

Demeze-Jouatsa, Ghislain-Herman

Article — Published Version

A complete folk theorem for finitely repeated games

International Journal of Game Theory

Provided in Cooperation with:

Springer Nature

Suggested Citation: Demeze-Jouatsa, Ghislain-Herman (2020) : A complete folk theorem for finitely repeated games, International Journal of Game Theory, ISSN 1432-1270, Springer, Berlin, Heidelberg, Vol. 49, Iss. 4, pp. 1129-1142, <https://doi.org/10.1007/s00182-020-00735-z>

This Version is available at:

<https://hdl.handle.net/10419/288449>

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.



<https://creativecommons.org/licenses/by/4.0/>



A complete folk theorem for finitely repeated games

Ghislain-Herman Demeze-Jouatsa¹ 

Accepted: 26 August 2020 / Published online: 28 September 2020

© The Author(s) 2020

Abstract

This paper analyzes 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. This model includes the special case of observable mixed strategies.

Keywords Finitely repeated games · Pure strategy · Observable mixed strategies · Subgame perfect Nash equilibrium · Limit perfect folk theorem

JEL Classification C72 and 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

I acknowledge DAAD and the DFG (Deutsche Forschungsgemeinschaft/German Research Foundation) via Grant Ri 1128-9-1 (Open Research Area in the Social Sciences, Ambiguity in Dynamic Environments) for funding support and thank Christoph Kuzmics, Frank Riedel, Lones Smith, Michael Greinecker, Karl Schlag, Roland Pongou, Tondji-Jean Baptiste and Olivier Gossner for useful comments. I also thank seminar and conference participants at Bielefeld University, Cardiff University, University of Yaoundé I, Stony Brook University, Lisbon School of Economics and Management, and University of Graz for positive inputs. I thank two anonymous reviewers and an associated editor for their comments and suggestions.

✉ Ghislain-Herman Demeze-Jouatsa
demeze_jouatsa@uni-bielefeld.de

¹ Center for Mathematical Economics, Bielefeld University, Bielefeld, Germany

payoff defined by Wen (1994). The main theorem nests earlier results of Benoit and Krishna (1985), Smith (1995), and Demeze-Jouatsa and Wilson (2019).

Whether non-Nash outcomes of the stage-game can be sustained via subgame perfect Nash equilibria of the finitely repeated game depends on whether players 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. On the other hand, for stage-games in which all players receive different Nash equilibrium payoffs such as the battle of sexes, the limit perfect folk theorem holds: 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 (1985) 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.¹ 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.² The distinct Nash payoffs condition and the full dimensionality of the set of feasible payoff vectors as in Benoit and Krishna (1985) 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 finitely repeated 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 (1985) 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 stage-games that do 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?³

¹ 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.

² 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.

³ Consider the following 4-player stage-game. Player 1 is a social planner who wishes to maximize the welfare of a representative agent, player 2 is a monopolist in the market, and players 3 and 4 are potential entrants to the market. The latter entrants need both approval and support of the planner (player 1) to enter the

If the stage-game has an incomplete Nash decomposition, then the set of players naturally breaks up into two blocks where the first block contains all the players whose behavior can recursively be leveraged near the end of the finitely repeated game (see Footnote 2 for details). In contrast, it is not possible to control short run incentives of players of the second block. 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 via pure strategy subgame perfect Nash equilibria of the finitely repeated game, it has to be in the convex hull of the set of payoffs to profiles of actions that are Nash equilibria of the effective one shot game. I introduce the concept of enforceable payoff vector. I call a payoff vector enforceable 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

Footnote 3 continued

market and make profit. Player 1 can choose either to open the market to the potential entrants (OM) or to keep it closed (CM), player 2 can choose either to cooperate (C) or to deviate (D) and maximize her profit in the detriment of consumers. Player 3 can choose up (U) or down (D) and player 4 can choose L or R. The Payoffs of the game are given by the following table:

