J}(i)} \max_{a'_j \in A_j} \left[\alpha_{ij} \cdot u_j(a'_j, a_{-j}) + \beta_{ij}\right]$$

and $\widetilde{\mu} = (\widetilde{\mu}_1, \ldots, \widetilde{\mu}_n)$. The payoff $\widetilde{\mu}_i$ is the **recursive effective minimax** of player i in the stage-game G and the n-tuple $\widetilde{\mu}$ is the **recursive effective minimax payoff vector** of the stage-game G.

Proof of Lemma 7.

I proceed by induction on the time horizon T.

At $T = 1$, pure strategy subgame perfect Nash equilibria of the game $G(T)$ are pure Nash equilibria of the stage-game G and $u^T(\sigma)$ dominates $\widetilde{\mu}$.[10]

Assume that $T > 1$ and that the average payoff vector to any pure strategy subgame perfect Nash equilibrium of the finitely repeated game $G(t)$ with $0 < t < T$ dominates the recursive effective minimax payoff vector $\widetilde{\mu}$. Let us show that the payoff vector $u^T(\sigma)$ dominates $\widetilde{\mu}$.

Let $\pi_1(\sigma)$ be the action profile played in the first round of the game $G(T)$ according to σ. The restriction $\sigma_{|\pi_1(\sigma)}$ of the strategy σ to the history $\pi_1(\sigma)$ is a pure strategy subgame perfect Nash equilibrium of the finitely repeated game $G(T-1)$ and by induction hypothesis, we have that the payoff vector $u^{T-1}(\sigma_{|\pi_1(\sigma)})$ dominates $\widetilde{\mu}$. Suppose now that $u^T(\sigma)$ does not dominates $\widetilde{\mu}$. Then there exists a player $i \in N$ such that $u_i^T(\sigma) < \widetilde{\mu}_i$. It follows that $u_i[\pi_1(\sigma)] < \widetilde{\mu}_i$ since $u_i^T(\sigma)$ is a convex combination of $u_i[\pi_1(\sigma)]$ and $u_i^{T-1}(\sigma_{|\pi_1(\sigma)})$. Moreover, as $\pi_1(\sigma) \in \text{Nash}(G^h)$, we have $u_j[\pi_1(\sigma)] < \widetilde{\mu}_j$ for all $j \in \mathcal{J}(i)$. From the definition of $\widetilde{\mu}$, there exists a player $i_0 \in \mathcal{J}(i)$ and a pure action $a_{i_0} \in A_{i_0}$ of player i_0 such that $u_{i_0}[a_{i_0}, \pi_1(\sigma)_{-i_0}] \geq \widetilde{\mu}_{i_0}$. Consider the pure strategy one shot deviation σ'_{i_0} of player i_0 from σ in which she plays a_{i_0} in the first round of the finitely repeated game $G(T)$ and conforms to her strategy σ_{i_0} from the second round on. We have

$$u_{i_0}^T(\sigma'_{i_0}, \sigma_{-i_0}) = \tfrac{1}{T} u_{i_0}[a_{i_0}, \pi_1(\sigma)_{-i_0}] + \tfrac{T-1}{T} u_{i_0}^{T-1}(\sigma_{|(a_{i_0}, \pi_1(\sigma)_{-i_0})})$$

which is greater than or equal to $\widetilde{\mu}_{i_0}$. Indeed, since $\sigma_{|(a_{i_0}, \pi_1(\sigma)_{-i_0})}$ is a pure strategy subgame perfect Nash equilibrium play path of the finitely repeated game $G(T-1)$, the induction hypothesis implies that $u(\sigma_{|(a_{i_0}, \pi_1(\sigma)_{-i_0})})$ dominates $\widetilde{\mu}$. ■

[10] Indeed, as each pure Nash equilibrium of the stage-game G is a pure Nash equilibrium of the game G^h and each player plays a best response in Nash equilibrium, the Nash equilibrium payoff of any player is greater than or equal to her recursive effective minimax payoff. It follows that any pure Nash equilibrium payoff vector weakly dominates the recursive effective minimax payoff vector.

6.4 Sufficiency of the recursive feasibility and the recursive effective individual rationality

From Corollary 1 and Lemma 7, the set of pure strategy subgame perfect Nash equilibrium payoff vectors of any finite repetition of the stage-game G is included in the set of recursively feasible and recursively individually rational payoff vectors. To complete the proofs of Theorem 1, it is left to show that any recursively feasible and recursively individually rational payoff vector belongs to the limit set E. In what follows, I prove that any recursively feasible and recursively individually rational payoff vector is approachable by means of pure strategy subgame perfect Nash equilibria of the finitely repeated game. This will conclude the proof of Theorem 1 as well as the proof of Theorem 3, see Lemma 5. I proceed with 3 lemmata. The message of the first lemma is that in the finitely repeated game, players of the block N_h receive distinct payoffs at pure strategy subgame perfect Nash equilibria.

The sequence of subsets $(N_l)_{l\geq 0}$ defined in Section 2.1 induces a separation of the set of players into two blocks N_h and $N\backslash N_h$. As a corollary of Lemma 6, each player of the block $N\backslash N_h$ (if any) receives her unique stage-game pure Nash equilibrium payoff at each round of a pure strategy subgame perfect Nash equilibrium of any finite repetition of the stage-game G. The underlying reason is that there is no way to credibly leverage the behavior of any player of the latter block near the end of the game. The next lemma says that each player of the block N_h receives distinct payoffs at pure strategy subgame perfect Nash equilibria of the finitely repeated game. The construction of this lemma is inspired by Smith (1995).

Let G be a compact normal form game that has at least two distinct pure Nash equilibrium payoff vectors. Let

$$\emptyset = N_0 \subsetneq N_1 \subsetneq ... \subsetneq N_h$$

be the Nash decomposition of G.

