ESSAYS ON

FINITELY REPEATED GAMES

Ghislain-Herman DEMEZE-JOUATSA, M. Sc.

Universität Bielefeld

Submitted to the Faculty of Business Administration and Economics of Bielefeld University in partial fulfillment of the requirements for the doctoral degree Dr. rer. pol.

Bielefeld, January 2018

Essays on finitely repeated games

Dissertation zur Erlangung des Doktorgrades
der Wirtschaftswissenschaften (Dr. rer. pol.)
der Universität Bielefeld
vorgelegt von

Ghislain-Herman DEMEZE-JOUATSA, M. Sc.

Gutachter:
Prof. Dr. Frank Riedel
Prof. Christoph Kuzmics, PhD

Gedruckt auf alterungsbeständigem Papier nach DIN-ISO 9706

To My Family

To the DEMEZE

Acknowledgments

Firstly, I would like to express my gratitude to my advisors, Christoph Kuzmics and Frank Riedel for their patience, their motivation and unconditional support during my PhD.

I am also grateful to Olivier Gossner, Roland Pongou, Vitoria Barham, and Moyouwou Issofa who invited me to join their research teams where I had the chance to discuss my research with them and their colleagues.

I thank all the professors of the Economic faculty and specially those of the Bilefeld Graduate School of Economics and Management for their supports and for the organization of our weekly Economics seminars and more importantly for inviting famous speakers from all over the world to Bielefeld.

I thank my classmates and friends from the BiGSEM and IMW for all the nice moments we had together and for the multiple responses and comments I received from them. I also thank Michael Greinecker, Patrick Beissner, Igor Muraviev, Victor Fenou Nzengang and Tolulope Rhoda Fadina for their comments.

I am grateful to DAAD and to the Welcome Center of the University of Bielefeld for financing me during my PhD. I also thanks the BGTS for financing my research stays at the University of Ottawa and University of Yaoundé I, as well as IMW and BiGSEM for sponsoring me for international conferences. I thank Bernhard Eckwert, Nikolai Brandt, and Tobias Hellman for their support for my first international talk in Cardiff.

Contents

Chapter 1

Introduction

Economic situations involve interactions of agents with diverse interests. Normal form games appear as a good formal representation of such situations in case of conflicting interests. Game theorists have provided many solution concepts that try as much as possible to predict the behavior of agents involved in a given normal form game. Some solutions are the dominant strategy equilibrium, the Nash equilibrium due to Nash (1951), the perfect equilibrium due to Selten (1975) and the minimax due to Von Neumann, the correlated equilibrium and so on. The application of those concepts provides interesting previsions in economics but also in political sciences, biology and psychology.

In some other games as insurance agreements and Cournot duopoly, the solution concepts above mentioned fail to explain the behavior of agents as well as the achieved outcomes. Indeed, in such games, participants do no obey their short-term incentives and aim to optimize their long-run payoffs.

THE THEORY OF REPEATED GAMES give insights to understand and explain how rational and compelling can the behavior of agents who have engaged in a long-run relationship differ from those of agents who interact only once. The main message of this theory is that repetition facilitates cooperation: In a long-run relationship, an agent may abandon her short term interests and cooperate with her vis-a-vis because she fears future penalties or because she expects some future rewards. Trivial examples are the prisoners' dilemma and the Cournot oligopoly where participants achieve efficiency in the long-run.
A repeated game is obtained as a finite or an indefinite repetition of a given normal form game, the stage-game. A strategy of a player in a repeated game

is a contingent plan which specifies the action that player intends to choose in the first round of that repeated game as well as the action she intends to choose at any subsequent round given an observed history.

The most popular result of the theory of repeated games is the folk theorem. It characterizes the set of payoffs that agents can achieve when they repeatedly play a stage-game. The folk theorem is said to hold if the set of equilibrium payoffs of the repeated game includes any feasible and individually rational payoff of the stage-game. In the folk theorem, a feasible payoff is a convex sum of stage-game payoff vectors and, it is called individually rational if each player's entry is greater than or equal to her stage-game reservation value which is her stage-game minimax payoff.
Under different classes of assumptions, game theorists have proved the folk theorem for both finitely repeated games where the time horizon is finite and commonly known and infinitely repeated games where the time horizon is either uncertain or infinite, each model describing as much as possible some specific long-run relationships. Two models of repeated games may differ from the monitoring structure available to agents (perfect, public, private), the characteristics of agents involved in the game (players with bounded memories, sophisticated players, short-lived or long-lived players, patient players), the type of actions available (pure actions or mixed actions), the solution concepts (Nash equilibrium, subgame perfect Nash equilibrium, renegotiation-proofness), the information available to players about the game (complete or incomplete) and so on.

This thesis provides an analysis of finitely repeated games of complete information and perfect monitoring and aim to answer two questions.
1) The folk theorem for finitely repeated game hold under some necessary and sufficient conditions; see Benoit and Krishna (1984) for a sufficient condition and Smith (1995) for a necessary and sufficient condition. In the case that the folk theorem does not hold, what is the exact range of payoffs that players can achieve using Nash equilibria or subgame perfect Nash equilibria of the finitely repeated game?
2) Classic models of repeated games assume that players can employ only pure and mixed actions and, in the equilibrium, at each point of time and given any observed history, each player specify her pure action or the probability distribution her action for the next period will be issued from. This assumption is

in contrast with most of observed behaviors as people more often bargain and agree on what one could call an incomplete strategy, that is a specification of the collusive path to be followed but not the enforcing schemes. In a long-run relationship, players might not want to predetermine the action they will choose if some participants deviate from the collusive path. They might for instance think that the potential deviator might immune herself against the punishment schemes if she knows them in advance. Moreover, if the punishment scheme (off equilibrium path) is too severe, it might make it difficult to get an agreement as players could consider that they can make mistake and choose an unintended action, as Selten (1975) explains. Could the theory of repeated game explain why incomplete strategies are stable?

1.1 Contribution

Repeated game models considered in this thesis are of complete information and perfect monitoring. I implicitly assume that each player has the complete information about her preferences and strategies as well as those of her fellow players. Furthermore, at each point of time each player can observe the strategy played by her fellow players.

This thesis provides two contributions to the repeated games literature. The first contribution is a complete folk theorem. This theorem fully characterizes the set of payoffs that are approachable by means of pure strategy subgame perfect Nash equilibria of finite repetitions of an arbitrary normal form game. In contrast with the classic folk theorem which provides necessary and sufficient conditions on the stage-game which ensure that each feasible and individually rational payoff vector of the stage-game is approachable by means of subgame perfect Nash equilibria of the finitely repeated game, the complete folk theorem applies to any compact normal form game and provides a full characterization of the whole set of achievable payoffs.

The second contribution is a new model of finitely repeated game where players can employ imprecise probabilistic devices to conceal their intentions. The analysis of that model gives insights to understand why incomplete contracts where participants agree on a collusive path and do not specify any enforcing scheme are widespread and stable. This new model of finitely repeated games also allows to explain the emergence of cooperation without relaxing the assumptions the information structure available to players (see Kreps et al. (1982) and Kreps and Wilson (1982)), the perfection of the monitoring technology (see Abreu et al. (1990), Aumann et al. (1995)) and players' rationality (see Neyman (1985), Aumann and Sorin (1989)).

The findings of this thesis are presented in four chapters.

In Chapter 2 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 full characterization of the limit set, as the time horizon increases, of the set of pure strategy subgame perfect Nash equilibrium payoffs 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).

Smith (1995) presents a necessary and sufficient condition for the finite-horizon perfect folk theorem. In the proof of this result, the author constructs a family of five-phase strategy profiles to approach a feasible and individually rational payoff vector of the stage-game. In Chapter 3, I use a counter-example to show that these strategy profiles are not subgame perfect Nash equilibria of the discounted repeated game. Nevertheless, the characterization of attainable payoff vectors by Smith remains true. I provide an alternative proof.

In Chapter 4 I present a model of finitely repeated games in which players can strategically make use of objective ambiguity. In each round of a finite repetition of a given finite stage-game, in addition to the classic pure and mixed actions, players can employ objectively ambiguous actions by using imprecise probabilistic devices as Ellsberg urns to conceal their intentions. I follow Riedel and Sass (2014) and I call a Nash equilibrium of this extended stage-game an

Ellsberg equilibrium. The main finding is that, when each player has many continuation payoffs in Ellsberg actions, any feasible payoff vector of the original stage-game that dominates the mixed strategy maxmin payoff vector of the original stage-game is (ex-ante and ex-post) approachable by means of subgame perfect Ellsberg equilibrium strategies of the finitely repeated game with discounting. I prove that this condition is also necessary. The novelty of this model is that it allows to approach any equilibrium payoff by simple subgame perfect Nash equilibria, equilibria involving very few parameters. Furthermore, at the equilibrium players do not have to predetermine the profile of (possibly mixed) actions they will employ on punishment paths. Another finding is that adding an infinitesimal level of ambiguity to the classic model of finitely repeated games allows to explain the emergence of cooperation, even if the stage game has a unique mixed strategy Nash equilibrium.

In Chapter 5 I present a model of infinitely repeated game with complete information and perfect monitoring and where players are allowed to employ pure actions, mixed actions as well as an additional device which captures the willingness of a player to exercise her right to remain silent. I show that any feasible payoff that dominates the maxmin (modulo some players have equivalent utility functions) payoff vector is sustainable by means of pure strategy subgame perfect Nash equilibria of the infinitely repeated game with discounting.

1.2 A first example

In this section, use a four-player game to illustrate how the complete (perfect) folk theorem works. The chosen game does not satisfy neither the distinct Nash payoff condition of Benoit and Krishna (1984) nor the recursively distinct Nash payoffs of Smith (1995) so that the perfect folk theorem for finite time horizon does not hold. I show how to determine the exact range of payoff vectors that are approachable by means of pure strategies subgame perfect Nash equilibrium strategies of the finitely repeated game.

Consider the four-player game G whose payoff matrix is given by Table 1.1 and where the set of actions of players 1, 2, 3 and 4 are respectively given by $A_1 = \{a_1^1, a_1^2\}$, $A_2 = \{a_2^1, a_2^2, a_2^3\}$, $A_3 = \{a_3^1, a_3^2\}$ and $A_4 = \{a_4^1, a_4^2\}$.

		a_2^1, a_4^1	a_2^1, a_4^2	a_2^2, a_4^1	a_2^2, a_4^2	a_2^3, a_4^1	a_2^3, a_4^2
a_1^1	a_3^1	1 1 1 1	0 0 0 0	3 0 0 1	3 3 1 0	4 1 0 -1	0 0 0 0
	a_3^2	1 2 1 1	0 0 1 1	0 1 1 0	0 $\frac{1}{2}$ 0 1	0 0 -1 1	0 0 1 0
a_1^2	a_3^1	0 0 2 3	0 0 3 2	3 0 1 1	1 4 1 1	3 5 2 3	1 5 3 2
	a_3^2	0 0 3 2	0 3 2 3	0 0 0 0	1 2 1 1	0 0 3 2	0 0 2 3

Table 1.1: Payoff matrix of a game with an incomplete Nash decomposition and where players can achieve a partial level of cooperation in finite time.

This game admits two pure Nash equilibrium profiles $a^1 = (a_1^1, a_2^1, a_3^1, a_4^1)$ and $a^2 = (a_1^1, a_2^1, a_3^2, a_4^1)$. As any pure strategy subgame perfect Nash equilibrium play path of the finitely repeated game lasts with a phase where only pure Nash equilibrium action profiles of the stage-game are played, this phase will employ only the action profiles a^1 and a^2. As player 2 receives distinct payoffs at Nash equilibrium profiles a^1 and a^2, in a second to last phase of a pure strategy subgame perfect Nash equilibrium play path of the finitely repeated game, she is willing to conform to any sequence of pure actions of the stage-game given that (i) the last phase is long enough, (ii) the last phase pays her (in average) strictly more than her worst stage-game pure Nash equilibrium and (iii) the deviations during the second to last phase are punished by playing a^1 in every period of the last phase.

As player 1 (respectively player 3 and player 4) receives the same payoff at Nash equilibrium profiles a^1 and a^2 it is not possible to credibly punish her if she profitably deviates during the second to last phase. Therefore, player 1 (respectively player 3 and player 4) has to play a stage-game pure best response at any profile of actions played in the second to last phase. Consequently, for a pure action profile of the stage-game to be eligible for a second to last phase of a pure strategy subgame perfect Nash equilibrium of the finitely repeated game, it has to be a pure Nash equilibrium of the new stage-game G^1 that is obtained from the stage-game G by setting the utility function of player 2 equal to a constant, say γ; see Table 1.2.

The stage-game G^1 admits five pure Nash equilibrium profiles a^1, a^2, $a^3 = (a_1^1, a_2^1, a_3^2, a_4^2)$, $a^4 = (a_1^2, a_2^2, a_3^1, a_4^1)$ and $a^5 = (a_1^2, a_2^2, a_3^2, a_4^2)$ and only player 1 receives distinct payoffs at those profiles. Therefore, in a third to last phase of any pure strategy subgame perfect Nash equilibrium play path of any finite

		a_2^1		a_2^2		a_2^3	
		a_4^1	a_4^2	a_4^1	a_4^2	a_4^1	a_4^2
a_1^1	a_3^1	1 γ 1 1	0 γ 0 0	3 γ 0 1	3 γ 1 0	4 γ 0 -1	0 γ 0 0
	a_3^2	1 γ 1 1	0 γ 1 1	0 γ 1 0	0 γ 0 1	0 γ -1 1	0 γ 1 0
a_1^2	a_3^1	0 γ 2 3	0 γ 3 2	3 γ 1 1	1 γ 1 1	3 γ 2 3	1 γ 3 2
	a_3^2	0 γ 3 2	0 γ 2 3	0 γ 0 0	1 γ 1 1	0 γ 3 2	0 γ 2 3

Table 1.2: First transformation of the game G

repetition of the game G, player 1 is willing to conform to any sequence of play of action profiles of the stage-game G given that (j) the second to last phase is long enough, (jj) in the second to last phase she receives (in average) strictly more than her worst pure Nash equilibrium payoff in the stage-game G^1 and that (jjj) the deviations during the third to last phase are punished by playing a^3 in every period of the second to last phase. From (i), (ii) and (iii), player 2 will not find it profitable to deviate during the third to last phase of any pure strategy subgame perfect Nash equilibrium play path of any finite repetition of the game G. As player 3 (respectively player 4) receives the same payoff at any pure strategy Nash equilibrium of the game G^1, it is not possible to motivate her to stick to a path involving an action profiles where she is not playing a stage-game pure best response. The action profiles eligible for the third to last phase of a subgame perfect Nash equilibrium play path are therefore Nash equilibria of the game G^2 that is obtained from G by setting the utility functions of both players 1 and 2 equal to a constant, say γ; see Table 1.3.