		C		D	
		L	R	L	R
CM	U	3 3 0 0	3 3 0 0	0 6 0 0	0 6 0 0
	D	3 3 0 0	3 3 0 0	0 6 0 0	0 6 0 0
OM	U	2 0 3 1	2 0 1 3	2 2 0 0	1 0 0 0
	D	2 0 1 3	1 1 3 1	1 0 0 0	1 1 0 0

This game admits two pure Nash equilibrium profiles with respective payoff vectors (2,2,0,0) and (1,1,0,0), which are both strictly Pareto-dominated, for instance by the feasible payoff $(\frac{17}{8}, \frac{19}{8}, 1, 1)$. One might wonder if in the finite repetitions of this game players could cooperate and achieve efficiency via equilibrium strategies of the repeated game. As only players 1 and 2 receive distinct payoffs at pure Nash equilibria of the stage-game, the distinct Nash payoff condition of Benoit and Krishna (1985) does not hold. Furthermore, given any fixed profile of action of players 1 and 2, the 2-player induced game (played by players 3 and 4) does not admit a pure Nash equilibrium that has a payoff that is different from (0, 0). This means that the complete Nash decomposition condition of Smith (1995), which is necessary and sufficient for the finite horizon perfect folk theorem, does not hold. As a consequence, the finite horizon perfect folk theorem does not hold for this game. If it is immediate that players can cooperate and achieve some payoffs that are weakly Pareto-superior to any stage-game Nash equilibrium payoff, for instance (2.7, 2.7, 0, 0) in 10 repetitions, it is not clear what is the exact set of payoffs vectors that can be achieved via subgame perfect equilibria of the finitely repeated game.

strictly greater than her effective minimax payoff. I call this new reservation payoff the enforceable minimax payoff.

The main finding of this paper says that, as the time horizon increases, the set of payoff vectors to pure strategy subgame perfect Nash equilibria of the finitely repeated game converges to the set of enforceable payoff vectors that dominate the enforceable minimax payoff vector.

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

2 The model

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 $N = \{1, \dots, n\}$ is finite and where for all player $i \in N$ the set A_i of actions of player i is compact. Given a player $i \in N$ and an action profile $a = (a_1, \dots, a_n) \in A$, 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 $a'_i \in A_i$.

Each stage-game considered in this paper is *compact* in the sense that each A_i is compact, and u is continuous. A stage-game could for instance be finite, the mixed extension of another finite stage-game, or a game with a continuum of actions for some players.

Let γ be a real number that is strictly greater than any payoff a player might receive in the stage-game G .⁴ A player is said 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 $\tilde{A} = \text{Nash}(G^h)$ be the set of pure Nash equilibria of the game G^h . We call \tilde{A} the *enforceable action set*. The set of

⁴ 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 guarantees the existence of γ .

enforceable payoff vectors of the game G is defined as the convex hull $Conv[u(\tilde{A})]$ of the set $u(\tilde{A}) = \{u(a) \mid a \in \tilde{A}\}$. The sequence $0 \subsetneq N_1 \subsetneq \dots \subsetneq N_h$ is the *Nash decomposition* of the game G , and the Nash decomposition is complete if $N_h = N$.⁵

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 exist $\alpha_{ij} > 0$ and $\beta_{ij} \in \mathbb{R}$ such that for all $a \in \tilde{A}$, 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

$$\tilde{\mu}_i = \min_{a \in \tilde{A}} \max_{j \in \mathcal{J}(i)} \max_{a'_j \in A_j} [\alpha_{ij} \cdot u_j(a'_j, a_{-j}) + \beta_{ij}] \text{ and } \tilde{\mu} = (\tilde{\mu}_1, \dots, \tilde{\mu}_n).$$

The payoff $\tilde{\mu}_i$ is the *enforceable minimax* of player i in the stage-game G .⁶

Call a payoff vector *e-rational* if it dominates the enforceable minimax payoff vector $\tilde{\mu}$. Let $\tilde{I} = \{x = (x_1, \dots, x_n) \in \mathbb{R} \mid x_i \geq \tilde{\mu}_i \text{ for all } i \in N\}$ be the set of e-rational payoff vectors.