Lemma 8 *There exists T_0 such that for all $T \geq T_0$, each player of N_h receives at least two distinct payoffs at pure strategy subgame perfect Nash equilibria of the finitely repeated game $G(T)$.*

Proof of Lemma 8.
I prove that for all $g \leq h$, there exists $T_{0,g}$ such that for all $T \geq T_{0,g}$, each player of the block N_g receives distinct payoffs at pure strategy subgame perfect Nash equilibria of $G(T)$. Obviously this property holds for $g = 1$ since each player of the block N_1 receives

distinct payoffs at pure Nash equilibria of the stage-game G. Let $g \geq 1$ and assume that the property holds for g. For all $j \in N_g$, let $\pi^{j,g}$ and $\overline{\pi}^{j,g}$ be respectively the best and the worst pure strategy subgame perfect Nash equilibrium play path of player j in the game $G(T_{0,g})$. Let $\rho = 2\max_{a\in A} \|u(a)\|_\infty$ and $\psi > 0$ such that

$$-\rho + \psi \cdot T_{0,g} \cdot \sum_{j\in N_g} u_i^T(\pi^{j,g}) > \psi\,|N_g| \cdot T_{0,g} \cdot u_i^T(\overline{\pi}^{i,g})$$

for all $i \in N_g$. Each player $j \in N_g$ is willing to conform to any pure action profile followed by ψ cycles $(\pi^{i,g})_{i\in N_g}$ if deviations by player j are punished by switching each $\pi^{i,g}$ to $\overline{\pi}^{j,g}$. Let $i_0 \in N_{g+1}\backslash N_g$ and let $y^{i_0,g}$ and $z^{i_0,g}$ the best and respectively the worst pure strategy Nash equilibrium of player i_0 in the one shot game G^g. Player i_0 receives distinct payoffs at pure strategy subgame perfect Nash equilibrium play paths

$$\pi^{i_0} = \left(y^{i_0,g}, \underbrace{(\pi^{i,g})_{i\in N_g}, \cdots, (\pi^{i,g})_{i\in N_g}}_{\psi\text{times}} \right)$$

and

$$\overline{\pi}^{i_0} = \left(z^{i_0,g}, \underbrace{(\pi^{i,g})_{i\in N_g}, \cdots, (\pi^{i,g})_{i\in N_g}}_{\psi\text{times}} \right).$$

This guarantee the existence of $T_{0,g+1}$ such that each player of the block $N_{g+1}\backslash N_g$ receives distinct payoffs at pure strategy subgame perfect Nash equilibria of $G(T_{0,g+1})$. Repeatedly appending the same stage-game pure Nash equilibrium profile at each π^{i_0} and $\overline{\pi}^{i_0}$, we obtain for each $T \geq T_{0,g+1}$ and $i_0 \in N_{g+1}\backslash N_g$ two pure strategy subgame perfect Nash equilibrium play paths of $G(T)$ at which player i_0 receives distinct payoffs. This concludes the proof of the lemma. ■

The next lemma establishes the existence of a multi-level reward path function. In the case that the full dimensionality condition of Fudenberg and Maskin (1986) or the non-equivalent utility (NEU) condition of Abreu et al. (1994) does not hold, a multi-level reward path function can still be used to independently control the incentives of players of the block N_h and motivate them to be effective punishers during a punishment phase. This lemma also allows to leverage the behavior of players of the block N_h near the end of the game.

Lemma 9 *Let $\emptyset = N_0 \subsetneq N_1 \subsetneq ... \subsetneq N_h$ be the Nash decomposition of the game G. Then there exists $\phi > 0$ such that for all $p \geq 0$ there exists $r_p > 0$ and*

$$\theta^p : \{0,1\}^n \cup \{(-1,\cdots,-1)\} \to A^{r_p} := A \times \cdots \times A$$

such that for all $\alpha \in \{0,1\}^n \cup \{(-1,\cdots,-1)\}$, $\theta^p(\alpha)$ *is a play path generated by a pure strategy subgame perfect Nash equilibrium of the repeated game* $G(r_p)$. *Furthermore, for all* $i \in N_h$ *and* $\alpha, \alpha' \in \{0,1\}^n$, *we have*

$$u_i^{r_p}[\theta^p(1,\alpha_{-i})] - u_i^{r_p}[\theta^p(0,\alpha_{-i})] \geq \phi, \tag{1}$$

$$u_i^{r_p}[\theta^p(\alpha)] - u_i^{r_p}[\theta^p(-1,\cdots,-1)] \geq \phi \tag{2}$$

and

$$|u_i^{r_p}[\theta^p(\alpha)] - u_i^{r_p}[\theta^p(\alpha_{\mathcal{J}(i)}, \alpha'_{N\setminus\mathcal{J}(i)})]| < \frac{1}{2^p}. \tag{3}$$

Proof of Lemma 9. The set ASPNE(G) of feasible payoff vectors that are approachable by means of pure strategy subgame perfect Nash equilibria of finite repetitions of the stage-game G is non-empty and convex and therefore has a relative interior point x, see Lemma 3. Let $\phi > 0$ such that the relative ball $\widetilde{B}(x, 5\phi n)$ is included in ASPNE(G).[11] For all $\alpha \in \{-1,0,1\}^n$ and $j \in N_h$, let

$$\theta_j(\alpha) = x_j - \phi|\mathcal{J}(j)| + 3\phi \textstyle\sum_{j' \in \mathcal{J}(j)} \alpha'_j.$$

For all $j \notin N_h$, let

$$\theta_j(\alpha) = x_j.$$

I recall that if $j \notin N_h$, then x_j is the unique stage-game pure Nash equilibrium payoff of player j. For all $\alpha \in \{-1,0,1\}^n$, let

$$\theta(\alpha) = (\theta_1(\alpha), \cdots, \theta_n(\alpha)).$$

For all $\alpha \in \{0,1\}^n$ and $i \in N_h$ we have

$$\theta_i(1,\alpha_{-i}) - \theta_i(0,\alpha_{-i}) = 3\phi;$$

$$\theta_i(\alpha) - \theta_i(-1,\cdots,-1) \geq 3\phi$$

and

$$\|\theta(\alpha) - x\| < 5n\phi.$$

Furthermore, since players of the block N_h receive distinct payoffs at pure strategy subgame perfect Nash equilibria of the finitely repeated game (see Lemma 8), each of them also receives distinct payoffs within the set ASPNE(G) (see Lemma 4). It follows that