		a_2^1		a_2^2		a_2^3	
		a_4^1	a_4^2	a_4^1	a_4^2	a_4^1	a_4^2
a_1^1	a_3^1	γ γ 1 1	γ γ 0 0	γ γ 0 1	γ γ 1 0	γ γ 0 -1	γ γ 0 0
	a_3^2	γ γ 1 1	γ γ 1 1	γ γ 1 0	γ γ 0 1	γ γ -1 1	γ γ 1 0
a_1^2	a_3^1	γ γ 2 3	γ γ 3 2	γ γ 1 1	γ γ 1 1	γ γ 2 3	γ γ 3 2
	a_3^2	γ γ 3 2	γ γ 2 3	γ γ 0 0	γ γ 1 1	γ γ 3 2	γ γ 2 3

Table 1.3: Payoff matrix of the game G^2

The game G^2 admits six pure Nash equilibrium profiles a^1, a^2, a^3, a^4, a^5 and $a^6 = (a_1^2, a_2^2, a_3^1, a_4^2)$ and a unique pure Nash equilibrium payoff vector $(\gamma, \gamma, 1, 1)$. Two remarks follow. Firstly, the Nash decomposition of the game G is incomplete. Secondly, any pure strategy subgame perfect Nash equilibrium play path has only three phases: A first phase that employs action profiles

a^1, a^2, a^3, a^4, a^5 and a^6 which are pure Nash equilibria of the game G^2; a second phase that employs action profiles a^1, a^2, a^3, a^4, a^5 which are pure Nash equilibria of the game G^1; and a third phase which employs action profiles a^1, a^2 which are pure Nash equilibria of the original stage-game, the game G. Therefore, the set of action profile eligible for subgame perfect Nash equilibrium play paths of finite repetitions of the stage-game G is restricted to $\{a^1, \cdots a^6\}$. It follows that any subgame perfect Nash equilibrium payoff vector of any finite repetition of the stage-game G has to be in the set of recursively feasible payoff vectors of the game G which is the convex hull of the set

$$\{(1,1,1,1),(1,2,1,1),(0,0,1,1),(3,0,1,1),(1,4,1,1),(1,2,1,1)\}.$$

Players 3 and 4 will therefore receive their unique stage-game pure Nash equilibrium payoff at any pure strategy subgame perfect Nash equilibrium of the finitely repeated game. Furthermore, within the set of eligible actions $\{a^1, \cdots, a^6\}$, player 2 can not be pushed down by her fellow players to a payoff that is strictly less than $\frac{1}{2}$. Indeed each pure strategy subgame perfect Nash equilibrium average payoff vector of the finitely repeated game weakly dominate the payoff vector $(0, \frac{1}{2}, 1, 1)$ which in turns weakly dominates the effective minimax payoff vector $(0,0,0,0)$ of the game G. I call the payoff vector $(0, \frac{1}{2}, 1, 1)$ the recursive effective minimax payoff vector.

The above reasoning teaches that the set of pure strategy subgame perfect Nash equilibrium payoff vectors of any finite repetition of the game G is included in the convex hull of the set

$$\{(\tfrac{1}{8}, \tfrac{1}{2}, 1, 1), (\tfrac{11}{4}, \tfrac{1}{2}, 1, 1), (1,4,1,1)\}$$

which is the set of recursively feasible payoff vectors that dominate the recursive effective minimax payoff vector. This set is a lower-dimension subset of the set of feasible and individually rational payoff vectors of the game G.

Theorem 1 in page 23 says that, 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 recursively feasible payoff vectors that dominate the recursive effective minimax payoff vector. In this example, this limit set equals the convex hull of the payoff set

$$\{(\tfrac{1}{8}, \tfrac{1}{2}, 1, 1), (\tfrac{11}{4}, \tfrac{1}{2}, 1, 1), (1,4,1,1)\}$$

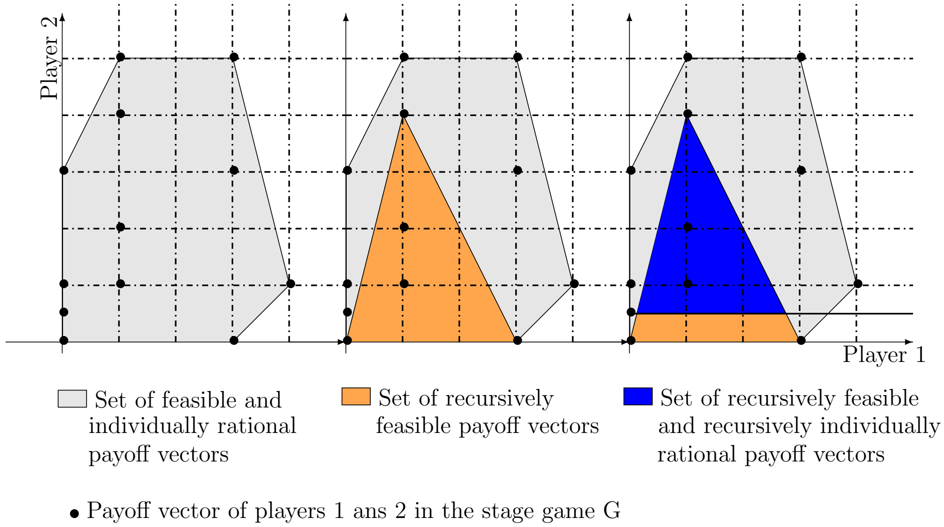

Figure 1: Equilibrium payoff vectors of players 1 and 2.

1.3 A second example

In this section, I use a four-player normal form game to illustrate that when the finite time horizon (Nash) folk theorem does not hold, not all stage game action profiles are eligible for pure strategy Nash play paths. I then show how to discriminate action profiles that are eligible for pure strategy Nash play paths. Those action profiles turn out to be Nash equilibria of the stage game or Nash equilibria of a degenerated game obtained from the stage game by making some of the players indifferent across their set of actions. I call a payoff Nash-feasible if it belongs to the convex hull of the set of payoffs to eligible action profiles. Theorem 4 in page 26 says that a payoff is approachable by means of pure strategy Nash equilibria of the finitely repeated game if and only if it is Nash-feasible and individually rational.

In the four-player normal form game G whose payoff matrix is described by Table 1.4 and where player 4 chooses the row of matrices (a_4^1 or a_4^2), player 3 chooses the column of matrices (a_3^1 or a_3^2), player 2 chooses the column (a_2^1 or a_2^2) and player 1 chooses the row (a_1^1 or a_1^2), the unique pure Nash equilibrium payoff is $(2, 1, 1, 1)$ and the pure minimax payoff of each player is equal to 1. As we will observe, as the time horizon increases, the set of pure strategy Nash equilibrium payoffs of the finite repetitions of the stage-game G converges to the convex hull of the set $\{(1, 1, 1, 1), (4, 1, 1, 1), (1, 4, 1, 1)\}$, a lower-dimension proper subset of the set of feasible and individually rational payoff vectors.

		a_3^1		a_3^2	
		a_2^1	a_2^2	a_2^1	a_2^2
a_4^1	a_1^1	2 1 1 1	1 1 0 4	1 0 0 1	0 1 4 0
	a_1^2	1 1 0 0	0 0 0 0	0 5 1 1	1 0 1 0
a_4^2	a_1^1	1 1 1 0	6 6 4 0	1 2 0 0	1 0 0 1
	a_1^2	0 1 0 1	5 0 1 1	0 2 1 0	1 1 1 1

Table 1.4: Payoff matrix of the game G

In a finite repetition of the game G, any pure strategy Nash play path will end with a phase where the unique stage-game pure Nash equilibrium profile $a^1 = (a_1^1, a_2^1, a_3^1, a_4^1)$ is repeatedly played. The average payoff of player 1 in that last phase equals 2 and is strictly greater than her pure minimax payoff which is equal to 1. Thus, in a second to last phase of a pure strategy Nash play path, player 1 is willing to conform to any sequence of non-Nash equilibrium action profiles given that the last phase is long enough and that deviations from an ongoing path are threaten by the grim trigger strategy profile, that is, after a unilateral deviation is observed, the author of the deviation is minimaxed so that she received at most her minimax payoff in each subsequent period of the repeated game. As players 2, 3 and 4 receive their pure minimax payoffs in the last phase of all pure strategy Nash play path, in the second to last phase, there is not a way to simultaneously make players 2, 3 and 4 play an action where they are not at their stage game best response. Therefore, the set of actions eligible for the second to last phase is the set of Nash equilibria of a new game G^{*1}, game obtain from G by setting the utility function of player 1 equal to a constant, let's say γ (see Table 1.5).

		a_3^1		a_3^2	
		a_2^1	a_2^2	a_2^1	a_2^2
a_4^1	a_1^1	γ 1 1 1	γ 1 0 4	γ 0 0 1	γ 1 4 0
	a_1^2	γ 1 0 0	γ 0 0 0	γ 5 1 1	γ 0 1 0
a_4^2	a_1^1	γ 1 1 0	γ 6 4 0	γ 2 0 0	γ 0 0 1
	a_1^2	γ 1 0 1	γ 0 1 1	γ 2 1 0	γ 1 1 1

Table 1.5: Payoff matrix of the game G^{*1}.

The game G^{*1} has two pure Nash equilibrium profiles, $a^1 = (a_1^1, a_2^1, a_3^1, a_4^1)$ and $a^2 = (a_1^2, a_2^1, a_3^2, a_4^1)$ and the associated payoff vectors in the original game G are respectively $(2,1,1,1)$ and $(0,5,1,1)$. The play path

$$(\underbrace{a^1, a^2, a^2}_{2^{nd}\text{ to last phase}}, \underbrace{a^1, a^1, a^1}_{\text{last phase}})$$

is an example of two-phase pure strategy Nash play path. At such two-phase play path, both of players 1 and 2 receive (8/6 for player 1 and 14/6 for player 2) strictly more than their pure minimax payoffs while player 3 and 4 receive their pure strategy minimax payoffs. Thus, in the third to last phase of a pure strategy Nash play path, players 3 and 4 need to be at their stage game best response at any action profile played whereas players 1 and 2 are willing to conform to any sequence of action profiles given that the two last phases are long enough and that deviations from an ongoing path are threaten by the grim trigger strategy. Action profiles eligible for the third to last phase are therefore Nash equilibria of the game G^{*2} obtained from G by setting the utility functions of both players 1 and 2 equal to the constant γ (see Table 1.6).

		a_3^1		a_3^2	
		a_2^1	a_2^2	a_2^1	a_2^2
a_4^1	a_1^1	$\gamma\ \gamma\ 1\ 1$	$\gamma\ \gamma\ 0\ 4$	$\gamma\ \gamma\ 0\ 1$	$\gamma\ \gamma\ 4\ 0$
	a_1^2	$\gamma\ \gamma\ 0\ 0$	$\gamma\ \gamma\ 0\ 0$	$\gamma\ \gamma\ 1\ 1$	$\gamma\ \gamma\ 1\ 0$
a_4^2	a_1^1	$\gamma\ \gamma\ 1\ 0$	$\gamma\ \gamma\ 4\ 0$	$\gamma\ \gamma\ 0\ 0$	$\gamma\ \gamma\ 0\ 1$
	a_1^2	$\gamma\ \gamma\ 0\ 1$	$\gamma\ \gamma\ 1\ 1$	$\gamma\ \gamma\ 1\ 0$	$\gamma\ \gamma\ 1\ 1$

Table 1.6: Payoff matrix of the game G^{*2}

The game G^{*2} has four pure Nash profiles: $a^1 = (a_1^1, a_2^1, a_3^1, a_4^1)$, $a^2 = (a_1^2, a_2^1, a_3^2, a_4^1)$, $a^3 = (a_1^2, a_2^2, a_3^1, a_4^2)$ and $a^4 = (a_1^2, a_2^2, a_3^2, a_4^2)$ and the associated payoff vectors in the original game G are respectively $(2,1,1,1)$, $(0,5,1,1)$, $(5,0,1,1)$ and $(1,1,1,1)$. An example of three-phase pure strategy Nash play path is

$$(\underbrace{a^3, a^3, a^2, a^4}_{3^{rd}\text{ to last phase}}, \underbrace{a^1, a^2, a^2}_{2^{nd}\text{ to last phase}}, \underbrace{a^1, a^1, a^1}_{\text{Last phase}}).$$

At such three-phase play path, both players 3 and 4 still receive their pure strategy minimax payoffs. Therefore, in the game G, there is no way to leverage the behavior of players 3 and 4. It follows that, a pure strategy Nash play path of the finite repetition of G will have at most three phases. A first phase where action profiles a^1, a^2, a^3 and a^4 are played; a second phase where actions a^1 and a^2 are played; and a third phase where only the action profile a^1 is played. At any occurrence of any other action profile, either player 3 or 4 will have incentive to deviate and there is no way to prevent such deviation. This reasoning suggests that, the set of feasible payoff vectors eligible for pure strategy Nash play paths is included in the convex hull of the set $\{(2,1,1,1),$ $(0,5,1,1)$, $(5,0,1,1), (1,1,1,1)\}$, the set of Nash-feasible payoff vectors.

Theorem 2 in page 25 says that, the set of pure strategy Nash equilibrium payoffs of finite repetitions of the game G converges to the set of Nash-feasible and individually rational payoff vectors.

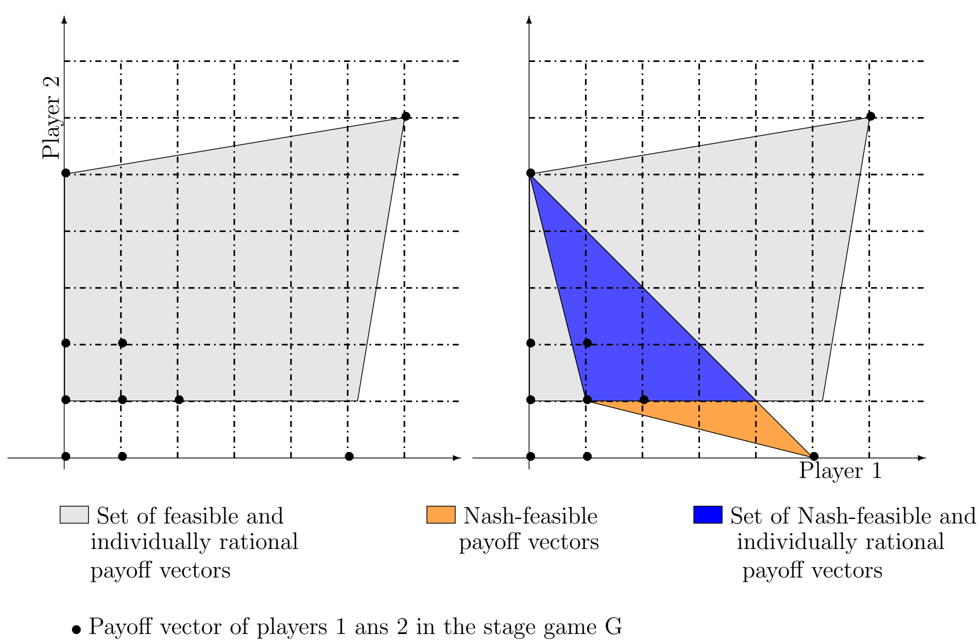

Figure 2: Equilibrium payoff vectors of players 1 and 2.