The name “enforceable action” comes from Fudenberg et al. (2009) notion of enforceability, which requires an action to be incentive compatible given some set of continuation payoffs. The concepts of enforceable payoff and enforceable minimax are respective generalizations of the classic concepts of feasible payoff and effective minimax, viewed as indicators to derive the perfect folk theorem for finitely repeated games. If each player receives (recursively) distinct payoffs at Nash equilibria of the stage-game, then the behavior of each player can be leveraged if the game is finitely repeated. In that case, the enforceable action set \tilde{A} equals the whole set A of action profiles, the enforceable minimax equals the classic effective minimax, and the set of enforceable payoff vectors equals the set of feasible payoff vectors. In the other case, the set of enforceable actions \tilde{A} is a proper subset of the whole set of profile of pure actions, the set of enforceable payoff vectors is a proper subset of the classic set of feasible payoff vectors, and the enforceable minimax of a player can be strictly greater

⁵ While being equivalent to Smith’s (1995) definition of Nash decomposition, ours is simpler and requires to analyse no more than n simple transformations of the stage-game, while Smith’s definition requires, in many cases, to analyse at least 2^{n-1} subgames, n being the number of players. 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 enforceable payoff vectors equals the classic set of feasible payoff vectors and the enforceable minimax payoff vector equals the classic effective minimax payoff vector. In that case, the main result (see Theorem 1) says that any feasible payoff vector that dominates the effective minimax payoff vector is approachable via pure strategy subgame perfect Nash equilibria of the finitely repeated game. That is the message of the limit perfect folk theorem. Benoit and Krishna (1985) 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 the main result of this paper, Theorem 1. Indeed, under the distinct stage-game Nash equilibrium payoffs condition of Benoit and Krishna (1985), the Nash decomposition of the stage-game equals $\emptyset \subsetneq N_h = N$ which is complete and therefore the set of the enforceable payoff vectors equals the classic set of the feasible payoff vectors and the enforceable 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.

⁶ 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 \tilde{A} is non-empty and $\tilde{\mu}$ is well defined.




Fig. 1 Equilibrium payoff vectors of players 1 and 2

than her effective minimax.⁷ Figure 1 uses the example of Footnote 3 to illustrate the differences between our newly introduced concepts and the classic ones. Only payoffs of players 1 and 2 are displayed. In that game, pure action profiles where player 1 chooses OM and player 2 chooses C are not enforceable. The effective minimax payoff vector equals $(1, 0, 0, 0)$ and the enforceable minimax payoff vector equals $(1, 1, 0, 0)$.⁸

2.2 The finitely repeated game

Let G be the stage-game. Given $T > 0$, let $G(T)$ denote the T -fold 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 function $\sigma_i : \bigcup_{t=1}^T A^{t-1} \rightarrow A_i$ where A^0 contains only the empty history.⁹ The strategy profile $\sigma = (\sigma_1, \dots, \sigma_n)$ of $G(T)$ generates a **play path** $\pi(\sigma) = [\pi_1(\sigma), \dots, \pi_T(\sigma)] \in A^T$ 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 payoff $u_i^T(\sigma) = \frac{1}{T} \sum_{t=1}^T u_i[\pi_t(\sigma)]$. A strategy profile $\sigma = (\sigma_1, \dots, \sigma_n)$ 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 $\sigma = (\sigma_1, \dots, \sigma_n)$ is a *pure strategy subgame perfect Nash equilibrium of $G(T)$* if given any $t \in \{1, \dots, T\}$ 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 $G(T - t + 1)$.

⁷ Exceptions are degenerated games where at least one player has a constant utility function.

⁸ Note that the enforceable minimax payoff vector weakly Pareto-dominates the effective minimax payoff vector.

⁹ If the stage-game G is the mixed extension of another finite stage-game and mixed actions are observable, then players can condition their actions on the mixed actions themselves instead of outcomes of randomization devices.