[11] For simplicity and as ASPNE(G) is convex, one can take $\widetilde{B}(y, 5\phi n) = \{x \in \text{ASPNE}(G) \mid d(x,y) < 5\phi n\}$.

$$\theta(\alpha) \in \widetilde{B}(x, 5\phi n) \subseteq \text{ASPNE}(G).$$

For all $p \geq 0$, let $\varepsilon_p = \frac{1}{2}\min\{\phi, \frac{1}{2^p}\}$. For all $\alpha \in \{0,1\}^n \cup \{(-1, \cdots, -1)\}$, let $T_{0,\alpha,p} < \infty$ and for all $T \geq T_{0,\alpha,p}$, let $\sigma^{\alpha,p}$ be a pure strategy subgame perfect Nash equilibrium of the repeated game $G(T)$ such that $\|u^T(\sigma^{\alpha,p}) - \theta(\alpha)\| < \varepsilon_p$.
Let $r_p = \max\{T_{0,\alpha,p} \mid \alpha \in \{0,1\}^n \cup \{(-1, \cdots, -1)\}\}$. For all $\alpha \in \{0,1\}^n \cup \{(-1, \cdots, -1)\}$, let $\theta^p(\alpha)$ be the pure strategy subgame perfect Nash equilibrium play path generated by the pure strategy subgame perfect Nash equilibrium $\sigma^{\alpha,p}$ of the repeated game $G(r_p)$. ■

Lemma 10 *Let G be a compact normal form game. We have $\widetilde{F} \cap \widetilde{I} \subseteq$ ASPNE(G).*

Proof of Lemma 10.
Let G be a compact normal form game. If G admits no pure Nash equilibrium, then $\widetilde{F} = \emptyset$ and $\widetilde{F} \cap \widetilde{I} \subseteq$ ASPNE(G). If G admits a unique pure Nash equilibrium payoff vector x, then $\widetilde{F} = \{x\} =$ ASPNE(G) and $\widetilde{F} \cap \widetilde{I} \subseteq$ ASPNE(G). Now suppose that G admits at least two distinct pure Nash equilibrium payoff vectors. Normalize the game such that the recursive effective minimax of each player equals 0 and such that two equivalent players have the same utility function on $\widetilde{A}$. Consider

$$F_1 = \{\tfrac{1}{p}\textstyle\sum_{1 \leq l \leq p} u(a^l) \mid p > 0, a^l \in \widetilde{A}\ \forall l \leq p\}$$

and

$$I_1 = \{x \in \mathbb{R}^n \mid x_i > 0 \text{ if } i \in N_h \text{ and } x_i = 0 \text{ otherwise}\}.$$

It is immediate that the closure of $F_1 \cap I_1$ is equal to the set $\widetilde{F} \cap \widetilde{I}$. From Lemma 3, ASPNE$(G)$ is closed. Therefore, it is enough to show that $F_1 \cap I_1 \subseteq$ ASPNE(G). Let

$$y = \tfrac{1}{k}\textstyle\sum_{1 \leq l \leq k} u(a^l) \in F_1 \cap I_1$$

and

$$\pi^y = (a^1, ..., a^k).$$

For all $i \in N_h$, let

$$\widetilde{m}^i \in arg\min_{a \in \widetilde{A}} \max_{j \in \mathcal{J}(i)} \max_{a'_i \in A_i} u_i(a'_i, a_{-i}).^{12}$$

Obtain $\phi,\ r_1$ and θ^1 with $p = 1$ from the Lemma 9. Let $q_1 > 0$ and $q_2 > 0$ such that

$$0 < q_1 u_i(\widetilde{m}^i) + q_2 r_1 u_i^{r_1}[\theta^1(1, \cdots, 1)] < \frac{q_1 + q_2 r_1}{2} y_i \tag{4}$$

and

$$-2\rho + \frac{q_1}{2} y_i > 0 \ for\ all\ i \in N_h. \tag{5}$$

[12]Few comments on $\widetilde{m}^i$ are provided in footnote 5.

Given q_1, q_2 and r_1, choose r such that

$$-2(q_1 + q_2 r_1)\rho + r\phi > 0. \tag{6}$$

Given q_1 q_2, r_1 and r, choose $p_0 > 0$ such that

$$\frac{q_2 r_1}{2} y_i - \frac{r}{2^{p_0}} > y_i - \frac{r}{2^{p_0}} > 0 \tag{7}$$

Apply the Lemma 9 to p_0 and obtain r_{p_0} and θ^{p_0}. Update $q_1 \leftarrow r_{p_0} q_1; q_2 \leftarrow r_{p_0} q_2 r_1;\ r \leftarrow r_{p_0} r$. The parameters $\phi,\ \theta^1,\ q_1,\ q_2,\ r,\ r_1$ and θ^{p_0} are such that

$$0 < q_1 u_i(\widetilde{m}^i) + q_2 u_i^{r_1}[\theta^1(1, \cdots, 1)] < \frac{q_1 + q_2}{2} y_i \tag{8}$$

$$-2(q_1 + q_2)\rho + r\phi > 0 \tag{9}$$

$$-2\rho + \frac{q_1 + q_2}{2} y_i - \frac{r}{2^{p_0}} > 0 \tag{10}$$

and

$$y_i - \frac{r}{2^{p_0}} > 0 \ for\ all\ i \in N_h. \tag{11}$$

Let

$$\widehat{\pi}^s = (\underbrace{\pi^y, ..., \pi^y}_{s \text{ times}}, \theta^{p_0}(1, \cdots, 1)).$$

Assume that for all $s \geq 0$ there exists σ^s a pure strategy subgame perfect Nash equilibrium of the finitely repeated game $G(sk + r)$ such that the play path $\pi(\sigma^s)$ generated by σ^s equals $\widehat{\pi}^s$. Since the limit of $u^{sk+r}(\widehat{\pi}^s)$ as s goes to infinity equals the payoff vector y and k is finite, there exists $s_\varepsilon > 0$ such that for all $T > s_\varepsilon k + r$, the finitely repeated game $G(T)$ has a pure strategy subgame perfect Nash equilibrium whose average payoff vector is within ε of y. This will conclude the proof of Lemma 10.