1.4 A third example: The benefit of the ambiguity

In this section I present an example of a game in which the classic model of finitely repeated games with pure and mixed strategies can not explain the emergence of cooperation while the introduction of an infinitesimal level of ambiguity in the model allows for sustaining cooperation.

Consider the three-player normal form game G whose payoff matrix is given by table 1.7 in which player 1 chooses the columns (L_1, RH_1 or RT_1), player 2 chooses the rows (H_2 or T_2) and player 3 chooses the matrix (L_3 or R_3). In this game, the strategy L_1 of player 1 is strictly dominated and therefore player 1 will play L_1 with probability 0 at any Nash equilibrium. Given that player 1 plays L_1 with probability 0, player 3 will find it strictly dominant to play R_3 with probability 1. The resulting restricted game is the well known 2×2 matching pennies game played by players 1 and 2, game that has a unique mixed strategy Nash equilibrium profile where player 1 plays RH_1 and RT_1 each with probability $\frac{1}{2}$ and player 2 plays H_2 and T_2 each with probability $\frac{1}{2}$. Consequently, the game G has a unique Nash equilibrium profile $s^* = \left(\{\frac{1}{2}RH_1 \oplus \frac{1}{2}RT_1\}, \{\frac{1}{2}H_2 \oplus \frac{1}{2}T_2\}, \{R_3\}\right)$ where player 1 plays L_1 with probability 0 and plays RH_1 and RT_1 with the same probability $\frac{1}{2}$, player 2 plays H_2 and T_2 with the same probability $\frac{1}{2}$ and player 3 plays R_3 with probability 1.

	L_1	RH_1	RT_1
H_2	-2 -2 4	6 6 0	6 6 0
T_2	-2 -2 4	6 6 0	6 6 0

L_3

	L_1	RH_1	RT_1
H_2	-2 -2 4	-1 1 1	1 -1 1
T_2	-2 -2 4	1 -1 2	-1 1 2

R_3

Table 1.7: Payoff matrix of the stage-game G.

As the game G has a unique Nash equilibrium in mixed strategy, any finite repetition of G in which players are allowed to employ only pure and mixed actions admits a unique subgame perfect equilibrium payoff which is $u(s^*) = (0, 0, \frac{3}{2})$ (see Benoit and Krishna (1984)). Now assume that players are ambiguity averse and are allowed to use sophisticated devices as Ellsberg urns to conceal their intentions. For all $\varepsilon_1, \varepsilon_2 \in [0, \frac{1}{2}]$, let

$$\overline{s}(\varepsilon_1, \varepsilon_2) = \left(\{\tfrac{1}{2}RH_1 \oplus \tfrac{1}{2}RT_1\}, \{pH_2 \oplus (1-p)T_2, \tfrac{1}{2} - \varepsilon_1 \leq p \leq \tfrac{1}{2} + \varepsilon_2\}, \{R_3\}\right)$$

be the profile in which player 1 plays L_1 with probability 0 and RH_1 and RT_1 with the same probability $\frac{1}{2}$, player 3 plays R_3 with probability 1 while player 2 issues her action from a device whose unique known property is that the probability to issue H_2 is between $\frac{1}{2} - \varepsilon_1$ and $\frac{1}{2} + \varepsilon_2$. At any profile $\overline{p} = \left(\{\frac{1}{2}RH_1 \oplus \frac{1}{2}RT_1\}, \{pH_2 \oplus (1-p)T_2\}, \{R_3\}\right)$ of probability distribution where $\frac{1}{2} - \varepsilon_1 \leq p \leq \frac{1}{2} + \varepsilon_2$, player 1 and player 2 receive each 0 while player 3 receives $2-p$. At the profile $\overline{s}(\varepsilon_1, \varepsilon_2)$, as player 3 is ambiguity averse and does not know the value of p, she ex-ante receives her worst expected payoff, that is $\frac{3}{2} - \varepsilon_2$. The ex-ante payoff to the profile $\overline{s}(\varepsilon_1, \varepsilon_2)$ is therefore $(0, 0, \frac{3}{2} - \varepsilon_2)$. Note that at the profile $\overline{s}(\varepsilon_1, \varepsilon_2)$, no ambiguity averse player can profitably deviate. Indeed, if player 1 plays L_1 with probability 0, then R_3 is a strictly dominant action of player 3. The expected payoff of player 2 is independent of her chosen action (possibly mixed) if player 1 plays RH_1 and RH_2 with the same probability $\frac{1}{2}$ and player 3 plays R_3 with probability 1. Furthermore, if player 3 plays R_3 with probability 1 and player 2 plays $\{pH_2 \oplus (1-p)T_2, \frac{1}{2} - \varepsilon_1 \leq p \leq \frac{1}{2} + \varepsilon_2\}$, the worst expected payoff of player 1 is maximal if she plays RH_1 and RT_1 with the same probability $\frac{1}{2}$.

At the equilibrium profile $\overline{s}(\varepsilon_1, \varepsilon_2)$, player 3 receives a payoff that is strictly less than her mixed Nash equilibrium payoff. Therefore, in the repeated game, she is willing to conform to a play of the pure action profile (RH_1, H_2, L_3) if it is followed by sufficiently many plays of the unique stage-game mixed Nash equilibrium s^* and deviations by player 3 are punished by switching each s^* to $\overline{s}(\varepsilon_1, \varepsilon_2)$. As players 1 and 2 play best responses at the profile (RH_1, H_2, L_3), the above described path and the associated mechanism constitute a subgame perfect equilibrium of the finitely repeated game. At that equilibrium, player 1 (as well as player 2) receives an average payoff that is strictly greater than her expected payoff at s^*. Thus, the behavior of players 1 and 2 can also credibly be leveraged near the end of the finitely repeated game. This allows to approximate collusive payoffs via subgame perfect equilibrium strategies of the finitely repeated game. For instance the Pareto superior payoff vector $(2, 2, 2)$ can be approximated by the following subgame perfect equilibrium strategy of the finitely repeated game.

1. For any $t \in \{0, \dots T_1\}$, play $s_1 = (LH_1, H_2, L_3)$ at time $2t$ and play $s_2 = (RL_1, L_2, L_3)$ at time $2t+1$.

2. For any $t \in \{2T_1 + 2, \dots, 2T_1 + 3 + \lceil \frac{2}{\varepsilon_1} \rceil\}$, play s^*.

3. If any player deviates, play $\overline{s}(\varepsilon_1, \varepsilon_2)$ till the end of the game.

As we observe in this example, when the classic model of finitely repeated games where players are allowed to employ only pure and mixed actions fail to explain the emergence of cooperation, allowing players to be objectively imprecise about the probability distribution they intend to used to issue their actions in each round of the finitely repeated game can allow to sustain cooperation. This observation still holds if players are allowed to used a relatively small level of ambiguity(that is if the upper bound of the level of imprecision of each player approaches zero). This is counter-intuitive as the set of stage-game actions with zero noises equals the set of mixed actions and, as in our example, the classic models of finitely repeated game with mixed actions predict no cooperation at all.

Chapter 2

A complete folk theorem for finitely repeated games

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.

2.1 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 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 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 (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.[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; 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

[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.

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?

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 two 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. 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.

2.2 Model and definitions

2.2.1 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.[2] 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.

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.

[2] 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 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 $\widetilde{I} = \{x = (x_1, \ldots, x_n) \in \mathbb{R} \mid x_i \geq \widetilde{\mu}_i$ for all $i \in N\}$ be the set of recursively individually rational payoff vectors.

2.2.2 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 $(\mathbb{R}^n, d)$.[3] 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 $d(x, Y) = \inf_{y \in Y} d(x, y)$.

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

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.

2.3 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 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 $\widetilde{m}^i$.[4] 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 $\mathbf{P}(i)$.

[4] 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.

2.4 Discussion and extension

2.4.1 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.

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).

2.4.3 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 $\underline{\delta}$.

2.4.4 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 2.2.1 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 2.4.1 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.

2.5 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.

2.6 Appendix 1: Proof of the Complete perfect folk theorem

2.6.1 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 $\text{ASPNE}(G)$ 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 $\text{ASPNE}(G)$ is compact and convex.*

Proof of Lemma 3.

The reader can check that $\text{ASPNE}(G)$ is a closed subset of the set of feasible payoff vectors which is compact. The set $\text{ASPNE}(G)$ is therefore compact. Since $\text{ASPNE}(G)$ is closed, its convexity holds if $z = \frac{1}{2}(x+y) \in \text{ASPNE}(G)$ for all $x, y \in \text{ASPNE}(G)$. Let $x, y \in \text{ASPNE}(G)$ 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 $a \in \text{Nash}(G)$ 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 $z \in \text{ASPNE}(G)$. ■

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

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).[5]

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

[5]The convergence in this lemma uses the Hausdorff distance. See Section 2.2.2.

a finite covering of ASPNE(G).[6] For all $l = 1, ..., L$ take T_0^l given by the definition of "$x^l \in$ ASPNE(G)" with $\frac{\varepsilon}{2}$.[7] 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 $T_\varepsilon = T_0$. ■

2.6.2 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

[6] $B(x, \varepsilon) = \{y \in \mathbb{R}^n \ / \ d(x, y) < \varepsilon\}$

[7] See Definition 4.

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 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.[8] 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 $\widetilde{F}$.*

2.6.3 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.

[8] 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.

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.

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 $\widetilde{\mu}$.[9]
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

[9] 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.

$u_i^{T-1}(\sigma_{|\pi_1(\sigma)})$. 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 $\widetilde{\mu}$. ■

2.6.4 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 2.2.1 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

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

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*

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

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

$$u_i^{r_p}[\theta^p(1, \alpha_{-i})] - u_i^{r_p}[\theta^p(0, \alpha_{-i})] \geq \phi, \tag{2.1}$$

$$u_i^{r_p}[\theta^p(\alpha)] - u_i^{r_p}[\theta^p(-1, \cdots, -1)] \geq \phi \tag{2.2}$$

and

$$|u_i^{r_p}[\theta^p(\alpha)] - u_i^{r_p}[\theta^p(\alpha_{\mathcal{J}(i)}, \alpha'_{N \backslash \mathcal{J}(i)})]| < \frac{1}{2^p}. \tag{2.3}$$

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).[10] For all $\alpha \in \{-1,0,1\}^n$ and $j \in N_h$, let

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

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

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

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).*

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

Proof of Lemma 10.

Let G be a compact normal form game. If G admits no pure Nash equilibrium, then $\widetilde{F} = \emptyset$ and $\widetilde{F} \cap \widetilde{I} \subseteq \text{ASPNE}(G)$. If G admits a unique pure Nash equilibrium payoff vector x, then $\widetilde{F} = \{x\} = \text{ASPNE}(G)$ and $\widetilde{F} \cap \widetilde{I} \subseteq \text{ASPNE}(G)$. Now suppose that G admits at least two distinct pure Nash equilibrium payoff vectors. Normalize the game such that the 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

$$\widetilde{m}^i \in arg\min_{a \in \widetilde{A}} \max_{j \in \mathcal{J}(i)} \max_{a'_i \in A_i} u_i(a'_i, a_{-i}).$$[11]

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

$$0 < q_1 u_i(\widetilde{m}^i) + q_2 r_1 u_i^{r_1}[\theta^1(1, \cdots, 1)] < \frac{q_1 + q_2 r_1}{2} y_i \tag{2.4}$$

and

$$-2\rho + \frac{q_1}{2} y_i > 0 \ for\ all\ i \in N_h. \tag{2.5}$$

Given q_1, q_2 and r_1, choose r such that

$$-2(q_1 + q_2 r_1)\rho + r\phi > 0. \tag{2.6}$$

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

$$\frac{q_2 r_1}{2} y_i - \frac{r}{2^{p_0}} > y_i - \frac{r}{2^{p_0}} > 0 \tag{2.7}$$

[11] Few comments on $\widetilde{m}^i$ are provided in footnote 4.

Apply the Lemma 9 to p_0 and obtain r_{p_0} and θ^{p_0}. Update $q_1 \leftarrow r_{p_0}q_1; q_2 \leftarrow r_{p_0}q_2r_1;\ r \leftarrow r_{p_0}r$. The parameters $\phi,\ \theta^1,\ q_1,\ q_2,\ r,\ r_1$ and θ^{p_0} are such that

$$0 < q_1u_i(\widetilde{m}^i) + q_2u_i^{r_1}[\theta^1(1,\cdots,1)] < \frac{q_1+q_2}{2}y_i \tag{2.8}$$

$$-2(q_1+q_2)\rho + r\phi > 0 \tag{2.9}$$

$$-2\rho + \frac{q_1+q_2}{2}y_i - \frac{r}{2^{p_0}} > 0 \tag{2.10}$$

and

$$y_i - \frac{r}{2^{p_0}} > 0\ for\ all\ i \in N_h. \tag{2.11}$$

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 **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 ϕ, θ^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).[12] 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**;

2. some payoff U_j till period sk;

[12] 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.

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 (2.1). 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 $a \in \widetilde{A}$.[13].

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

If after $l_1 k + l_2$ (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 $l_1 k + l_2$ periods of the Phase $\mathbf{P}(i)$.
2. $q_1 u_i(\widetilde{m}^i) + q_2 u_i^{r_1}[\theta^1(1, \cdots, 1)]$ in the new phase ***P(i)*** and the following ***SPE*** phase;
3. some payoff U_i till period sk;

[13] 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.

4. the payoff $ru_i^{r_{p_0}}[\theta^{p_0}(\alpha)] + \frac{r}{2^{p_0}}$ in the End-game.

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 $\mathbf{P(}i\mathbf{)}$ 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 $\mathbf{P(}i\mathbf{)}$, 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 $\mathbf{P(}i\mathbf{)}$ 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 (2.8). 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**;[14]
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**.

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

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 (2.3).

Early deviations from the main path are therefore deterred by inequality (2.10).

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 (2.2)), and $-2(q_1 + q_2)\rho + r\phi > 0$, no late deviation is profitable. This concludes the proof. ■

2.7 Appendix 2: Proof of the complete Nash folk theorem

2.7.1 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 2.6.1.

Let G be a compact normal form game and let $\mathrm{ANE}(G)$ 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 $\mathrm{NE}(T)$ 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* $\mathrm{ANE}(G)$ *is a compact and convex set.*

Proof of Lemma 11. It is immediate that $\mathrm{ANE}(G)$ 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 $\mathrm{ANE}(G)$ is also compact. The convexity of the set $\mathrm{ANE}(G)$ 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* $T > 0$, $\mathrm{NE}(T) \subseteq \mathrm{ANE}(G)$.