For any $T > 0$, let $E(T)$ be the set of pure strategy 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.¹⁰ The set E is the Hausdorff limit of the set of pure strategy subgame perfect Nash equilibrium payoff vectors of the finitely repeated game. As I show later in the Appendix A, the limit set E exists, is nonempty, convex, and compact.

3 Main result

Theorem 1 *Let G be a compact stage-game. As the time horizon increases, the set of pure strategy subgame perfect Nash equilibrium payoff (average payoff) vectors of the finitely repeated game converges (in the Hausdorff sense) to the set of enforceable and ϵ -rational payoff vectors.*

A constructive proof of Theorem 1 is provided in the appendix. It uses four main lemmata. Lemma 3 states that as the time horizon increases, the set of pure strategy subgame perfect Nash equilibrium payoffs of the finitely repeated game converges to a well defined set, $ASPNE(G)$, which is the set of payoffs that are approachable via pure strategy subgame perfect Nash equilibria. Lemmata 4 and 5 together say that the limit set of the set of pure strategy subgame perfect Nash equilibrium payoff vectors, which equals the set $ASPNE(G)$, is included in the set of enforceable and ϵ -rational payoff vectors. Lemma 6 states that every enforceable and ϵ -rational payoff vector belongs to the set $ASPNE(G)$. The enforceability and the ϵ -rationality can therefore be observed as necessary and sufficient conditions on feasible payoffs to be approachable via pure strategy subgame perfect Nash equilibria of the finitely repeated game.

Theorem 1 assumes no discounting. This assumption is without loss of generality. One can indeed check that as the discount factor goes to 1, the discounted average converges to the average payoff. Therefore, if the average payoff of a player to a path π is strictly greater than her average payoff to another path π' , then the discounted average payoff of that player to the path π is strictly greater than her discounted average to π' , given that the discount factor is high enough. One can also make use of the payoff continuation lemma for finitely repeated games and prove a stronger result.¹¹ With a fixed discount factor, one can show that the limit set of the set of pure strategy subgame perfect Nash equilibrium payoffs of the discounted finitely repeated game equals the set of enforceable and ϵ -rational payoffs, given that the discount factor exceeds a threshold δ .

¹⁰ 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) . 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 $d(x, Y) = \inf_{y \in Y} d(x, y)$.

¹¹ Fudenberg and Maskin (1991) provides a payoff continuation lemma for infinitely repeated games with discounting. The payoff continuation lemma for finitely repeated games say that: for any $\epsilon > 0$, there exists $k > 0$ and $\delta < 1$ 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 $\delta \geq \delta$. A proof of this lemma is provided in Demeze-Jouatsa (2019).

4 Conclusion

This paper analysed 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 compact, 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. This finding exhibits the exact range of cooperative payoffs that players can achieve via subgame perfect Nash equilibria of the finitely repeated game.

One point of this work is that it provides a full characterization of the optimal punishment payoff of finitely repeated games with complete information and perfect monitoring (Benoit and Krishna 1985, Gossner and Hörner 2010).

The method of this paper applies to the Nash equilibrium case. In this particular case, to leverage the behaviour of a player near the end of the finitely repeated game, it is necessary and sufficient that the latter player either has a pure Nash equilibrium payoff that is strictly greater than her pure minimax payoff, or that there exists a recursive Nash equilibrium in which the latter player receives a payoff that is different from her pure minimax payoff. Pure actions eligible for pure strategy Nash equilibrium play paths of the finitely repeated game are therefore the pure Nash equilibrium profiles of a new stage-game obtained from the original stage-game by setting the utility functions of players whose behaviour can be leverage near the end of the finitely repeated game to a constant. I refer to the convex hull of the set of original payoffs to eligible profiles as the set of Nash-feasible payoffs. As the time horizon increases, the set of pure strategy Nash equilibrium payoff vectors of the finitely repeated game converges to the set of Nash-feasible payoffs vectors that dominate the pure minimax payoff vector. This characterization of the limit set of the set of Nash equilibrium payoff vectors of the finitely repeated game nests early results of Benoit and Krishna (1987) and González-Díaz (2006).