Let $s \geq 0$. Let us construct a pure strategy subgame perfect Nash equilibrium σ^s of the finitely repeated game $G(sk + r)$ such that the play path $\pi(\sigma^s)$ generated by σ^s equals $\widehat{\pi}^s$.

In the following, a deviation from a strategy profile of the finitely repeated game $G(sk + r)$ is called "late" if it occurs during the last $q_1 + q_2 + r$ periods of the game $G(sk + r)$. In the other case the deviation is called "early". Set $\alpha = (1, \cdots, 1)$ and consider the pure strategy profile σ^s described by the following 5 phases.

P$_0$ (Main play path): In this phase, players are required to play the $(sk + r - t + 1)$th to last profile of actions of the path $\widehat{\pi}^s$ at time t, $1 \leq t \leq sk + r$.

$$\left[\begin{array}{c}\text{If player } i \in N_h \text{ deviates early, start the Phase } \mathbf{P}(i);\\ \text{if } j \in N_{g'}\backslash N_{g'-1} \text{ deviates late, then start Phase } \mathbf{LD}.\\ \text{Ignore any deviation by a player } i \notin N_h\end{array}\right]$$

P(i) (Punish player i): Reorder the profile of actions in each upcoming cycle of length k of the main play path according to player i's preferences, starting from her best profile.

During this phase, each player of the block $\mathcal{J}(i) \cup (N\backslash N_h)$ is required to play as in the action profile $\widetilde{m}^i$ while players of the block $N_h\backslash\mathcal{J}(i)$ can play whatever pure action they want. This phase last for q_1 periods. [If any player $j \in \mathcal{J}(i)$ deviates early, restart **P**(i) ; if player $j \in \mathcal{J}(i)$ deviates late, start **LD**; Ignore any deviation by a player $i \notin N_h$.]

At the end of this phase and for all $j \in N_h\backslash\mathcal{J}(i)$, set $\alpha_j = 0$ if there is at least one period of the punishment phase **P**(i) where player j played an action different to $\widetilde{m}^i_j$. In the other case, set $\alpha_j = 1$. Go to phase **SPE**.

SPE (Compensation): Follow $\frac{q_2}{r_1}$ times a pure strategy SPNE of the game $G(r_1)$ whose play path is $\theta^1(1, \cdots, 1)$.
Go to Phase $\mathbf{P}_0$

LD (Late deviation): Each player can play whatever action she wants till period sk. At period sk, set Set $\alpha = (-1, \cdots, -1)$. Go to **EG**.

EG (End-game): Follow $\frac{r}{r_{p_0}}$ times a pure strategy subgame perfect Nash equilibrium of the finitely repeated game $G(r_{p_0})$ that supports the equilibrium play path $\theta^{p_0}(\alpha)$.

The strategy profile σ^s is a pure strategy subgame perfect Nash equilibrium of the finitely repeated game $G(sk+r)$. To see this, I show that parameters $\phi,\ \theta^1,\ q_1,\ q_2,\ r,\ r_1$ and θ^{p_0} are chosen in such a way to deter any deviation from the main play path as well as any deviation from the minimax phase.

I first show that a utility maximizing player $j \in N_h\backslash\mathcal{J}(i)$ will find it strictly dominant to be effective punisher during any punishment phase **P**(i).[13] The underlying reason is that for each player $j \in N_h$, the average utility $u_j^{r_{p_0}}(\theta^{p_0}(\alpha_j, \alpha_{-j}))$ is strictly increasing in α_j. Indeed, if player $j \in N_h\backslash\mathcal{J}(i)$ is effective punisher during the Phase **P**(i), she gets at least

1. $-(q_1 + q_2)\rho$ in the phases **P**(i) and **SPE**;

[13] I call player $j \in N_h\backslash\mathcal{J}(i)$ effective punisher during the punishment phase **P**(i) if $\alpha_j = 1$ at the end of the latter phase. After the punishment phase **P**(i), if $\alpha_{-\mathcal{J}(i)} = (1, \ldots, 1)$, then the average payoff of player i during the punishment phase **P**(i) is less than or equal to 0, independently of the value of α_i.

2. some payoff U_j till period sk;

3. $ru_i^{r_{p_0}}[\theta^{p_0}(1,\alpha_{-j})]$ in the last r periods of the repeated game $G(sk+r)$.

That is in total $-(q_1+q_2)\rho + U_j + ru_i^{r_{p_0}}[\theta^{p_0}(1,\alpha_{-j})]$. If she is not effective punisher, she get at most

1. $(q_1+q_2)\rho$ in the phases $\mathbf{P}(i)$ and **SPE**;

2. the same payoff U_j till period sk;

3. $ru_i^{r_{p_0}}[\theta^{p_0}(0,\alpha_{-j})]$ in the last r periods of the repeated game $G(sk+r)$.

That is in total $(q_1+q_2)\rho + U_j + ru_i^{r_{p_0}}[\theta^{p_0}(0,\alpha_{-j})]$ which is less than or equal to $(q_1+q_2)\rho + U_j + ru_i^{r_{p_0}}[\theta^{p_0}(1,\alpha_{-j})] - r\phi$, see inequality (1). Since $-2(q_1+q_2)\rho + r\phi > 0$, we have

$$-(q_1+q_2)\rho + U_j + ru_i^{r_{p_0}}[\theta^{p_0}(1,\alpha_{-j})] > (q_1+q_2)\rho + U_j + ru_i^{r_{p_0}}[\theta^{p_0}(0,\alpha_{-j})]$$

Thus, it is strictly dominant for any player of the block $N_h \backslash \mathcal{J}(i)$ to be effective punisher during the punishment phase $\mathbf{P}(i)$. No player of the block $N \backslash N_h$ will have any incentive to deviate given that players of the block $N_h \backslash \mathcal{J}(i)$ are effective punisher. Indeed, every player of the block $N \backslash N_h$ plays a stage-game pure best response at each profile of actions $a \in \widetilde{A}$.[14].