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* $\mathrm{ANE}(G)$.

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

2.7.2 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* $\mathrm{Nash}(G^{*h})$ *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 $\pi(\sigma)$.[15]
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. ■

[15] 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.

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.*

2.7.3 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 $T \geq 1$.

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 2.4.1 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.[16] 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.[17] Then there exists $T_0 > 0$*

[16]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.

[17]See Section 2.4.1 for the definition of the sequence $(N_l^*)_{l\geq 0}$.

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.[18]
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 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 $d(x, u^T(\sigma^{\varepsilon,x,T})) < \varepsilon$.

[18] 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 $\pi^{N_g^*}$.

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

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

and $x'_i > \mu_i$ for all $i \in N_h^*$.[19]

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 Nash(G^{*h}) such that

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

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

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

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 $T_{\varepsilon,x} = (s_{\varepsilon,x} + 1)q + kT_0$. ■

[19] 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.

2.8 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'. The proof of Lemma 1 is postponed to Section 4.5.2.

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 $G(\delta, r_p)$.[20] *Furthermore, for all $i \in N_h$ and $\alpha, \alpha' \in \{0,1\}^n$ and $\delta \in (\delta_p, 1)$, we have*

$$u_i^{r_p,\delta}[\theta^p(1, \alpha_{-i})] - u_i^{r_p,\delta}[\theta^p(0, \alpha_{-i})] \geq \phi, \tag{2.12}$$

$$u_i^{r_p,\delta}[\theta^p(\alpha)] - u_i^{r_p,\delta}[\theta^p(-1, \cdots, -1)] \geq \phi \tag{2.13}$$

and

$$|u_i^{r_p,\delta}[\theta^p(\alpha)] - u_i^{r_p,\delta}[\theta^p(\alpha_{\mathcal{J}(i)}, \alpha'_{N\setminus\mathcal{J}(i)})]| < \frac{1}{2^p}. \tag{2.14}$$

[20] 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 $\theta^p(\alpha)$.

Chapter 3

A note on "Necessary and sufficient conditions for the perfect finite horizon folk theorem" [Econometrica, 63 (2): 425-430, 1995.][1]

Abstract: Smith (1995) presented a necessary and sufficient condition for the finite-horizon perfect folk theorem. In the proof of this result, the author constructed a family of five-phase strategy profiles to approach feasible and individually rational payoff vectors of the stage-game. These strategy profiles are not subgame perfect Nash equilibria of the finitely repeated game. I illustrate this fact with a counter-example. However, the characterization of attainable payoff vectors by Smith remains true. I provide an alternative proof.

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

JEL classification: C72, C73.

[1] I am grateful to Christoph Kuzmics, Frank Riedel and Tim Hellmann for their helpful comments. Special thanks to Olivier Gossner who highlighted a mistake in an earlier version version of this note.

3.1 Introduction

Benoit and Krishna (1984) proved a finite-horizon perfect folk theorem under two sufficient conditions on the stage-game. The first condition is the full-dimensionality defined in Fudenberg and Maskin (1986). A stage-game meets the full-dimensionality condition if the dimension of the set of feasible payoff vectors equals the number of players. The second condition of Benoit and Krishna (1984) is that each player receives at least two distinct payoffs at stage-game Nash equilibria. Smith (1995) generalized the result of Benoit and Krishna (1984) and provided a necessary and sufficient condition for the finite-horizon perfect folk theorem. Smith's (1995) condition is that the stage-game has recursively distinct Nash payoffs. This basically means that there exists a time horizon T such that each player receives at least two distinct payoffs at subgame perfect Nash equilibria of the T-fold repeated game.

In the proof of this result, and under the assumption that the stage-game has recursively distinct Nash payoffs, Smith constructed a family of five-phase strategy profiles to approximate feasible payoff vectors that dominate the effective minimax payoff vector of the stage-game. These strategy profiles are not subgame perfect Nash equilibria of the finitely repeated game. I illustrate this fact with a counter-example. However, the characterization of attainable payoff vectors by Smith remains true. I provide an alternative proof.

This note is organized as follows. Section 3.2 provides a counter-example and discusses the failure of Smith's (1995) proof. Section 3.3 recalls the model and formally states the finite-horizon perfect folk theorem of Smith (1995) and Section 3.4 provides an alternative proof the later result.

3.2 The counter-example

3.2.1 The stage-game

Consider the three-player stage-game G whose payoff matrix is given in Table 3.1. In the game G, player 1 chooses lines (a_1^1 or a_1^2), player 2 chooses columns (a_2^1 or a_2^2) and player 3 chooses matrices (a_3^1 or a_3^2).

The pure action profiles (a_1^2, a_2^1, a_3^2) and (a_1^1, a_2^2, a_3^2) are Nash equilibria of the stage-game G and each player receives distinct payoffs at those action pro-

	a_3^1		a_3^2	
	a_2^1	a_2^2	a_2^1	a_2^2
a_1^1	0 0 0	2 2 0	2 2 2	3 3 3
a_1^2	0 0 0	1 1 0	2 2 1	2 2 2

Table 3.1: Payoff matrix of the stage-game G.

files. Therefore, this game has recursively distinct Nash payoffs, see Definition 6. Players 1 and 2 have the same utility function and are therefore equivalent.[2] The pure effective minimax payoff of player 1 (respectively player 2) equals 1 and is uniquely provided by the action profile $w^1 = w^2 = (a_1^2, a_2^2, a_3^1)$.[3]

The payoff vector $u = (\frac{3}{2}, \frac{3}{2}, \frac{3}{2})$ is feasible and strictly dominates the effective minimax payoff vector $\widetilde{\mu} = (1, 1, 1)$. The payoff vector u is therefore approachable by means of subgame perfect Nash equilibria of the finitely-repeated game with discounting; see Theorem 5.

In the proof of Theorem 5 of Smith (1995) which is stated in page 59 of this note, to approach the feasible payoff vector u, the author used a five-phase strategy. I recall it below and show that it is not a subgame perfect Nash equilibrium profile.

3.2.2 The five-phase strategy of Smith

The strategy profile used by Smith (1995) employs the concept of payoff asymmetry family that I briefly recall below.

The payoff asymmetry family

The concept of payoff asymmetry family is introduced by Abreu et al. (1994). Such a family allows to suitably reward effective punishers after a punishment phase. In our example, the payoff vectors $x^1 = x^2 = (1.3, 1.3, 1.3)$ and $x^3 = (1.4, 1.4, 1.2)$ form a payoff asymmetry family relatively to u. Indeed, the

[2] Player i is equivalent to player j in the game G if the utility function of player i is a positive affine transformation of the utility function of player j.

[3] The mixed effective minimax payoff of both players 1 and 2 also equals 1 and is uniquely provided by the pure action profile w^1.

payoff family $\{x^1, x^2, x^3\}$ meets the following requirements:

(A1) $x^i >> \widetilde{\mu}$ for all $i \in \{1, 2, 3\}$, [strict individually rationality]

(A2) $x_i^i < u_i$ for all $i \in \{1, 2, 3\}$, [target payoff domination]

and

(A3′) $x_i^i < x_i^j$ for all $i, j \in \{1, 2, 3\}$, $i \nsim j$.[4] [payoff asymmetry]

I should notice that (A3′) is an adjusted version of the original requirement (3) in Abreu et al. (1994) where the game meets the NEU (non-equivalent utility) property.

Length of phases

Let β^i be the best payoff vector of player i in the game G.
Let ω^i be worst payoff vector of player i in the game G.

Choose q such that for all i, $\omega_i^i + qx_i^i > \beta_i^i + 1$. Take $q = 4$.

Given q, choose r such that for all j with $j \nsim i$,

$q\omega_j^j + rx_j^i > \beta_j^j + rx_j^j + (q-1)u_j + 1.$ Take $r = 86$.

Take $t_h(q + r) = 3(q + r)$.

Smith's strategy

Let a be the outcome of a public randomization device that has an average payoff of $u = (\frac{3}{2}, \frac{3}{2}, \frac{3}{2})$.

Let $T \geq t_h(r+q)$ and σ be the strategy profile of the finitely-repeated game $G(T)$ described by the following five phases.

1. MAIN PATH: Play a until period $T - t_h(r + q)$. [If any player i deviates

[4] The notation $i \nsim j$ means that player i is not equivalent to player j in the game G.

early, start 3; if some player deviates late, start 5.[5]]

2. GOOD RECURSIVE PHASE: Play the stage-game Nash equilibrium profile (a_1^1, a_2^2, a_3^2) till the end of the finitely-repeated game $G(T)$.

3. MINIMAX PHASE: Play w^i for q periods. [If player j (with $j \nsim i$) deviates, start 4.] Set $j \leftarrow i$.

4. REWARD PHASE: Play x^j for r periods. [If i deviates early, restart 3; if some player deviates late, start 5.]

5. BAD RECURSIVE NASH PHASE: Play the stage-game Nash equilibrium (a_1^2, a_2^1, a_3^2) until the end of the game.

A profitable deviation from σ

For all $k \geq 0$, let $T(k) = k + r + q + t_h(r+q)$. Let σ_1' be a strategy of player 1 in which player 1 deviates from a in the first period of the repeated game as well as at the beginning of each REWARD PHASE and plays her stage-game best response in each period of the MINIMAX PHASE. This deviation is profitable for large k. Indeed, if player 1 does not deviate from σ, she gets at most an expected payoff of $A(k) = \frac{1}{T(k)}\{\beta^1 + \frac{3(k+r+q-1)}{2} + 3t_h(r+q)\}$.

If she deviates and plays σ_1', she gets at least $B(k) = \frac{1}{T(k)}\{2(k - \lceil\frac{k-1}{q+1}\rceil - 2)\}$ where $\lceil\frac{k-1}{q+1}\rceil$ is the smallest integer greater than or equal to $\frac{k-1}{q+1}$.

As k goes to ∞, $A(k)$ goes to $\frac{3}{2}$ and $B(k)$ goes to $\frac{8}{5}$.

This means that for sufficiently long time horizon T and sufficiently high discount factor δ, the strategy profile σ is not a Nash equilibrium of the finitely-repeated game $G(\delta, T)$ and therefore not a subgame perfect Nash equilibrium of $G(\delta, T)$.

[5] A deviation is called late if it occurs during the final $q + r + t_h(r+q)$ periods of the repeated game; all others are called early deviations.

3.2.3 Intuition behind the failure of Smith's proof

Denote by w^i the profile of stage-game mixed actions at which player i receives her effective minimax payoff.[6]

If the utility function of player i in the stage-game G is equivalent to that of another player, say player j, then the effective minimax payoff of player i might be strictly greater than her minimax payoff and player i might even have a strict incentive to deviate from w^i. Indeed, it might be the case that only player j plays a stage-game best response at the profile w^i. In that case, it is not convenient to use the five-phase strategy profile of Smith (1995) to approximate a payoff vector in which player i receives strictly less than her best response payoff at w^i.

Indeed, during the third phase of the five-phase strategy of Smith (1995), player i is minimaxed using the stage-game action profile w^i where she might not be at a stage-game best response. In addition, during this phase, deviations by any player who is equivalent to player i (including player i) are ignored. As in the counter-example above, player i might find it profitable to deviate from an ongoing path (either from the MAIN PATH or from the REWARD PHASE) to push her fellow players to start the MINIMAX PHASE where she is punished.

This failure of is not minor in the sense that for any specification of the action profile to be used in the MINIMAX PHASE where $i = 1$, at least one player will find it strictly profitable to deviate from the five-phase strategy of Smith (1995).

Denote a MINIMAX PHASE where $i = 1$ by MP(1).

Indeed, if for a given specification $\overline{w}^1$ of the stage-game profile to be repeatedly played in the phase MP(1) the strategy profile σ is a subgame perfect Nash equilibrium of the finitely-repeated game $G(T)$, then at $\overline{w}^1$ player 3 has to play a_3^1 with strictly positive probability. Otherwise the punishment payoff

[6] The strategy profile defined in page 63 has a slightly different interpretation. Indeed at that profile, a player whose utility function is equivalent to that of player i might have incentive to deviate. If she does so, she receives at most her stage-game effective minimax payoff.

of player 1 in the minimax phase MP(1) will be strictly greater than player 1's entry in the target payoff vector $u = (\frac{3}{2}, \frac{3}{2}, \frac{3}{2})$. Given the choice of player 3 in $\overline{w}^1$, player 2 has to play a_2^2 with probability 1 at the profile $\overline{w}^1$. Otherwise she will find it strictly profitable to deviate from σ and repeatedly play her best response at $\overline{w}^1$ during the phase MP(1), as she will not be punished if she does so. Given that player 2 plays a_2^2 with probability 1 in the profile $\overline{w}^1$, player 1 has to play a_1^1 with probability 1 in the profile $\overline{w}^1$. Otherwise she will find it strictly profitable to deviate and to play a_1^1 with probability 1 in each round of the phase MP(1). Therefore, only convex sums of payoff vectors $(2, 2, 0)$ and $(3, 3, 3)$ are possible payoff to the profile of actions $\overline{w}^1$. This implies that player 1 receives an average payoff greater than or equal to 2 in each round of the minimax phase MP(1), which is strictly greater than her entry in the target equilibrium payoff u. This contrasts the idea of punishment behind a minimax phase, which is to deter deviations. A player should not find it profitable to start a minimax phase.

The above reasoning teaches that the incentives of any player who is not at her stage-game best response at the profile $\overline{w}^1$ have to be controlled during a minimax phase. Note that this reasoning is not possible in case the stage-game meets the NEU (non-equivalent utility) property of Abreu et al. (1994) or the full dimensionality property of Fudenberg and Maskin (1986). Under those conditions, no player is equivalent to another and therefore any stage-game profile at which player i plays a stage-game best response and receives her minimax payoff is suitable for a minimax phase, see Benoit and Krishna (1984) and Smith (1993) for the finite-horizon perfect folk theorem under those properties.

The methods of Benoit and Krishna (1984) and Smith (1993) do not easily extend to games where some players have equivalent utility functions. But still, the finite-horizon perfect folk theorem for games that possibly violate the NEU condition as stated in Smith (1995) holds. This note provides a clear proof.

In the next section I recall Smith's (1995) model and state his finite-horizon perfect folk theorem. I provide the proof of the latter theorem in Section 3.4.

3.3 Smith's model

3.3.1 The stage-game

Let $G = \langle A_i, \pi_i; i = 1, ..., n\rangle$ be a finite normal form n-player game, where A_i is player i's finite set of actions, and $\pi_i : A = \times_{i=1}^{n} A_i \to \mathbb{R}$ is player i's utility function. Let M_i be player i's mixed action set and let $M = \times_{i=1}^{n} M_i$. For any profile of actions $a \in A$, set $\pi(a) = (\pi_1(a), \ldots, \pi_n(a))$. For any profile of mixed actions $\mu = (\mu_1, \ldots, \mu_n) \in M$, denote by $\pi(\mu) = (\pi_1(\mu), \ldots, \pi_n(\mu))$ the vector of expected payoffs of players.