One might wonder if similar method applies in the case that players can employ unobservable mixed strategies (Gossner 1995), in the case that the monitoring technology is imperfect (see Fudenberg et al. 2007 for the infinite horizon case, and with public monitoring), or in the case that equilibrium strategies are protected against renegotiation (Benoit and Krishna 1993).

Funding Open Access funding provided by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Appendix

A. Existence of the limit set of the set of equilibrium payoffs

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 compact stage-game, the set of feasible payoff vectors that are approachable via 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-henceforth, Lemma CBK-established by Benoit and Krishna (1985). Lemma CBK states 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.

Let G be a compact stage-game and let $ASPNE(G)$ be the set of all feasible payoff vectors of G that are approachable via pure strategy subgame perfect Nash equilibrium payoff vectors of the finitely repeated game.¹²

Lemma 1 *The set $ASPNE(G)$ is compact and convex.*

Proof of Lemma 1 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 $x, y \in ASPNE(G)$, which follows from Lemma CBK. \square

Lemma 2 *For all $T > 0$, $E(T) \subseteq ASPNE(G)$.*

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

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

be a play path of $G((s + 1)T - t + 1)$. From Lemma CBK, $\pi(s, t)$ 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 $(u^{(s+1)T-t+1}[\pi(s, t)])_{s \geq 0}$ converges to x . \square

Lemma 3 *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 $ASPNE(G)$.*

¹² A feasible payoff vector x is approachable via 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 .

Proof of Lemma 3 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 $\{B(x^l, \frac{\varepsilon}{2}) \mid x^l, l = 1, \dots, L\}$ be a finite covering of $\text{ASPNE}(G)$, where $B(x, \varepsilon) = \{y \in \mathbb{R}^n \mid d(x, y) < \varepsilon\}$. For all $l = 1, \dots, L$ take T_0^l given by the definition of “ $x^l \in \text{ASPNE}(G)$ ” with $\frac{\varepsilon}{2}$. Pose $T_0 = \max_{l \leq L} T_0^l$. Let $T > T_0$ and let $x \in \text{ASPNE}(G)$. Let $x^{l_0} \in \text{ASPNE}(G)$ 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, $\sup_{x \in \text{ASPNE}(G)} d(x, E(T)) \leq \varepsilon$. Furthermore, from Lemma 2, $d(y, \text{ASPNE}(G)) = 0$ 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 $T_\varepsilon = T_0$. \square

B. Necessity of enforceability

Lemma 4 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 4 If $N_h = N$, then $\text{Nash}(G^h) = A$ and $\text{Supp}(\pi(\sigma)) \subseteq \text{Nash}(G^h)$. Now assume that $N \setminus N_h \neq \emptyset$.

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 , and $\text{Nash}(G) = \text{Nash}(G^0) \subseteq \text{Nash}(G^h)$.

Suppose that $T > 1$ and that the support of any pure strategy 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 $N \setminus N_h$ 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 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 $\text{Nash}(G^h)$. Therefore, at the profile (σ'_i, σ_{-i}) player i receives the sequence of stage-game payoffs

$\{u_i(\pi^1) + e, n_i, \dots, n_i\}$ where n_i is her unique stage-game pure Nash equilibrium payoff. Since player i receives $\{u_i(\pi_1(\sigma)), n_i, \dots, n_i\}$ 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. \square

Let \tilde{F} be the set of enforceable 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 \tilde{F} .*

C. Necessity of the e-rationality

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 sharper lower bound of the set of equilibrium payoffs of the finitely repeated game. The lemma says that, any pure strategy subgame perfect Nash equilibrium payoff vector of the finitely repeated game weakly dominates the enforceable minimax payoff vector.

Lemma 5 *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 enforceable minimax payoff vector of the stage-game.*

Proof of Lemma 5 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 $\tilde{\mu}$.¹³

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 enforceable minimax payoff vector $\tilde{\mu}$. Let us show that the payoff vector $u^T(\sigma)$ dominates $\tilde{\mu}$.