1) No early deviation from the phase $\mathbf{P}(i)$ is profitable

If after $l_1k + l_2$ (where $l_2 < k$) periods in the Phase $\mathbf{P}(i)$ a player $j \in \mathcal{J}(i)$ deviates unilaterally, the strategy profile σ^s prescribes to start a new punishment phase ***P(i)*** followed by a ***SPE*** phase, to reorder the profiles of the target path, and to go back to the Phase $\mathbf{P}_0$. Such deviation is not profitable. Indeed, if player j deviates early, she receives at most:

1. 0 in the first $l_1k + l_2$ periods of the Phase $\mathbf{P}(i)$.
2. $q_1u_i(\widetilde{m}^i) + q_2u_i^{r_1}[\theta^1(1,\cdots,1)]$ in the new phase ***P(i)*** and the following ***SPE*** phase;
3. some payoff U_i till period sk;
4. the payoff $ru_i^{r_{p_0}}[\theta^{p_0}(\alpha)] + \frac{r}{2^{p_0}}$ in the End-game.

[14] In the finitely repeated game $G(sk+r)$, after any history h, the strategy profile σ^s prescribes to play the stage-game action profile $\sigma^s(h)$ which belongs to $\widetilde{A} = \text{Nash}(G^h)$, see Lemma 6. As every player of the block $N \backslash N_h$ plays a stage-game pure best response in any profile $a \in \widetilde{A}$, no player of the block $N \backslash N_h$ can profitably deviate from the strategy profile σ^s.

If player i does not deviate, she receives at least:

1. $q_1 u_i(\widetilde{m}^i) + q_2 u_i^{r_1}[\theta^1(1, \cdots, 1)]$ in the Phases $\mathbf{P}(i)$ and **SPE**;
2. $l_1 k y_i + l_2 y_i$ till the end of the phase ***SPE***;
3. the same payoff U_i till period sk;
4. the payoff $ru_i^{r_{p_0}}[\theta^{p_0}(\alpha)]$ in the End-game.

As $y_i - \frac{r}{2^{p_0}} > 0$ and $l_1 k y_i + l_2 y_i \geq 1$, no early deviation from the phase $\mathbf{P}(i)$ is profitable.

2) No early deviation during phase $\mathbf{P}_0$ is profitable

If from the phase $\mathbf{P}_0$ a player let's say i deviates early, then the strategy profile σ^s prescribes to start phase $\mathbf{P}(i)$, to update α and to go to the phase **SPE**. Such a deviation is not profitable. Indeed, if player i deviates early from the phase $\mathbf{P}_0$, she receives at most

1. ρ in the deviation period;
2. $q_1 u_i(\widetilde{m}^i) + q_2 u_i^{r_1}[\theta^1(1, \cdots, 1)]$ in the phase $\mathbf{P}(i)$ and the following **SPE** phase;
3. some payoff U_i till the period sk;
4. the payoff $ru_i^{r_{p_0}}[\theta^{p_0}(\alpha_{\mathcal{J}(i)}, 1, \cdots, 1)]$ till the end of the game.

In total $\rho + q_1 u_i(\widetilde{m}^i) + q_2 u_i^{r_1}[\theta^1(1, \cdots, 1)] + U_i + ru_i^{r_{p_0}}[\theta^{p_0}(\alpha_{\mathcal{J}(i)}, 1, \cdots, 1)]$ which is strictly less than $\rho + \frac{q_1+q_2}{2} y_i + U_i + ru_i^{r_{p_0}}[\theta^{p_0}(\alpha_{\mathcal{J}(i)}, 1, \cdots, 1)]$, see inequality (8). If player i does not deviates, she get at least

1. $-\rho$ in that deviation period;
2. Followed by $(q_1 + q_2) y_i$ corresponding to the phases $\mathbf{P}(i)$ and **SPE**;[15]
3. the same payoff U_i till period sk;
4. the payoff $ru_i^{r_{p_0}}[\theta^{p_0}(\alpha)]$ in the phase **EG**.

That is in total $-\rho + (q_1 + q_2) y_i + U_i + ru_i^{r_{p_0}}[\theta^{p_0}(\alpha)]$ which is greater than $-\rho + (q_1 + q_2) y_i + U_i + ru_i^{r_{p_0}}[\theta^{p_0}(\alpha_{\mathcal{J}(i)}, 1, \cdots, 1)] - \frac{r}{2^{p_0}}$, see inequality (3).

Early deviations from the main path are therefore deterred by inequality (10).

[15] Indeed there is no loss of generality to consider that q_1 and q_2 are multiple of k.

3) No late deviation is profitable.

If from an ongoing phase ($\mathbf{P}_0$ or $\mathbf{P}(i)$) a player let's say $j \in N_h$ deviates late, she receives at most

1. $(q_1 + q_2)\rho$ till the beginning of the phase **EG**;
2. $ru_j^{r_{p_0}}[\theta^{p_0}(-1, \cdots, -1)]$ in the phase **EG**.

If player j does not deviates, she receives at least

1. $-(q_1 + q_2)\rho$ till the beginning of the phase **EG**;
2. $ru_i^{r_{p_0}}[\theta^{p_0}(\alpha)]$ in the phase **EG**, where $\alpha \in \{0.1\}^n$.

As $ru_i^{r_{p_0}}[\theta^{p_0}(\alpha)]$ is grater than or equal to $ru_i^{r_{p_0}}[\theta^{p_0}(-1, \cdots, -1)] + r\phi$ (see inequality (2)), and $-2(q_1+q_2)\rho+r\phi > 0$, no late deviation is profitable. This concludes the proof. ■

7 Appendix 2: Proof of the complete Nash folk theorem

7.1 On the existence of the limit set of the set of pure strategy Nash equilibrium payoff vectors of the finitely repeated game

In this section, I show that the limit set of the set of pure strategy Nash equilibrium payoff vectors of the finitely repeated game is well defined. Namely, I show that for any compact stage-game, this limit set equals the set of feasible payoff vectors that are approachable by means of pure strategy Nash equilibria of the finitely repeated game (see Definition 5). I proceed with lemmata. These lemmata as well as their proofs are very similar to those used in Section 6.1.

Let G be a compact normal form game and let ANE(G) be the set of feasible payoff vectors that are approachable by means of pure strategy Nash equilibria of the finitely repeated game. For any $T > 0$, let NE(T) be the set of pure strategy Nash equilibrium payoff vectors of the finitely repeated game $G(T)$. Let NE be the Hausdorff limit of the set of pure strategy Nash equilibrium payoff vectors of the finitely repeated game.