Let $\mathcal{J} = \{1, \ldots, n\}$ be the set of players. Let $\mathcal{J}(i)$ be the set of players whose utility function is a positive affine transformation of π_i. Let

$$\widetilde{\mu}_i = \min_{\mu \in M} \max_{j \in \mathcal{J}(i)} \max_{\mu'_j} \pi_i(\mu'_j, \mu_{-j})$$

be the effective minimax payoff of player i. Normalize the utilities functions of players such that $\widetilde{\mu}_i = 0$ for all i. Let $F = co\{\pi(\mu) :\ \mu \in M\}$ be convex hull of the set of expected payoff vectors. Let $F^* = \{u \in F : u_i > 0,\ \text{for all } i\}$ be the feasible and strictly rational set.

Given a subset of players $\mathcal{J}' = \{j_1, \ldots, j_m\} \subset \mathcal{J}$ and their mixed actions profile

$$a_{\mathcal{J}'} = (a_{j_1}, \ldots, a_{j_m}) \in M_{j_1} \times M_{j_2} \times \cdots \times M_{j_m} \equiv M_{\mathcal{J}'}, \tag{3.1}$$

let $G(a_{\mathcal{J}'})$ be the induced $(n-m)-$player game for players $\mathcal{J}\backslash\mathcal{J}'$ obtained from G when the actions of players $\mathcal{J}'$ are fixed to $a_{\mathcal{J}'}$.

Define a Nash decomposition of the game G as an increasing sequence of $h \geq 0$ nonempty subset of players from $\mathcal{J}$, namely

$$\{\varnothing = \mathcal{J}_0 \subsetneq \mathcal{J}_1 \subsetneq \cdots \subsetneq \mathcal{J}_h \subseteq \mathcal{J}\}, \tag{3.2}$$

so that for $g = 1, \ldots, h$, actions $e_{\mathcal{J}_{g-1}}, f_{\mathcal{J}_{g-1}} \in M_{\mathcal{J}_{g-1}}$ exist with a pair of Nash payoff vectors $y(e_{\mathcal{J}_{g-1}})$ of $G(e_{\mathcal{J}_{g-1}})$ and $y(f_{\mathcal{J}_{g-1}})$ of $G(f_{\mathcal{J}_{g-1}})$ different exactly for players in $\mathcal{J}_g\backslash\mathcal{J}_{g-1}$, ie

$$y(e_{\mathcal{J}_{g-1}})_i \neq y(f_{\mathcal{J}_{g-1}})_i \tag{3.3}$$

for all $i \in \mathcal{J}_g\backslash\mathcal{J}_{g-1}$.

Definition 6 *The game G has recursively distinct Nash payoffs if there is a Nash decomposition with $\mathcal{J}_h = \mathcal{J}$.*

3.3.2 The finitely-repeated game

Let $G(\delta, T)$ be the $T-$fold repeated game in which players discount the futur with the parameter $\delta < 1$. Smith (1995) assumed that the monitoring structure is perfect so that each player can condition her current action on the past actions of all players.

A strategy behavioral strategy of player i in the repeated game $G(\delta, T)$ is a $T-$tuple $\alpha_i = (\alpha_{i1}, \ldots, \alpha_{i,T})$ where for all $t \in \{1, \ldots, T\}$ and past history $h^t \in A^{t-1}$ (with $A^0 = \varnothing$), $\alpha_{it}(h^t)$ is the (possibly mixed) action that player i intends to play at time t if she observes h^t. The objective function of player i in the finitely-repeated game $G(\delta, T)$ is the expected discounted sum of her stage-game payoffs:

$$\pi^{\delta}_{iT}(\alpha) := \tfrac{1-\delta}{1-\delta^T} \textstyle\sum_{t=1}^{T} \delta^{t-1} \pi_{it}(\alpha)$$

where $\pi_{it}(\alpha)$ is player i's expected payoff at period t with the strategy profile $\alpha = (\alpha_1, \ldots, \alpha_n)$. The strategy profile α is a Nash equilibrium of the finitely-repeated game $G(\delta, T)$ if for all player i, α_i maximizes the objective function $\pi^{\delta}_{iT}(\cdot, \alpha_{-i})$ of player i.

The strategy profile α is a subgame perfect Nash equilibrium of the finitely-repeated game $G(\delta, T)$ if after any history h^t, the restriction $\alpha_{|h^t}$ of α to the history h^t is a Nash equilibrium of the remaining game.

Let

$$V(\delta, T) = \{\pi^{\delta}_{T}(\alpha) = (\pi^{\delta}_{1T}(\alpha), \cdots, \pi^{\delta}_{nT}(\alpha)) \mid \alpha \text{ is a subgame perfect Nash equilibrium of } G(\delta, T)\}$$

be the set of subgame perfect Nash equilibrium payoff vectors of the finitely-repeated game $G(\delta, T)$.

Theorem 5 (See Smith (1995)) *Suppose that the stage-game G has recursively distinct Nash payoffs. Then for the finitely-repeated game $G(\delta, T)$, $\forall u \in F^*$ and $\forall \varepsilon > 0$, $\exists T_0 < \infty$ and $\delta_0 < 1$ so that $T \geq T_0$ and $\delta \in [\delta_0, 1] \Rightarrow \exists v \in V(\delta, T)$ with $\|u - v\| < \varepsilon$.*

3.4 A proof of Smith's folk theorem

I follow Smith (1995) and assume that players condition their choices on the outcome of a publicly observed exogenous continuous random variable. For simplicity, I also assume that the discount factor equals 1. The later assumption is without loss of generality as it does not change the incentives of players if those are strict.

The main ingredient of the proof of Theorem 5 is a multi-level reward path function whose existence is guaranteed by the recursively distinct Nash payoffs condition, see Lemma 18. The multi-level reward path function allows to independently leverage the behavior of players near the end of the finitely-repeated game, no matter if there are or not players who have equivalent utility functions. In addition, and backwardly, this multi-level reward path function allows to leverage the behavior of players at any stage of the finitely-repeated game.

Gossner (1995) used similar method to prove a finite-horizon perfect folk theorem with unobservable mixed strategies. The advantage of Lemma 18 is that it does not require the dimension of the set of feasible payoff vectors to equal the number of players neither each player to have at least two distinct payoffs at Nash equilibria of the stage-game.

Denote by $G(T)$ the T-fold finitely repeated game $G(\delta, T)$ where the discount factor δ equals 1. In the game $G(T)$, the utility of player i at the behavioral strategy α is

$$\pi_i^T(\alpha) := \lim_{\delta \to 1} \pi_{iT}^{\delta}(\alpha)$$

which is equal to the payoff average $\frac{1}{T}\sum_{t=1}^{T} \pi_{iT}(\alpha)$. Let

$$V(1,T) := \{\pi^T(\alpha) = (\pi_1^T(\alpha), \cdots, \pi_n^T(\alpha)) \mid \alpha \text{ is a SPNE of } G(T)\}$$

be the set of subgame perfect Nash equilibrium payoff vectors of the finitely repeated game $G(T)$ and let

$$AP = \{u \in F \mid \forall \varepsilon > 0, \exists T_0 < \infty \text{ so that } T \geq T_0 \;\Rightarrow\; \exists v \in V(1,T) \text{ with } \|u - v\| < \varepsilon\}$$

be the set of feasible payoff vectors that are approachable by means of subgame perfect Nash equilibria of finite repetitions of the stage-game G. To prove Theorem 5, we will show that $F^* \subseteq AP$.

Lemma 18 *Suppose that the stage-game G has recursively distinct Nash payoffs. Then there exists $\phi > 0$ such that for all $p \geq 0$, there exists $r_p > 0$ and*

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

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

$$\pi_i^{r_p}[\theta^p(1, \gamma_{-i})] - \pi_i^{r_p}[\theta^p(0, \gamma_{-i})] \geq \phi \tag{3.4}$$

$$\pi_i^{r_p}[\theta^p(\gamma)] - \pi_i^{r_p}[\theta^p(-1, \cdots, -1)] \geq \phi \tag{3.5}$$

$$|\pi_i^{r_p}[\theta^p(\gamma)] - \pi_i^{r_p}[\theta^p(\gamma_{\mathcal{J}(i)}, \gamma'_{\mathcal{J}\setminus\mathcal{J}(i)})]| < \frac{1}{2^p}. \tag{3.6}$$

This lemma says that, if the stage-game G has recursively distinct Nash payoffs, then we can (almost-) independently leverage the behavior of each player near the end of the game. This lemma also allows to construct credible punishment schemes and to approximate any feasible payoff vector that dominates the effective minimax payoff vector by means of SPNE of the finitely repeated game.

Assume that the finitely repeated game will last with a reward phase where players are rewarded with respect to their behavior in the earlier stage of the repeated game, that players are informed that the reward path to be used is a SPNE path $\theta^p(\gamma)$ of the repeated game $G(r_p)$. Furthermore, assume that γ is initialized to the value $(1, \cdots, 1)$ and that each player has the possibility to update her entry in the vector γ each time where a player whose utility function is not equivalent to her deviates. Inequality (3.4) says that, given the profile γ_{-i} of players of the block $\mathcal{J}\setminus\{i\}$, player i strictly prefers the path $\theta^p(1, \gamma_{-i})$ to the path $\theta^p(0, \gamma_{-i})$. Inequality (3.6) ensures that the incentives of players of different equivalence classes are almost independent for sufficiently large p. The strategy constructed in the proof of Theorem 5 does not allow a player to strategically improve her payoff by giving to players whose utilities' function are equivalent to her a chance to update their entries in the vector γ.

Consider for instance the stage-game whose payoff matrix is given by Table 3.1. In that game, player 1 and player 2 have the same utility function and are therefore equivalent. Figure 3 below displays the relative position of the payoff vectors $\pi^{r_p}[\theta^p(\gamma)]$ where $\gamma \in \{0,1\}^3 \cup \{(-1,-1,-1)\}$. The path $\theta^p(-1,-1,-1)$ will allow to deter deviations that occurs near the end of the game.

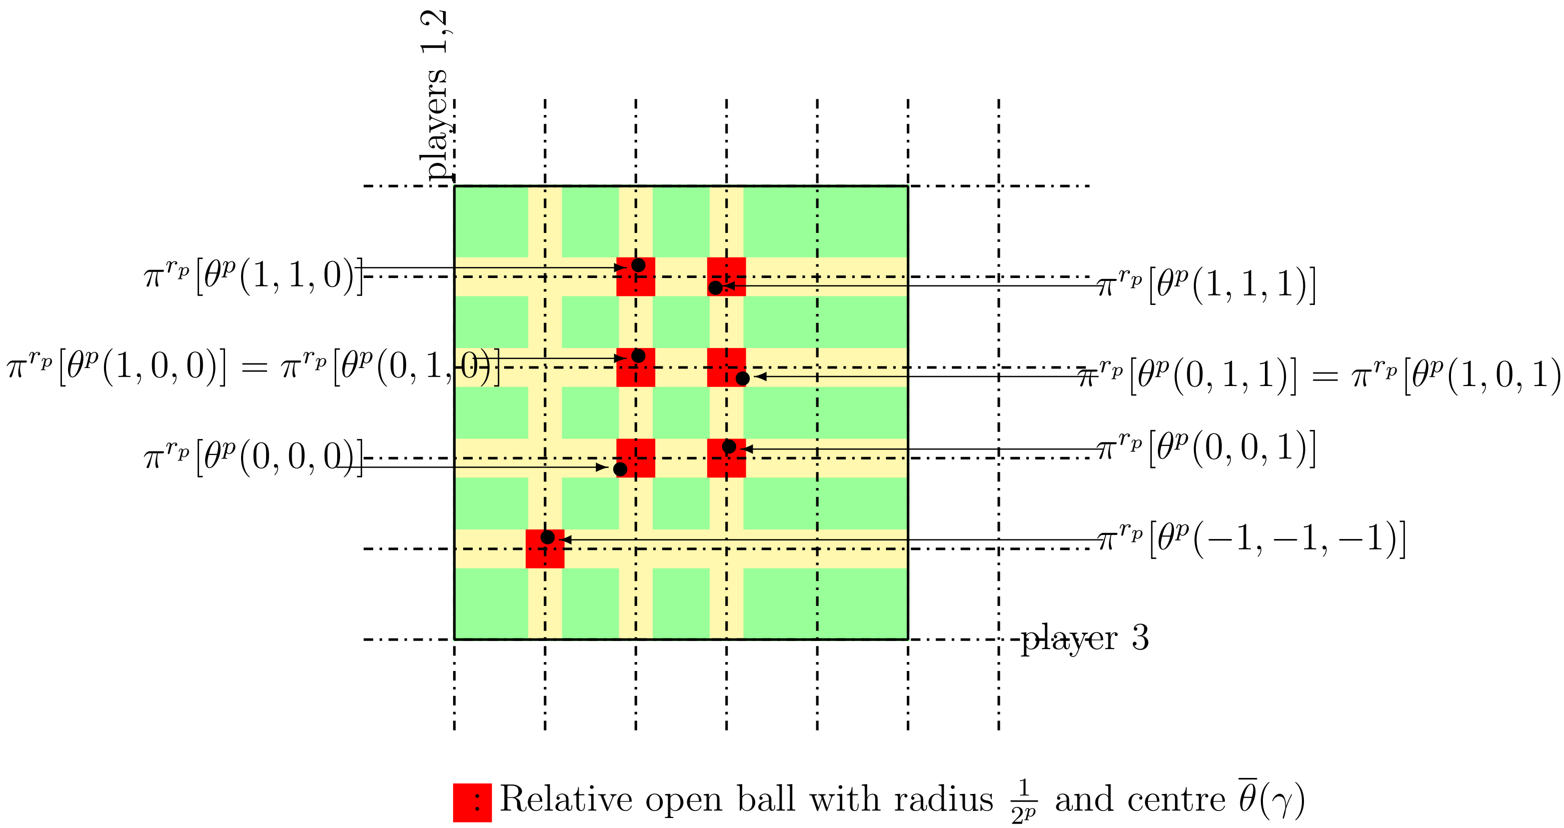

Figure 3: An example of relative position of the payoff vectors $\pi^{r_p}[\theta^p(\gamma)]$.

A detailed proof of Lemma 18 is presented in Section 3.5.

Proof of Smith's (1995) folk theorem.

Let u be a feasible payoff vector that lies in the relative interior of F^*, and let a be the outcome of a public randomization device that has an expected payoff vector of u.