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 $\tilde{\mu}$. Suppose now that $u^T(\sigma)$ does not dominates $\tilde{\mu}$. Then there exists a player $i \in N$ such that $u_i^T(\sigma) < \tilde{\mu}_i$. It follows that $u_i[\pi_1(\sigma)] < \tilde{\mu}_i$ 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 $u_j[\pi_1(\sigma)] < \tilde{\mu}_j$ for all $j \in \mathcal{J}(i)$. From the definition of $\tilde{\mu}$, 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 $u_{i_0}[a_{i_0}, \pi_1(\sigma)_{-i_0}] \geq \tilde{\mu}_{i_0}$. Consider the pure strategy one shot deviation σ'_{i_0} of player i_0 from σ in which she

¹³ 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 enforceable minimax payoff. It follows that any pure Nash equilibrium payoff vector weakly dominates the enforceable minimax payoff vector.

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

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

which is greater than or equal to $\tilde{\mu}_{i_0}$. 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 $u^{T-1}(\sigma_{|(a_{i_0}, \pi_1(\sigma)_{-i_0})})$ dominates $\tilde{\mu}$. It follows that the deviation σ'_{i_0} of player i_0 is profitable. This contradicts the fact that σ is a subgame perfect Nash equilibrium. \square

D. Proof of the complete perfect folk theorem

From Corollary 1 and Lemma 5, 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 enforceable and e-rational payoff vectors. To complete the proof of Theorem 1, it is left to show that any enforceable and e-rational payoff vector belongs to the limit set E . In what follows, I prove that any enforceable and e-rational payoff vector is approachable via pure strategy subgame perfect Nash equilibria of the finitely repeated game (see Lemma 3).

Lemma 6 *Let G be a compact normal form game. We have $\tilde{F} \cap \tilde{I} \subseteq \text{ASPNE}(G)$.*

Proof of Lemma 6 Let G be a compact normal form game. If G admits no pure Nash equilibrium, then $\tilde{F} = \emptyset$ and $\tilde{F} \cap \tilde{I} \subseteq \text{ASPNE}(G)$. If G admits a unique pure Nash equilibrium payoff vector x , then $\tilde{F} = \{x\} = \text{ASPNE}(G)$ and $\tilde{F} \cap \tilde{I} \subseteq \text{ASPNE}(G)$. Now suppose that G admits at least two distinct pure Nash equilibrium payoff vectors. Normalize the game such that the enforceable minimax of each player equals 0 and such that two equivalent players have the same utility function on \tilde{A} . Consider

$$F_1 = \left\{ \frac{1}{p} \sum_{1 \leq l \leq p} u(a^l) \mid p > 0, a^l \in \tilde{A} \forall l \leq p \right\} \text{ and}$$

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

It is immediate that the closure of $F_1 \cap I_1$ is equal to the set $\tilde{F} \cap \tilde{I}$. From Lemma 1, $\text{ASPNE}(G)$ is closed. Therefore, it is enough to show that $F_1 \cap I_1 \subseteq \text{ASPNE}(G)$. Let $y = \frac{1}{p} \sum_{1 \leq l \leq p} u(a^l) \in F_1 \cap I_1$. For all $i \in N_h$, let $w^i \in \arg \min_{a \in \tilde{A}} \max_{j \in \mathcal{J}(i)} \max_{a'_j \in A_j} u_i(a'_j, a_{-j})$.¹⁴

I use a suitable but minimal adjustment of the 5-phase strategy presented in Smith (1995) and Demeze-Jouatsa and Wilson (2019) to cases where the Nash decomposition is incomplete to approximate the target payoff vector y .

¹⁴ At the profile of actions w^i , player i does not have to be at a pure best response. If she plays a pure best response to w^i , she receives at least her stage-game pure minimax payoff but no more than her stage-game enforceable minimax payoff.