Lemma 11 *The set* ANE(G) *is a compact and convex set.*

Proof of Lemma 11. It is immediate that ANE(G) is a closed subset of the set of feasible payoff vectors of the stage-game G. As the set of feasible payoff vectors is compact, the set ANE(G) is also compact. The convexity of the set ANE(G) follows from the fact that the conjunction of two pure strategy Nash equilibrium play paths remains a pure Nash equilibrium play path. ■

Lemma 12 *For all $T > 0$,* NE(T) $\subseteq$ ANE(G).

Proof of Lemma 12. Let σ be a pure strategy Nash equilibrium of the finitely repeated game $G(T)$ and $\pi(\sigma) = (\pi_1(\sigma), \cdots, \pi_T(\sigma))$ be the play path generated by σ. Let $x = u^T(\sigma)$. For all $s \geq 0$ and $t \in \{2, \ldots, T\}$, the play path

$$\pi(s,t) = (\pi_t(\sigma), \cdots, \pi_T(\sigma), \underbrace{\pi(\sigma), \cdots, \pi(\sigma)}_{s \text{ times}})$$

is a pure strategy Nash equilibrium play path of the finitely repeated game $G((s+1)T - t + 1)$ and the sequence $\left(u^{(s+1)T-t+1}[\pi(s,l)]\right)_{s\geq 0}$ converges to x. ■

Lemma 13 *As the time horizon increases, the set of pure strategy Nash equilibrium payoff vectors of the finitely repeated game converges to the set* ANE(G).

The proof of this lemma is similar to the one of Lemma 5 and therefore omitted.

7.2 On the Nash feasibility of pure strategy Nash equilibrium payoff vectors of the finitely repeated game

Lemma 14 *For any $T > 0$ and any pure strategy Nash equilibrium σ of the finitely repeated game $G(T)$, the support $\{\pi_1(\sigma), \cdots, \pi_T(\sigma)\}$ of the play path $\pi(\sigma) = (\pi_1(\sigma), \cdots, \pi_T(\sigma))$ generated by σ is included in the set* Nash(G^{*h}) *of pure Nash equilibria of the one shot game G^{*h}.*

Proof of Lemma 14. I proceed by induction on the time horizon T. For $T = 1$, σ is a pure Nash equilibrium of the stage-game G. As the sequence of sets $(\text{Nash}(G^{*l}))_{l\geq 0}$ is increasing, we have $\text{Nash}(G) = \text{Nash}(G^{*0}) \subseteq \text{Nash}(G^{*h})$ and the support $\{\pi_1(\sigma)\}$ of the play path $\pi(\sigma)$ is included in $\text{Nash}(G^{*h})$.
Assume that $T > 1$ and that the support of the play path generated by any pure strategy Nash equilibrium of the finitely repeated game $G(t)$ with $0 < t < T$ is included in $\text{Nash}(G^{*h})$ and let's show that $\{\pi_1(\sigma), \cdots, \pi_T(\sigma)\} \subseteq \text{Nash}(G^{*h})$.
The restriction $\sigma_{|\pi_1(\sigma)}$ of the strategy profile σ to the history $\pi_1(\sigma)$ is a pure strategy Nash equilibrium of the finitely repeated game $G(T-1)$ and by induction hypothesis, the support $\{\pi_2(\sigma), \cdots, \pi_T(\sigma)\}$ of $\sigma_{|\pi_1(\sigma)}$ is included in the set $\text{Nash}(G^{*h})$. It remains to

prove that $\pi_1(\sigma) \in \text{Nash}(G^{*h})$. Suppose that $\pi_1(\sigma) \notin \text{Nash}(G^{*h})$. Then there exists a player $i \in N$ who has an incentive to deviate from the pure action profile $\pi_1(\sigma)$ in the game G^{*h}. Player i has to be a member of the block $N \backslash N_h^*$ since each player of the block N_h^* has a constant utility function in the game G^{*h}.
Let σ_i' be the pure strategy of player i in the finitely repeated game $G(T)$ in which player i plays a stage-game pure best response at each round of the finitely repeated game. There is no lost if we assume that σ is the grim trigger strategy profile associated to the path $\pi(\sigma)$.[16]
At the pure strategy profile (σ_i', σ_{-i}), player i receives the sequence of stage-game payoffs

$$\{u_i(\pi_1(\sigma)) + e, n_i^*, \cdots, n_i^*\}$$

whereas at σ she receives

$$\{u_i(\pi_1(\sigma)), n_i^*, \cdots, n_i^*\}$$

where $e > 0$ and n_i^* is her unique pure Nash equilibrium payoff in the stage-game G. This implies that $u^T(\sigma_i', \sigma_{-i}) > u^T(\sigma)$. The pure strategy σ_i' is therefore a profitable deviation of player i from σ. This contradicts the fact that σ is a pure strategy Nash equilibrium of the finitely repeated game $G(T)$. It follows that $\pi_1(\sigma) \in \text{Nash}(G^{*h})$, which concludes the proof. ■

From Lemma 14, it follows that only the payoff vectors of the convex hull F of the set $u(Nash(G^{*h})) = \{u(a) \mid a \in \text{Nash}(G^{*h})\}$ can be sustainable by pure strategy Nash equilibria of the finitely repeated game. We have the following corollary.

Corollary 2 *For any $T > 0$ and for all pure strategy Nash equilibrium σ of the finitely repeated game $G(T)$, the average payoff vector $u^T(\sigma)$ belongs to the set F of Nash-feasible payoff vectors of the stage-game G.*

7.3 Proof of Theorem 2

From Corollary 2, any pure strategy Nash equilibrium payoff vector of any finite repetition of the stage-game has to be Nash-feasible. Denoting by I the set of payoff vectors that dominate the minimax payoff vector μ, we have that $NE(T) \subseteq F \cap I$ for all $T \geq 1$.

[16]The grim trigger strategy profile associated to a path $\pi \in A^T$ is a strategy profile σ^π of the finitely repeated game $G(T)$ in which players follow the path π until a unique player deviates. After a unilateral deviation has been observed, the grim trigger strategy profile prescribes to punish the deviator by pushing her down to her minimax payoff till the end of the game. It is straightforward to see that a path is a pure strategy Nash equilibrium play path of the finitely repeated game if and only if the grim trigger strategy profile associated to that path is a pure strategy Nash equilibrium of that finitely repeated game.