Obtain ϕ, r_1 and θ^1 with $p = 1$ from the Lemma 18. Let $q_1 > 0$ and $q_2 > 0$ such that

$$0 < q_1 \pi_i(w^i) + q_2 r_1 \pi_i^{r_1}[\theta^1(1, \cdots, 1)] < \frac{q_1 + q_2 r_1}{2} u_i \tag{3.7}$$

and

$$-2\rho + \frac{q_1}{2}u_i > 0 \ for \ all \ i \in N. \tag{3.8}$$

Given q_1, q_2 and r_1, choose r such that

$$-2(q_1 + q_2 r_1)\rho + r\phi > 0. \tag{3.9}$$

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

$$\frac{q_2 r_1}{2}u_i - \frac{r}{2^{p_0}} > u_i - \frac{r}{2^{p_0}} > 0 \tag{3.10}$$

Apply the Lemma 18 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 quantities $\phi,\ \theta^1,\ q_1,\ q_2,\ r,\ r_1$ and θ^{p_0} are such that

$$0 < q_1 \pi_i(w^i) + q_2 \pi_i^{r_1}[\theta^1(1, \cdots, 1)] < \frac{q_1 + q_2}{2}u_i \tag{3.11}$$

$$-2(q_1 + q_2)\rho + r\phi > 0 \tag{3.12}$$

$$-2\rho + \frac{q_1 + q_2}{2}u_i - \frac{r}{2^{p_0}} > 0 \tag{3.13}$$

and

$$u_i - \frac{r}{2^{p_0}} > 0 \ for \ all \ i \in N. \tag{3.14}$$

The $T-$period equilibrium outcome sequence is

$$a, \ldots, a; \theta^{p_0}(1, \cdots, 1)$$

where a is played for $T - r$ periods and the path $\theta^{p_0}(1, \cdots, 1)$ is of length r. Now I describe the subgame perfect Nash equilibrium σ of the finitely-repeated game that supports the equilibrium path. For all $i \in \mathcal{J}$, let w^i be a stage-game action profile such that

$$\max_{j \in \mathcal{J}(i)} \max_{m_j \in M_j} u_i(m_j, w^i_{-j}) = 0.$$

At the action profile w^i, no player of the block $\mathcal{J}(i)$ has to be at a best response. Playing a best response to the action profile w^i, a player of the block $\mathcal{J}(i)$ receives at most her effective minimax payoff.

Set $\gamma = (1, \cdots, 1)$. From now on, call a deviation late if it occurs during the final $q_1 + q_2 + r$ periods of the finitely-repeated game $G(T)$; all others are called early deviations. The strategy profile σ involves 5 phases and can be graphed as follows:

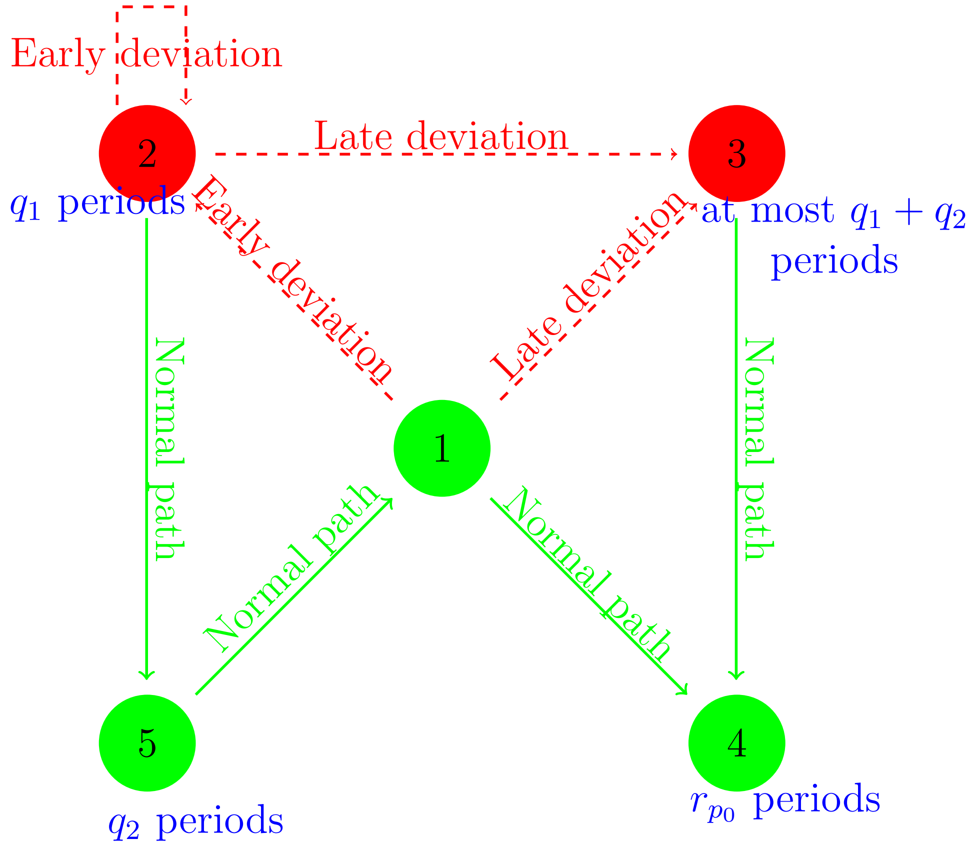

1. MAIN PATH: Play a until period $T - r$. [If any player i deviates early, start 2; if some player deviates late, start 3.] Go to 4.

2. MINIMAX PHASE $P(i)$: During this phase, each player $j \in \mathcal{J}(i)$ has to play her action w_j^i while each player of the block $N \backslash \mathcal{J}(i)$ can play whatever action she wants. This phase last for q_1 periods. [If any player $j \in \mathcal{J}(i)$ deviates early, restart 2.; if any player $j \in \mathcal{J}(i)$ deviates late, start 3.]
At the end of this phase, for all player $j \notin \mathcal{J}(i)$, set $\gamma_j = 0$ if there is at least one period of the MINIMAX PHASE where player j played an action different to w_j^i and set $\gamma_j = 1$ otherwise. Go to 5.

3. LATE DEVIATION: Each player can play whatever action she wants till period $T - r$. At period $T - r$, set $\gamma = (-1, \cdots, -1)$. Go to 4.

4. END OF THE GAME: Follow $\frac{r}{r_{p_0}}$ times a SPNE that supports the equilibrium path $\theta^{p_0}(\gamma)$.

5. SPE PHASE: Follow $\frac{q_2}{r_1}$ times the SPNE of the game $G(r_1)$ whose play path is $\theta^1(1, \cdots, 1)$. Go back to 1.

For sufficiently large time horizon T, the strategy profile σ is a subgame perfect Nash equilibrium of the finitely repeated game $G(T)$

In the following, call a player $j \notin \mathcal{J}(i)$ effective punisher if $\gamma_j = 1$ at the end of the MINIMAX PHASE $P(i)$. I prove the following:

A) It is strictly dominant for any player $j \notin \mathcal{J}(i)$ to be effective punisher during a MINIMAX PHASE $P(i)$

B) No early deviation from the MINIMAX PHASE is profitable

C) No early deviation from the MAIN PATH is profitable

D) No late deviation is profitable

A) It is strictly dominant to be effective punisher during a MINIMAX PHASE

If player $j \notin \mathcal{J}(i)$ is effective punisher, she receives at least:

1. $-(q_1 + q_2)\rho$ during the MINIMAX PHASE and the SPE PHASE;
2. some payoff U_j till period $T - r$;
3. $r \cdot \pi_i^{r_{p_0}}[\theta^{p_0}(1, \gamma_{-j})]$ in the last r periods of the repeated game.

In total she receives at least $-(q_1 + q_2)\rho + U_j + r \cdot \pi_i^{r_{p_0}}[\theta^{p_0}(1, \gamma_{-j})]$.

If player j is not effective punisher, she receives at most:

1. $(q_1 + q_2)\rho$ during the MINIMAX PHASE and the SPE PHASE;
2. the same payoff U_j till period $T - r$;
3. $r \cdot \pi_i^{r_{p_0}}[\theta^{p_0}(0, \gamma_{-j})]$ in the last r periods of the repeated game.

In total $(q_1 + q_2)\rho + U_j + r \cdot \pi_i^{r_{p_0}}[\theta^{p_0}(0, \gamma_{-j})]$ which is less than or equal to $(q_1 + q_2)\rho + U_j + r \cdot \pi_i^{r_{p_0}}[\theta^{p_0}(1, \gamma_{-j})] - r\phi$, see inequality (3.4). As $-2(q_1 + q_2)\rho + r\phi > 0$, it is strictly dominant for any player $j \notin \mathcal{J}(i)$ to be effective punisher.

B) No early deviation from the MINIMAX PHASE is profitable

If player $i \in \mathcal{J}(i)$ deviates early from the MINIMAX PHASE, she receives at most:

1. 0 in the deviation period;
2. $q_1\pi_i(w^i) + q_2\pi_i^{r_1}[\theta^1(1,\cdots,1)]$ in the new MINIMAX PHASE and the following SPE PHASE;
3. some payoff U_i till the end of the game.

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

1. $q_1\pi_i(w^i) + q_2\pi_i^{r_1}[\theta^1(1,\cdots,1)] + u_i$ till the end of the SPE PHASE;
2. the payoff $U_i - \frac{r}{2^{p_0}}$ till the end of the game.

As $u_i - \frac{r}{2^{p_0}} > 0$, no early deviation from the MINIMAX PHASE is profitable.

C) No early deviation from the MAIN PATH is profitable

If player i deviates early from the MAIN PATH, she receives at most:

1. ρ in the deviation period;
2. $q_1\pi_i(w^i) + q_2\pi_i^{r_1}[\theta^1(1,\cdots,1)]$ in the MINIMAX PHASE and the SPE PHASE;
3. some payoff U_i till period $T-r$;
4. $r \cdot \pi_i^{r_{p_0}}[\theta^{p_0}(\gamma_{\mathcal{J}(i)},1,\cdots,1)]$ in phase 4.

In total $\rho + q_1\pi_i(w^i) + q_2\pi_i^{r_1}[\theta^1(1,\cdots,1)] + U_i + r\cdot\pi_i^{r_{p_0}}[\theta^{p_0}(\gamma_{\mathcal{J}(i)},1,\cdots,1)]$ which is strictly less than $\rho + \frac{q_1+q_2}{2}u_i + U_i + r\cdot\pi_i^{r_{p_0}}[\theta^{p_0}(\gamma_{\mathcal{J}(i)},1,\cdots,1)]$, see inequality (3.11).

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

1. $-\rho + (q_1+q_2)\cdot u_i$ till the end of the SPE PHASE;
2. the same payoff U_i till period $T-r$;
3. $r\cdot\pi_i^{r_{p_0}}[\theta^{p_0}(\gamma)]$ in phase 4.

In total $-\rho + (q_1 + q_2) \cdot u_i + U_i + r \cdot \pi_i^{r_{p_0}}[\theta^{p_0}(\gamma)]$ which is strictly greater than $-\rho + (q_1 + q_2) \cdot u_i + U_i + r \cdot \pi_i^{r_{p_0}}[\theta^{p_0}(\gamma_{\mathcal{J}(i)}, 1, \cdots, 1)] - \frac{r}{2^{p_0}}$, see inequality (3.6).

As $-2\rho + \frac{q_1+q_2}{2} u_i - \frac{r}{2^{p_0}} > 0$, no early deviation from the MAIN PATH is profitable.

D) No late deviation is profitable

If from an ongoing path (MAIN PATH or MINIMAX PHASE) player i deviates late, then she receives at most:

1. $(q_1 + q_2)\rho$ till the beginning of the END OF THE GAME;
2. $r \cdot \pi_i^{r_{p_0}}[\theta^{p_0}(-1, \cdots, -1)]$ in the END OF THE GAME.

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

1. $-(q_1 + q_2)\rho$ till the beginning of the END OF THE GAME;
2. $r \cdot \pi_i^{r_{p_0}}[\theta^{p_0}(\gamma)]$ in the END OF THE GAME, where $\gamma \in \{0, 1\}^n$.

As $r \cdot \pi_i^{r_{p_0}}[\theta^{p_0}(\gamma)] \geq r \cdot \pi_i^{r_{p_0}}[\theta^{p_0}(-1, \cdots, -1)] + r\phi$ and $r\phi > 2(q_1 + q_2)\rho$, no late deviation is profitable. This concludes the proof. ■

3.5 Proof of intermediate results

In this section I proceed to the proof of Lemma 18. I first show that under the recursively distinct Nash payoffs condition, each player has many continuation equilibrium payoffs, which is necessary for the construction of our multi-level reward path function.

Lemma 19 *Suppose that the stage-game G has recursively distinct Nash payoffs. Then there exists $T_0 > 0$ such that for all $T > T_0$, each player receives at least two distinct payoffs at SPNE of $G(T)$.*

Proof of Lemma 19.

Let $\{\varnothing = \mathcal{J}_0 \subsetneq \mathcal{J}_1 \subsetneq \cdots \subsetneq \mathcal{J}_h = \mathcal{J}\}$ be the Nash decomposition of G.
I show by induction that for all $l \leq h$ there exists $T_{0,l} > 0$ such that each player of $\mathcal{J}_l$ receives distinct payoffs at SPNE of $G(T)$ for all $T > T_{0,l}$.
Players of the block $\mathcal{J}_1$ receive distinct payoffs at Nash equilibria of G. Therefore, the property holds for $l = 1$. Let $l < h$ such that $T_{0,l}$ exists. Let $i \in \mathcal{J}_{l+1} \backslash \mathcal{J}_l$, and let μ be a Nash equilibrium profile of $G(\mu_{\mathcal{J}_l})$ in which player

i receives a payoff that is different to her unique stage-game Nash equilibrium payoff. Let η_1 and η_2 be two SPNE play paths of $G(T_{0,l}+1)$ where each player of $\mathcal{J}_l$ receives distinct payoffs. The path $\eta^i = (\mu, \eta_1, \eta_2 \cdots, \eta_1, \eta_2)$ is a SPNE play path. At η^i, player i receives a payoff that is different to her stage-game Nash equilibrium payoff which is also a SPNE payoff. The conjunction lemma (see Benoit and Krishna (1984)) guarantee the existence of $T_{0,l+1}$. ■

Proof of Lemma 18.

The set AP is non-empty and convex and therefore has a relative interior point x. Let $\phi > 0$ such that the relative ball $\widetilde{B}(x, 5\phi n)$ is included in AP. For all $\gamma \in \{-1, 0, 1\}^n$ and $j \in \mathcal{J}$, let

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

and

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

For all $\gamma \in \{0, 1\}^n$, we have

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

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

and

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

Furthermore, since each player receives distinct payoffs within the set AP and players within the same equivalence class $\mathcal{J}(i)$ have equal entry at the payoff vector $\theta(\gamma)$, we have that

$$\theta(\gamma) \in \widetilde{B}(x, 5\phi n) \subseteq AP.$$[7]

Let $p \geq 0$ and $\varepsilon = \frac{1}{2}\min\{\phi, \frac{1}{2^p}\}$. For each $\gamma \in \{0, 1\}^n \cup \{(-1, \cdots, -1)\}$, there exists $T_{0\gamma p} < \infty$ so that for all $T \geq T_{0\gamma p}$, there exists $\alpha^{\gamma p}$ a subgame perfect Nash equilibrium of the repeated game $G(T)$ such that $\|\pi_T(\alpha^{\gamma p}) - \theta(\gamma)\| < \varepsilon$.