For each $g \in \{1, \dots, h\}$, let $e_{N_{g-1}}, f_{N_{g-1}} \in \times_{i \in N_{g-1}} A_i$ be two profiles of actions of players of the bloc N_{g-1} such that there exists two Nash equilibria $z(e_{N_{g-1}})$ and $z(f_{N_{g-1}})$ respectively for games $G(e_{N_{g-1}})$ and $G(f_{N_{g-1}})$, with distinct payoff for each player of the block $N_g \setminus N_{g-1}$, where $G(e_{N_{g-1}})$ (respectively $G(f_{N_{g-1}})$) is a stage-game with players $N \setminus N_{g-1}$ obtained from G by fixing the actions of players of the block N_{g-1} to $e_{N_{g-1}}$ (respectively $f_{N_{g-1}}$). Define $c_g = \min_{i \in N_g \setminus N_{g-1}} |u(z(e_{N_{g-1}})) - u(z(f_{N_{g-1}}))|$. Let y^g denote alternating between the action profiles $z(e_{N_{g-1}})$ (in even periods) and $z(f_{N_{g-1}})$ (in odd periods). Let $z^{i,g}$ be the Nash equilibrium profile among $z(e_{N_{g-1}})$ and $z(f_{N_{g-1}})$ which is the worst for player $i \in N_g$.

The 5-phase strategy profile is adjusted as follows. The phase length variables - namely q (Phase 3), r (Phase 4), and $t_g(qp + rp)$ ($g = 1, 2, \dots, h$, Phases 2 and 5)- will be chosen at the end of the construction, along with the reward vectors x^j ($\forall j \in N_h$) used in Phase 4. Early (late) deviations are those occurring up to (after) period $T - t_h(qp + rp) - (qp + rp)$, ie deviation is ‘‘early’’ if there is still time to run Phases 3 and 4 before period $T - t_h(qp + rp) + 1$.

STRATEGY PROFILES.

1. (Main Path) Play a^l at period $t = l[\text{mod } p] + t_h(qp + rp)$ until period $T - t_h(qp + rp)$. [After an early deviation by $i \in N_h$, go to Phase 3; after a late deviation by $i \in N_{g'}$, go to Phase 5.] Go to Phase 2.
2. (Good Recursive Nash) For $g = h, \dots, 1$: Play y^g in periods $T - t_g(qp + rp) + 1, \dots, T - t_{g-1}(qp + rp)$. [After a deviation by $i \in N_{g'}$ with $g' < g$, start Phase 5.]
3. (Minmax Phase for i): Play w^i for qp periods. [If any $j \in N \setminus \mathcal{J}(i)$ deviates early, start Phase 4; if any $j \in N_{g'}$ deviates late, start Phase 5 with $i \leftarrow j$. If any $j \in \mathcal{J}(i)$ deviates early, set $i \leftarrow j$ and restart Phase 3.] Then set $j \leftarrow i$ and start Phase 4.
4. (Reward Phase) Repeat the path $\pi^{p,j}$ for r rounds. [If any $i \in N_h$ deviates early, restart Phase 3; if any $i \in N_{g'}$ deviates late, start Phase 5.] Then return to Phase 1.
5. (Bad Recursive Nash) Play $z^{g',i}$ until period $T - t_{g'}(qp + rp)$. [If $j \in N_{g''}$ deviates, where $g'' < g'$, set $g' \leftarrow g''$ and $i \leftarrow j$ and restart Phase 5.] Then go to Phase 2.

So along the equilibrium path, the sequence of action profiles is

$$\underbrace{a^1, \dots, a^p; a^1, \dots, a^p; \dots; a^1, \dots, a^p}_{T - t_h(qp + rp) \text{ periods}}; \underbrace{y^h, \dots, y^h}_{s_h(qp + rp) \text{ periods}}; \underbrace{y^{h-1}, \dots, y^{h-1}}_{s_{h-1}(qp + rp) \text{ periods}}; \dots; \underbrace{y^1, \dots, y^1}_{s_1(qp + rp) \text{ periods}}$$