Lemma 16 says that any payoff vector $x \in F \cap I$ is approachable by means of pure strategy Nash equilibria of the finitely repeated game. This lemma concludes the proofs of both Theorem 2 and Theorem 4 as the limit set NE equals the set ANE(G) of payoff vectors that are approachable by means of pure strategy Nash equilibria of the finitely repeated game; see Lemma 13. I first construct an appropriate end-game strategy.

Similarly to the case of pure strategy subgame perfect Nash equilibrium solution, the sequence of subsets $(N_l^*)_{l\geq 0}$ defined in Section 4.1 induces a separation of the set of players into two blocks N_h^* and $N\backslash N_h^*$. As a corollary of Lemma 14, each player of the block $N\backslash N_h^*$ (if any) receives her unique stage-game pure Nash equilibrium payoff at each pure strategy Nash equilibrium of any finite repetition of the stage-game G.[17] The next lemma says that there exists a pure strategy Nash equilibrium of a finite repetition of the stage-game G where each player of the block N_h^* receives an average payoff that is strictly greater than her pure minimax payoff.

Lemma 15 *Let G be a compact normal form game and $\emptyset = N_0^* \subsetneq N_1^* \subsetneq \cdots \subsetneq N_h^*$ its decomposition.*[18] *Then there exists $T_0 > 0$ and a pure strategy Nash equilibrium of the repeated game $G(T_0)$ at which each player of the block N_h^* receives an average payoff that is strictly greater than her stage-game pure minimax payoff.*

Proof of Lemma 15. I will prove the following property by induction on g: for all $g \leq h$ and all $i \in N_g^*$, there exists $T_{i,g} > 0$ and a pure strategy Nash equilibrium of the repeated game $G(T_{i,g})$ at which player i receives an average payoff that is strictly greater than her stage-game pure minimax payoff.
For $g = 1$, take $T_{i,g} = 1$ for each $i \in N_1^*$.
Fix $g \in \{1, \cdots, h-1\}$ and assume that the property holds for g. Pose $N_g^* = \{j_1, \cdots, j_m\}$. For all $j \in N_g^*$, let $T_{j,g} > 0$ and let π^j be a play path generated by a pure strategy Nash equilibrium of the finitely repeated game $G(T_{j,g})$ at which player j receives an average payoff that is strictly greater than her stage-game pure minimax payoff. Let $\pi^{N_g^*} = (\pi^{j_1}, \cdots, \pi^{j_m})$. The trigger strategy associated to $\pi^{N_g^*}$ is a pure strategy Nash equilibrium of the repeated game $G(\sum_{j\in N_g^*} T_{j,g})$ and the average payoff of each player of the block N_g^* at that Nash equilibrium is strictly greater than her stage-game pure minimax payoff.[19]
Let $i \in N_{g+1}^* \backslash N_g^*$ and let $y^{i,g}$ be the best pure Nash equilibrium profile of player i in the

[17] Indeed, at any profile of action $a \in \text{Nash}(G^{*h})$, each player of the block N_h^* receives her unique stage-game pure Nash equilibrium payoff vector. This payoff equals her stage-game pure minimax payoff.

[18] See Section 4.1 for the definition of the sequence $(N_l^*)_{l\geq 0}$.

[19] Note that each player of the block $N\backslash N_g^*$ plays a stage-game pure best response at any profile of actions of the path $\pi^{N_g^*}$.

one shot game G^{*g}. There exists $k > 0$ such that the trigger strategy associated to the path

$$(y^{i,g}, \underbrace{\pi^{N_g^*}, \cdots, \pi^{N_g^*}}_{k \text{ times}})$$

is a pure strategy Nash equilibrium of the repeated game $G(1 + k \cdot \sum_{j \in N_g^*} T_{j,g})$. At the later Nash equilibrium, player i receives an average payoff that is strictly greater than her stage-game pure minimax payoff. Take $T_{i,g+1} = 1 + k \cdot \sum_{j \in N_g^*} T_{j,g}$. This concludes the proof of the lemma. ■

Lemma 16 *Let G be a compact normal form game. Any Nash-feasible and individually rational payoff vector is approachable by means of pure strategy Nash equilibria of the finitely repeated game.*

Proof of Lemma 16. Let x be a Nash-feasible and individually rational payoff vector and $\varepsilon > 0$. I wish to construct a time horizon $T_{\varepsilon,x}$ such that for all $T \geq T_{\varepsilon,x}$, the finitely repeated game $G(T)$ has a pure strategy Nash equilibrium $\sigma^{\varepsilon,x,T}$ satisfying $d(x, u^T(\sigma^{\varepsilon,x,T})) < \varepsilon$.

Let $x' \in F \cap I$ such that

$$d(x, x') \leq \tfrac{\varepsilon}{8}$$

and $x'_i > \mu_i$ for all $i \in N_h^*$.[20]

Since $\mathbb{Q}$ is dense in $\mathbb{R}$, there exists a sequence $(\gamma_t)_{1 \leq t \leq p}$ of strictly positive rationals numbers and a sequence $(a^t)_{1 \leq t \leq p}$ of elements of $\text{Nash}(G^{*h})$ such that

$$d(x', \textstyle\sum_{t=1}^p \gamma_t u(a^t)) < \tfrac{\varepsilon'}{8}$$

and $\sum_{t=1}^p \gamma_t = 1$ where

$$\varepsilon' = \min\{\tfrac{\varepsilon}{2}, \min_{i \in N_h^*}(x'_i - \mu_i)\}$$

is strictly positive. Let $x'' = \sum_{t=1}^p \gamma_t u(a^t)$. We have $u_i(a^t) = \mu_i$ for all $t, 1 \leq t \leq p$ and $i \notin N_h^*$. Thus, $x''_i = \mu_i$ for all $i \notin N_h^*$. We also have $x''_i > \mu_i$ for all $i \in N_h^*$. This holds since $d(x', x'') < x'_i - \mu_i$ for all $i \in N_h^*$. Consider a sequence of natural numbers $(q_t)_{1 \leq t \leq p}$ such that for all $t, t' \in \{1, ..., p\}$ we have $\frac{\gamma_t}{\gamma_{t'}} = \frac{q_t}{q_{t'}}$. Let $q = \sum_{t=1}^p q_t$ and

$$\pi = (\underbrace{a^1, a^1, \cdots, a^1}_{q_1 \text{ times}}, \cdots, \underbrace{a^p, a^p, \cdots, a^p}_{q_p \text{ times}}).$$

[20] One could take $x' = x + \frac{\varepsilon}{8 \cdot d(x,y)}(y - x)$ where y is the average payoff vector to the pure Nash equilibrium given by Lemma 15.