Take $r_p = \max\{T_{0\gamma p} \mid \gamma \in \{0, 1\}^n \cup \{(-1, \cdots, -1)\}\}$. For all $\gamma \in \{0, 1\}^n \cup \{(-1, \cdots, -1)\}$ and $p \geq 0$, let $\theta^p(\gamma)$ be the SPNE play path generated by the SPNE $\alpha^{\gamma p}$ of the repeated game $G(r_p)$. ■

[7] Indeed, from Lemma 19, each player receives distinct payoffs at subgame perfect Nash equilibria of finite repetitions of the stage-game G and, as corollary of the conjunction lemma (see Lemma 3.2 in Benoit and Krishna (1984)), each subgame perfect Nash equilibrium payoff vector of a finite repetition of the stage-game G with no discount belongs to AP.

Chapter 4

Repetition and cooperation: A model of finitely repeated games with objective ambiguity

Abstract: In this paper, we present a model of finitely repeated games in which players can strategically make use of objective ambiguity. In each round of a finite repetition of a finite stage-game, in addition to the classic pure and mixed actions, players can employ objectively ambiguous actions by using imprecise probabilistic devices such as Ellsberg urns to conceal their intentions. We find that adding an infinitesimal level of ambiguity can be enough to approximate collusive payoffs via subgame perfect equilibrium strategies of the finitely repeated game. Our main theorem states that if each player has many continuation equilibrium payoffs in ambiguous actions, any feasible payoff vector of the original stage-game that dominates the mixed strategy maxmin payoff vector is (ex-ante and ex-post) approachable by means of subgame perfect equilibrium strategies of the finitely repeated game with discounting. Our condition is also necessary.

Key words: Objective Ambiguity, Ambiguity Aversion, Finitely Repeated Games, Subgame Perfect Equilibrium, Ellsberg Urns, Ellsberg Strategies.

JEL classification: C72, C73, D81

4.1 Introduction

Contrary to the predictions of early models of repeated games with complete information and perfect monitoring which state that any finite repetition of a stage-game with a unique Nash equilibrium admits a unique subgame perfect Nash equilibrium payoff (see Benoit and Krishna (1984), Gossner (1995), Smith (1995)), the experimental evidence suggests at least a partial level of cooperation (see Kruse et al. (1994) and Sibly and Tisdell (2017)). This paper presents a new model of finitely repeated games with complete information and perfect monitoring that allows for an explanation of the emergence of cooperation in a larger class of normal form games. This class includes some stage-games with a unique Nash equilibrium.

The inconsistency of the predictions of the classic model of finitely repeated games with complete information and perfect monitoring with empirical evidence is subject to an extensive discussion and has led game theorists to relax their assumptions on the information structure available to players (see Kreps et al. (1982) and Kreps and Wilson (1982)), the perfection of the monitoring technology (see Abreu et al. (1990), Aumann et al. (1995)) and players' rationality (see Neyman (1985), Aumann and Sorin (1989)). However, the type of actions available to players also matters.

How well do pure and mixed actions capture the intentions of players involved in a dynamic game?

Greenberg (2000) argues that in a dynamic game, a player might want to exercise her right to remain silent. In the rock-paper-scissors game, a player might want to play "rock" with probability 0. These intentions are not captured by a pure or a mixed action, but rather by a set of lotteries over the set of the player's actions.

The strategies used in the proofs of the folk theorems to sustain equilibrium payoffs involve some punishment phases in which potential deviators are punished. In such phases, the player being punished responds to the punishment scheme settled by her fellow players, which is usually a minimax profile. In daily life, it is not always clear how precise a player would be when specifying what she intends to do in the event that her fellow player deviates from an agreement. An illustration of this situation can be found in incomplete contracts in which participants agree on the collusive paths to follow but are silent (totally ambiguous) about the enforcing mechanisms. In such cases, players might think that the deviator herself might be immune to the punishment scheme if she is aware of it in advance. Such behavior is not well-captured by pure or mixed strategies of the classic models of repeated games.

This paper presents a model of finitely repeated games with complete information and perfect monitoring in which players are allowed to use objectively ambiguous actions. In each period of the repeated game, in addition to the classic pure and mixed actions, players can employ objectively ambiguous actions by concealing their intentions in imprecise probabilistic devices, such as Ellsberg urns. I follow the work of Riedel and Sass (2014) and Riedel (2017) in referring to such imprecise action as an Ellsberg action. An Ellsberg action of a player can be thought of as a compact and convex set of probability distributions over the set of pure actions of that player. As in the related literature on ambiguity in games (see Riedel and Sass (2014), Riedel (2017), Greenberg (2000), Gilboa and Schmeidler (1989) and Ellsberg (1961)), I assume that players are ambiguity-averse and aim to maximize the worst payoff they expect to receive.

The main finding of this paper is that our model of finitely repeated games can explain the emergence of cooperation where the classic model with pure and mixed strategies fails to do so. We provide an example game to illustrate the idea that adding an infinitesimal level of ambiguity can be enough to approximate collusive payoffs via subgame perfect equilibrium strategies of the finitely repeated game. The main theorem states that if each player has many continuation equilibrium payoffs in Ellsberg actions, any feasible payoff vector that dominates the mixed strategy effective maxmin payoff vector is (ex-ante and ex-post) approachable by means of subgame perfect equilibrium strategies of the finitely repeated game with discounting. The existence of multiple continuation equilibrium payoffs in Ellsberg actions for each player is also a necessary condition for cooperation to arise in the finite horizon.

Earlier models of finitely repeated games assumed that players could employ only pure or mixed actions. Benoit and Krishna (1984), Benoit and Krishna (1987), and Smith (1995) provided conditions on the stage game that ensures that the set of equilibrium payoffs of the finitely repeated game includes any feasible payoff that dominates the minimax payoff vector. Gossner (1995) analyzed finitely repeated games in which players are allowed to use mixed actions, but do randomize privately.

Kreps et al. (1982) analyzed finite repetitions of the prisoners' dilemma and showed that the incompleteness of the information on players' options could generate a significant level of cooperation, and Kreps and Wilson (1982) showed that adding a small amount of incomplete information about players' payoffs could give rise to a reputation effect and therefore allow the monopolist to earn a relatively high payoff in finite repetitions of the Selten's chain-store game.

Neyman (1985) proved that in presence of complete information and perfect monitoring, utility-maximizing players can achieve cooperative payoffs in finite repetitions of the prisoners' dilemma given that there is a bound on the complexity of strategies available to them. Aumann and Sorin (1989) studied two-person games with common interests and demonstrated that if each player ascribes a positive probability to the event that her fellow player has a bounded recall, cooperative outcomes can be approximated by pure strategy equilibria.

Mailath et al. (2002) studied examples of finitely repeated games with imperfect public monitoring and illustrated that less informative signals about players' actions can allow for approximate Pareto superior payoffs by means of perfect equilibria of the repeated game, even if the stage game has a unique Nash equilibrium payoff. Sekiguchi (2002) studied the imperfect private monitoring case and provided a characterization of the stage-game whose finite repetitions admit non-trivial equilibrium outcomes.

The remainder of this paper is organized as follows: Section 2 presents the model as well as some preliminary results. The main theorem of the paper is presented and discussed in Section 3. Section 4 provides the proofs.

4.2 The Model

4.2.1 The stage-game

The initial stage-game

I represent a finite normal form game G by (N, S, u) where for all $i \in N$, S_i is the set of pure actions of player i. Both the set of players $N = \{1, ..., n\}$ and the set $S = \times_{i \in N} S_i$ of actions are finite. The utility of player i given $s = (s_1, ..., s_n) \in S$ is measured by $u_i(s)$. A mixed strategy of player $i \in N$ is a probability distribution p_i over the set S_i. Let ΔS_i be the set of mixed strategies of player i. We will abusively denote by $\Delta S = \Delta S_1 \times ... \times \Delta S_n$ the set of profiles of mixed strategies. At the profile $p = (p_1, .., p_n) \in \Delta S$, player i receives the expected payoff $u_i(p) = \sum_{s \in S} p(s) u_i(s)$ where for all $s \in S$, $p(s) = \prod_{i \in N} p_i(s_i)$, $p_i(s_i)$ being the probability that player i assigns to the action s_i according to the distribution p_i. For any $p = (p_1, ..., p_n) \in \Delta S$, $i \in N$ and $p'_i \in \Delta S_i$, (p'_i, p_{-i}) denotes the strategy profile in which all players except i behave the same as in p and the choice of i is p'_i. A profile of mixed strategy $p \in \Delta S$ is a **Nash equilibrium** of G ($p \in Nash(G)$) if for all $i \in N$ and $p'_i \in \Delta S_i$, $u_i(p'_i, p_{-i}) \leq u_i(p)$.

The payoff vector $x = (x_1, ..., x_n)$ is a **feasible** vector of the game G if it belongs to the convex hull of the set of payoff vectors of the game G. That is, if there exists a sequence $(\lambda_l)_{1 \leq l \leq L}$ of positive real numbers and a sequence $(a^l)_{1 \leq l \leq L}$ of pure actions' profile such that $\Sigma_{l=1}^{L} \lambda_l = 1$ and $x = \Sigma_{l=1}^{L} \lambda_l u(a^l)$. For all players $i, j \in N$, player i is equivalent to player j if there exists two real numbers β_{ij} and $\alpha_{ij} > 0$ such that $u_i(s) = \alpha_{ij} u_j(s) + \beta_{ij}$ for all $s \in S$. Denote by $\mathcal{J}(i)$ the set of players that are equivalent to player i. Let

$$\mu_i = \min_{p \in \Delta S} \max_{j \in \mathcal{J}(i)} \max_{p'_j \in \Delta S_j} u_i(p_{-j}, p'_j) = u_i(m^i)$$

be the **mixed strategy effective minimax payoff**[1] of player i and $\mu = (\mu_1, ..., \mu_n)$ be the effective minimax payoff vector of the game G. Let

$$\nu_i = \max_{j \in \mathcal{J}(i)} \max_{p_j \in \Delta S_j} \min_{p_{-j} \in \times_{k \neq j} \Delta S_k} u_i(p_{-j}, p_j)$$

be the **mixed strategy effective maxmin payoff** of player i and $\nu = (\nu_1, ..., \nu_n)$ be the effective maxmin payoff vector of the game G. Let V^* be

[1] The effective minimax has been introduced by Wen (1994). The effective minimax payoff of a player is her reservation value in the stage-game and equals her minimax payoff if she is not equivalent to any other player.

the set of feasible payoff vectors that strictly dominates the effective maxmin payoff vector ν.

The Ellsberg game

To ease the presentation of our results, we consider a very simple model of Ellsberg game and where players employ only reduced strategies. Riedel and Sass (2014) and Riedel (2017) provide a general model.

Let $G = (N, S, u)$ be a finite normal form game. An **Ellsberg strategy** P_i of player i is a compact set of probability distributions over the set S_i. Let $\mathcal{P}_i = \{P_i \subseteq \Delta S_i \mid P_i \text{ is compact}\}$ be the set of Ellsberg strategies of player i and $\mathcal{P}$ be the set of Ellsberg strategy profiles. Given a profile $P = (P_1, ..., P_n) \in \mathcal{P}$, the utility of player i is given by $u_i(P) = \min_{p \in P} u_i(p)$. The 3-tuple $(N, \mathcal{P}, u)$ is the Ellsberg extension of the game G. For any $P \in \mathcal{P}$, $i \in N$ and $P'_i \in \mathcal{P}_i$, (P'_i, P_{-i}) denotes the Ellsberg strategy profile in which all players except i behave the same as in P and the choice of i is P'_i. A profile of Ellsberg strategy $P \in \mathcal{P}$ is an **Ellsberg equilibrium** of G ($P \in E(G)$) if for all player $i \in N$ *and* Ellsberg strategy $P'_i \in \mathcal{P}_i$ of player i, $u_i(P'_i, P_{-i}) \leq u_i(P)$.

Priliminary results on the Ellsberg game

In the Ellsberg game, players have richer set of strategies and can even exercise their right to remain silent (totally ambiguous). Remaining silent can be more severe than employing a mixed strategy minimax profile. More importantly, remaining silent is an optimal punishment strategy profile in the Ellsberg game. We have the following lemma.

Lemma 20 *In the Ellsberg game, remaining silent is an optimal punishment strategy.*

Proof. of Lemma 20. Let $j \in N$ and $P_{-j} \in \mathcal{P}_{-j}$, be an Ellsberg profile of players of the block $-j$. We have $P_{-j} \subseteq \times_{k \neq j} \Delta S_k$ and therefore

$$u_i(\times_{k \neq j} \Delta S_k, P_j) \leq u_i(P_{-j}, P_j).$$

This means that, in the Ellsberg game, to punish an ambiguity averse players, it is optimal for her opponents to remain silent. ■

Intuitively, if on a punishment path all punishers exercise their right to remain silent, then, the target player, if she is ambiguity averse, will play a prudent strategy and will ex-ante receive at most her mixed strategy maxmin payoff. To illustrate how severe such punishment scheme can be, in comparison to the classic mixed strategy minimax, consider the three-player game whose payoff matrix is given by Table 4.1.

	c	d	c	d
a	$0 \quad 0 \quad 0$	$1 \quad -1 \quad 1$	$1 \quad 1 \quad -1$	$-1 \quad 1 \quad 1$
b	$-1 \quad 1 \quad 1$	$0 \quad 0 \quad -1$	$1 \quad -1 \quad 0$	$0 \quad 0 \quad 0$
	e		f	

Table 4.1: Payoff matrix of a three-player game where the use of Ellsberg strategies allow for severe punishment schemes.

In this game, player 1 chooses the rows (a or b), player 2 chooses the columns (c or d), and player 3 chooses the matrices (e or f). If only mixed strategies are allowed, each player can ensure herself the payoff 0. This is not possible under ambiguity. Indeed, under ambiguity, no player can ensure herself a payoff strictly greater than $-\frac{1}{2}$.

Suppose that player 2 plays c, and that player 3 plays e. Player 1 best responds playing a and receives a payoff equals 0. Moreover, given any mixed strategy profile of players 2 and 3, player 1 receives positive payoff if she plays a mixed strategy best response. Therefore, the mixed strategy minimax payoff of player 1 equals 0. Now suppose that player 2 and player 3 remain silent. Then, player 1, if she is ambiguity averse, will play a prudent strategy. She will mix a and b with equal probability and will ex-ante receive her mixed strategy maxmin payoff, $-\frac{1}{2}$. Using similar argument (this game is some how symmetric), the reader can check that the mixed strategy minimax payoff of both players 2 and 3 equal 0 and that the mixed strategy maxmin payoff of both players 2 and 3 equals $-\frac{1}{2}$. Thus, employing Ellsberg strategy allows to settle punishment schemes that are more severe than classic minimax strategies.