Length of phases: Let ρ be the largest gap between best and worst payoffs across all players in G . For all $i \in N_h$, let $\pi^{p',i}$ be a sequence of length p' of pure Nash equilibria of the effective game such that the sequence of average payoffs $x^i = u(\pi^{p',i})$, $i \in N_h$ satisfies $x^i \gg 0 \forall i \in N_h$, $x_i^i < x_i^j \forall j \notin (N \setminus N_h) \cup \mathcal{J}(i)$, $x^i = x^j \forall j \in \mathcal{J}(i)$, and $x_i^i < y_i \forall i \in N_h$. (Such vectors exist following Abreu et al. (1994)). There is no loss of generality to assume that $p' = p$. Choose q such that $\rho < qp \cdot x_i^i$ and r such that $\rho + \max \{0, qp \cdot (y_j - u_j(w^i))\} < rp \cdot (x_i^i - x_i^j)$ for all $i \in N_h$ and $j \in N_h \setminus \mathcal{J}(i)$.

For any number k , let $\psi_g(k)$ be the least even number above $2k\rho/c_g$, so that that a player $i \in N_g$ is willing to play k periods of any action followed by $\psi_g(k)$ periods of y^g , if deviations switch each y^g to $z^{g,i}$. Recursively define

$$\begin{aligned} s_h(m) &= \psi_h(m) \text{ and } (\forall g = 1, 2, \dots, h-1) s_g(m) \\ &= \psi_g(m + s_{g+1}(m) + \dots + s_h(m)) \end{aligned} \quad (1)$$

Then set $t_0(m) = 0$ and $t_g(m) = s_1(m) + \dots + s_g(m)$, for $g = 1, 2, \dots, h$.

Subgame perfect verification: By construction, no one-shot deviation by a player of the block N_h is profitable (see Smith (1995), Demeze-Jouatsa and Wilson (2019)). Observe that only Nash equilibria of the effective game appears on equilibrium paths of each subgame. Therefore, no player of the block $N \setminus N_h$ can profitably deviate from the constructed strategy. \square

References

- Abreu D, Dutta PK, Smith L (1994) The folk theorem for repeated games: a NEU condition. *Econometrica* 62(4):939–948
- Benoit JP, Krishna V (1985) Finitely repeated games. *Econometrica* 53(1):905–922
- Benoit JP, Krishna V (1987) Nash equilibria of finitely repeated games. *Int J Game Theory* 16(3):197–204
- Benoit JP, Krishna V (1993) Renegotiation in finitely repeated games. *Econom J Econom Soc* 62:303–323
- Demeze-Jouatsa GH (2019) Essays on finitely repeated games (doctoral dissertation). Universität Bielefeld, Bielefeld
- Demeze-Jouatsa GH, Wilson A (2019) Comment on “smith (1995): Perfect finite horizon folk theorem”. *Econometrica*. <https://www.econometricsociety.org/content/comment-%E2%80%9Csmith-1995-perfect-finite-horizon-folk-theorem> (online)
- Fudenberg D, Maskin E (1986) The folk theorem in repeated games with discounting or with incomplete information. *Econometrica* 54(3):533–554
- Fudenberg D, Maskin E (1991) On the dispensability of public randomization in discounted repeated games. *J Econ Theory* 53(2):428–438
- Fudenberg D, Levine DK, Takahashi S (2007) Perfect public equilibrium when players are patient. *Games Econ Behav* 61(1):27–49
- Fudenberg D, Levine D, Maskin E (2009) The folk theorem with imperfect public information. In: A long-run collaboration on long-run games, World Scientific, pp 231–273
- González-Díaz J (2006) Finitely repeated games: a generalized Nash folk theorem. *Games Econ Behav* 55(1):100–111
- Gossner O (1995) The folk theorem for finitely repeated games with mixed strategies. *Int J Game Theory* 24(1):95–107
- Gossner O, Hörner J (2010) When is the lowest equilibrium payoff in a repeated game equal to the minmax payoff? *J Econ theory* 145(1):63–84
- Smith L (1995) Necessary and sufficient conditions for the perfect finite horizon folk theorem. *Econometrica* 63(2):425–430
- Wen Q (1994) The “folk theorem” for repeated games with complete information. *Econometrica* 62(4):949–954

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.