Let π^h be a play path generated by a pure strategy Nash equilibrium of the repeated game $G(T_0)$ at which each player of the block N_h^* receives an average payoff that is strictly greater than her stage-game pure minimax payoff, see Lemma 15. There exists $k > 0$ such that the trigger strategy associated to the path

$$\pi(q) = (\pi, \underbrace{\pi^h, \cdots, \pi^h}_{k \text{ times}})$$

is a pure strategy Nash equilibrium of the repeated game $G(q + kT_0)$. Let $\widehat{\pi}(s, q)$ be the play path defined by

$$\widehat{\pi}(s, q) = (\underbrace{\pi, \cdots, \pi}_{s \text{ times}}, \pi(q)).$$

The grim trigger strategy profile $\sigma^{\widehat{\pi}(s,q)}$ associated to $\widehat{\pi}(s, q)$ is a pure strategy Nash equilibrium of the finitely repeated game $G(u^{(s+1)q+kT_0})$. As s increases, the payoff vector $u^{(s+1)q+kT_0}(\sigma^{\widehat{\pi}(s,q)})$ converges to x''. Therefore, there exists $s_{\varepsilon,x} > 0$ such that for all $s \geq s_{\varepsilon,x}$, $d(x', u^{(s+1)q+kT_0}(\sigma^{\widehat{\pi}(s,q)})) < \frac{\varepsilon}{8}$. Choose $s_{\varepsilon,x}$ large enough such that $\frac{\rho}{s} < \frac{\varepsilon}{8}$ for all $s > s_{\varepsilon,x}$ and take $T_{\varepsilon,x} = (s_{\varepsilon,x} + 1)q + kT_0$. ■

8 Appendix 3: In case there exists a discount factor

If there exists a discount factor, then one only has to adjust the proofs of Lemmata 10 and 16. In the proof of Lemma 10, one can apply Lemma 1 to y and obtain π^y and thereafter use the discounted version of Lemma 9, see Lemma 17 below. To adjust the proof of Lemma 16, one can apply Lemma 1 to $\varepsilon = \frac{\varepsilon'}{8}$ and obtain a deterministic path π whose discounted average is within ε of x'.

Lemma 17 *Let $\emptyset = N_0 \subsetneq N_1 \subsetneq ... \subsetneq N_h$ be the Nash decomposition of the stage-game G. Then there exists $\phi > 0$ such that for all $p \geq 0$ there exists $r_p > 0$, $\delta_p \in (0, 1)$ and*

$$\theta^p : \{0, 1\}^n \cup \{(-1, \cdots, -1)\} \to A^{r_p} := A \times \cdots \times A$$

such that for all $\alpha \in \{0, 1\}^n \cup \{(-1, \cdots, -1)\}$ and $\delta \in (\delta_p, 1)$, $\theta^p(\alpha)$ is a play path generated by a pure strategy subgame perfect Nash equilibrium of the repeated game with discounting $G(\delta, r_p)$.[21] *Furthermore, for all $i \in N_h$ and $\alpha, \alpha' \in \{0, 1\}^n$ and $\delta \in (\delta_p, 1)$, we have*

$$u_i^{r_p,\delta}[\theta^p(1, \alpha_{-i})] - u_i^{r_p,\delta}[\theta^p(0, \alpha_{-i})] \geq \phi, \tag{12}$$

[21] I recall that in the discounted repeated game $G(\delta, r_p)$, the utility of player i at the play path $\theta^p(\alpha)$ is $u_i^{r_p,\delta}[\theta^p(\alpha)] = \frac{1-\delta}{1-\delta^{r_p}} \sum_{t=1}^{r_p} \delta^{t-1} u_i(\theta_t^p(\alpha))$, where $\theta_t^p(\alpha)$ is the t th profile of action of $\theta^p(\alpha)$.

$$u_i^{r_p,\delta}[\theta^p(\alpha)] - u_i^{r_p,\delta}[\theta^p(-1,\cdots,-1)] \geq \phi \tag{13}$$

and

$$|u_i^{r_p,\delta}[\theta^p(\alpha)] - u_i^{r_p,\delta}[\theta^p(\alpha_{\mathcal{J}(i)}, \alpha'_{N\setminus\mathcal{J}(i)})]| < \frac{1}{2^p}. \tag{14}$$

References

Dilip Abreu, Prajit K Dutta, and Lones Smith. The folk theorem for repeated games: A NEU condition. *Econometrica*, 62(4): 939–948, 1994.

J.-P. Benoit and Vijay Krishna. Finitely repeated games. *Econometrica*, 53(1): 905–922, 1984.

J.-P. Benoit and Vijay Krishna. Nash equilibria of finitely repeated games. *International Journal of Game Theory*, 16(3): 197–204, 1987.

Drew Fudenberg and Eric Maskin. The folk theorem in repeated games with discounting or with incomplete information. *Econometrica*, 54(3): 533–554, 1986.

Julio González-Díaz. Finitely repeated games: A generalized Nash folk theorem. *Games and Economic Behavior*, 55(1): 100–111, 2006.

Olivier Gossner. The folk theorem for finitely repeated games with mixed strategies. *International Journal of Game Theory*, 24(1): 95–107, 1995.

Lones Smith. Necessary and sufficient conditions for the perfect finite horizon folk theorem. *Econometrica*, 63(2): 425–430, 1995.

Quan Wen. The "folk theorem" for repeated games with complete information. *Econometrica*, 62(4): 949–954, 1994.