Let

$$\mu_i^E = \min\nolimits_{P\in\mathcal{P}} \max\nolimits_{j\in\mathcal{J}(i)} \max\nolimits_{P'_j\in\mathcal{P}_j} u_i(P_{-j}, P'_j)$$

be the pure effective minimax payoff of player $i \in N$ in the Ellsberg game. We have the following lemma.

Lemma 21 *Let G be a finite normal form game. The pure strategy effective minimax payoff of a player in the Ellsberg game equals her mixed strategy effective maxmin payoff in the original game G.*

Proof. of Lemma 21. From Lemma 20, we have $\mu_i^E = \max_{j\in\mathcal{J}(i)} \max_{P_j\in\mathcal{P}_j} u_i(\times_{k\neq j}\Delta S_k, P_j)$. Let $j \in \mathcal{J}(i)$ and $p_j \in \Delta S_j$. We have

$$u_i(\times_{k\neq j}\Delta S_k, \{p_j\}) \leq \max_{P_j\in\mathcal{P}_j} u_i(\times_{k\neq j}\Delta S_k, P_j)$$

and therefore

$$\min_{p_{-j}\in\times_{k\neq j}\Delta S_k} u_i(p_{-j}, p_j) \leq \mu_i^E.$$

It follows that

$$\nu_i \leq \mu_i^E.$$

That is, the mixed strategy effective maxmin payoff of player i in the Ellsberg game is less than or equal to her effective minimax payoff in the Ellsberg game. The effective minimax payoff of player i in the Ellsberg game is less than or equal her mixed strategy effective maxmin payoff as well. Indeed,

$$\begin{aligned} \mu_i^E &= \max_{j\in\mathcal{J}(i)} \max_{P_j\in\mathcal{P}_j} u_i(\times_{k\neq j}\Delta S_k, P_j) \\ &= \max_{P_{j^*}\in\mathcal{P}_{j^*}} u_i(\times_{k\neq j^*}\Delta S_k, P_{j^*}) \\ &= u_i(\times_{k\neq j^*}\Delta S_k, P_{j^*}^*) \end{aligned}$$

for some $j^* \in \mathcal{J}(i)$ and $P_{j^*}^* \in \mathcal{P}_{j^*}$. We have

$$\begin{aligned} \mu_i^E &= \min_{p_{-j^*}\in\times_{k\neq j^*}\Delta S_k, p_{j^*}\in P_{j^*}^*} u_i(p_{-j^*}, p_{j^*}) \\ &= \min_{p_{-j^*}\in\times_{k\neq j^*}\Delta Sk} u_i(p_{-j^*}, p_{j^*}^*) \end{aligned}$$

for some $p_{j^*}^* \in P_{j^*}^*$. As $p_{j^*}^* \in P_{j^*}^* \subseteq \Delta S_{j^*}$, we have

$$\mu_i^E \leq \max_{p_{j^*}\in\Delta S_{j^*}} \min_{p_{-j^*}\in\times_{k\neq j^*}\Delta S_k} u_i(p_{-j^*}, p_{j^*}).$$

So,

$$\mu_i^E \leq \max_{j\in\mathcal{J}(i)} \max_{p_j\in\Delta S_j} \min_{p_{-j}\in\times_{k\neq j}\Delta S_k} u_i(p_{-j}, p_j).$$

We conclude that $\mu_i^E = \nu_i$. So, the reservation value of a player in the Ellsberg game equals her mixed strategy effective maxmin payoff. ■

Further notations

Let $G = (N, S, u)$ be a finite normal form game and let γ be a number that is strictly greater than any payoff a player might receive in the game G. Let $\tau(G) = (N, S, u')$ be the normal form game where the payoff function u'_i of player $i \in N$ is equals γ if i has distinct Ellsberg equilibrium payoff in the game G. In the case player i has a unique Ellsberg equilibrium payoff in the game G, $u'_i(s) = u_i(s)$ for all $s \in S$. For all $l > 0$, let N_l be the set of players who have their payoff function equal to the constant γ in the game $\tau^{(l)}(G)$, where $\tau^{(l)}$ is the l th compound of τ. Let h be minimal such that N_h is a maximal element of the sequence $\{N_l\}_{l=1}^{\infty}$.

Definition 7 *The sequence $N_0 = \emptyset \subseteq N_1 \subseteq ... \subseteq N_h$ is the Ellsberg decomposition of the game G.*

Definition 8 *The Ellsberg decomposition $N_0 = \emptyset \subseteq N_1 \subseteq ... \subseteq N_h$ is complete if $N_h = N$.*

4.2.2 The finitely repeated game

Let G be a finite normal form game which I will refer to as the stage game. Given $T > 1$ and $\delta < 1$, let $G(\delta, T)$ be the game obtained by repeating the stage game T times and where players' discount factor is δ. In the game $G(\delta, T)$, in every round, each player observes the properties of the profile of Ellberg strategies chosen (or equivalently the properties of the randomization devices chosen by players) as well as the realized action profile, receives her payoff as in the stage game and chooses her Ellsberg strategy for the next period. A player may therefore condition her behavior on the history of Ellsberg profiles used in the previous periods. Formally, a strategy of player i in the repeated game $G(\delta, T)$ is a map $\sigma_i : \cup_{t=1}^{T} \mathcal{P}^{t-1} \rightarrow \mathcal{P}_i$ where $\mathcal{P}^0$ is the empty set. Given a history $h^t = (h_1, .., h_{t-1}) \in \mathcal{P}^{t-1} = \mathcal{P} \times ... \times \mathcal{P}$, the strategy σ_i of player i recommends to play the Ellsberg strategy $\sigma_i(h^t)$ at period $t, 1 \leq t \leq T$. In the repeated game $G(\delta, T)$, the discounted average payoff of a player given a play path $(s^1, ..., s^T) \in S^T$ is

$$u_i^{\delta}(s^1, ..., s^T) = \tfrac{1-\delta}{1-\delta^T} \Sigma_{t=1}^{T} \delta^{t-1} u_i(s^t).$$

The strategy profile $\sigma = (\sigma_1, ..., \sigma_n)$ induces a set of probability distributions $P(\sigma)$ over the set S^T of play paths of length T. Players are ambiguity averse and aim to maximize the minimal expected payoff that they could get from the set $P(\sigma)$. That is, given σ_{-i}, player i chooses σ_i in order to maximize

$$u_i^\delta(\sigma_{-i}, \sigma_i) = \min_{p \in P(\sigma_{-i}, \sigma_i)} \sum_{h \in S^T} p(h) u_i^\delta(h)$$

where $p(h)$ is the probability with which the history h is observed according to the probability distribution p. The strategy profile σ is an Ellsberg equilibrium of $G(\delta, T)$ if for all player i, and given σ_{-i}, the strategy σ_i maximizes the minimal expected payoff of player i. The strategy profile σ is a **subgame perfect equilibrium** of $G(\delta, T)$ if for all $t < T$ and history $h^t \in S^{t-1}$, the restriction $\sigma_{|h^t}$ of the strategy profile σ to the observed history h^t is an Ellsberg equilibrium of the game $G(\delta, T - t + 1)$.

Any ex-ante payoff vector to a subgame perfect equilibrium strategy of the finitely repeated game with discounting dominates the mixed strategy effective maxmin payoff vector of the game G.

Lemma 22 *Let G be a finite normal form game, $\delta < 1$, $T > 0$, σ be a subgame perfect equilibrium of $G(\delta, T)$ and ν be the mixed strategy effective maxmin payoff vector of the game G. Then, $u_i^\delta(\sigma) \geq \nu_i$ for all $i \in N$.*

Indeed, playing a prudent strategy in each period of the finitely repeated game, at least one player of a given equivalence class can guarantee to herself (and therefore to the whole class) her effective maxmin payoff.

Lemma 23 *Let G be a finite normal form game. Any payoff vector that is ex-post approachable by means of subgame perfect equilibrium strategies of the finitely repeated game with discounting dominates the mixed strategy effective maxmin payoff vector of the game G.*

This lemma says that, if players are allowed to strategically make use of objective ambiguity, then, a necessary condition for a payoff vector to be ex-post approachable by means of subgame perfect equilibria of the finitely repeated game is that, the latter payoff vector dominates the mixed strategy effective maxmin payoff vector of the stage game G. Indeed, if a payoff vector is ex-post approachable by subgame perfect equilibria of the finitely repeated game, then, it is ex-ante approachable by subgame perfect equilibria and thus dominates the mixed strategy maxmin payoff vector.

4.3 Main result and discussion

In this section I present the main finding of this paper. It is convenient to introduce 2 definitions.

Definition 9 *Let G be a finite normal form game and σ be a strategy profile of the finitely repeated game with discounting $G(\delta, T)$. The support of $P(\sigma)$ is the set of histories $h \in S^T$ such that there exists a probability distribution in $P(\sigma)$ that assigns a strictly positive probability to the history h.*

Definition 10 *The support of a strategy profile of the finitely repeated game is the set of possible play paths.*

Definition 11 *Let G be a normal form game and x a payoff vector. The payoff vector x is ex-post approachable by means of subgame perfect equilibria of the finitely repeated with discounting if for any $\varepsilon > 0$, there exists $\underline{\delta} < 1$ and $\underline{T}$ such that, for all $\delta \geq \underline{\delta}$, $T \geq \underline{T}$, $G(\delta, T)$ has a subgame perfect equilibrium profile σ such that $\left\|u^{\delta}(h) - x\right\|_{\infty} < \varepsilon$ for all play path $h \in S^T$ in the support of $P(\sigma)$.*[2]

A payoff vector is ex-post approachable by mean of subgame perfect equilibria of the finitely repeated game if it can approached by subgame perfect equilibria that have the following property. the discounted payoff to any play path within the support of the strategy is closed enough to the given payoff vector.

4.3.1 Statement of the main result

Theorem 6 *Let G be a finite normal form game such that $V^* \neq \emptyset$. The following are equivalent.*

1. *G has a complete Ellsberg decomposition.*

2. *Any point of V^* is ex-post approachable by means of subgame perfect equilibria of the finitely repeated game with discounting.*

3. *The set of points of V that are approachable by means of subgame perfect equilibria of the finitely repeated with discounting has a relative interior point.*

[2]For all payoff vector $x = (x_1, ..., x_n)$, $\|x\|_{\infty} = max_{1 \leq i \leq n} |x_i|$

The most laborious part of the proof of Theorem 6 is to show that, under the statement 1) of Theorem 6, it is possible to ex-post approach any feasible payoff vector of the game G that dominates the mixed strategy effective maxmin payoff vector by means of subgame perfect equilibrium strategies of the finitely repeated game. The role of Statement 1) here is to leverage the behavior of players in the End-game, phase of equilibrium strategies of the finitely repeated game where essentially (recursive) equilibrium profiles of the stage game are played, see Lemma 25 and Lemma 26. As we do not assume that the dimension of the set of feasible payoff vectors equals the number of players, the block $\mathcal{J}(i)$ might contains more than one player. It is therefore not immediate to make use of the payoff asymmetry lemma of Abreu et al. (1994) to construct a suitable reward phase. Lemma 26 allows to independently reward players and motivate them to be effective punisher during a punishment phase.

Moreover, as the time horizon is finite, the powerful payoff continuation lemma of Fudenberg and Maskin (1991) does not apply. We obtain a version of the latter lemma for finitely repeated games with discounting which says that, for any positive ε, there exists 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 $\underline{\delta}$, see Lemma 24. Basically, the payoff continuation lemma for finitely repeated games provides an uniform integer k such that, any feasible payoff vector x can be approximated by a deterministic path of the same length k. Appending finitely many such deterministic paths, we obtain a deterministic path π whose discounted average is closed enough to the payoff vector x and, at any (sufficiently) early point of time, the continuation payoff of the path π is closed enough to the payoff vector x.

In Section 4.5.1, given a feasible payoff vector of G that dominates the mixed strategies effective maxmin payoff vector, I construct a sequence of subgame perfect equilibrium strategies of the finitely repeated game such that, ex-post, all the possible corresponding sequences of discounted payoff vectors converge to that target payoff vector.

4.3.2 Discussion

While both necessary and sufficient, Statement 1) of Theorem 6 is weaker than Smith's (1995) necessary and sufficient condition. Indeed, as mixed Nash equilibria of the stage-game are also Ellsberg equilibria, a complete Nash decomposition (see Smith (1995) for a formal definition of Nash decomposition) induces a complete Ellsberg decomposition. However, a complete Ellsberg decomposition does not necessarily induce a complete Nash decomposition. The three-player game whose payoff matrix is provided in Table 1.7 serves as an illustration. In that game, each player has a unique mixed Nash equilibrium payoff but many continuation equilibrium payoffs in Ellsberg actions (see Section 1.4 for details).

For the game in Table 1.7, the classic models of finitely repeated games in which players can employ only pure and mixed actions predict no cooperation at all. Our model predicts that any feasible payoff vector that dominates the mixed strategy effective maxmin payoff vector is approachable by means of subgame perfect equilibria of the finitely repeated game with discounting. Moreover, we are able to approximate the cooperative and Pareto superior payoff vector $(2, 2, 2)$ by means of a simple subgame perfect equilibrium of the finitely repeated game. Thus, the use of imprecise probabilistic devices in the finitely repeated game model can allow for an explanation of the emergence of cooperation in finite repetitions of a non-cooperative game where the classic models of finitely repeated games with pure and mixed strategies fail to do so.

As the Ellsberg extension of a finite normal form game is still a normal form game, it might appear logical to apply an existing limit perfect folk theorem [see, e.g., Benoit and Krishna (1984)] to the Ellsberg game and obtain the set of payoff vectors that are ex-ante approachable by means of the subgame perfect equilibrium strategies of the finitely repeated game. The previsions of Theorem 6 and Lemma 23 of this paper are different in the sense that they provide (under a weak condition) a characterization of the set of payoff vectors that are ex-post (and thus ex-ante) approachable by means of subgame perfect equilibrium strategies of the finitely repeated game. The difference between the former and the latter sets of payoff vectors can be clearly observed in the three-player game G, whose payoff matrix is given by Table 4.2.

In the Ellsberg extension Γ of the game G, each player has distinct Nash

	c	d
a	$0\ \ 0\ \ 0$	$1\ \ -1\ \ 1$
b	$-1\ \ 1\ \ 1$	$0\ \ 0\ \ -1$

e

c	d
$1\ \ 1\ \ -1$	$-1\ \ 1\ \ 1$
$1\ \ -1\ \ 0$	$0\ \ 0\ \ 0$

f

Table 4.2: Some Ellsberg payoff vectors are non feasible and an ex-ante approximation is vague.

equilibrium payoffs and no two players have equivalent utility functions. The limit perfect folk theorem of Benoit and Krishna (1984) states that any payoff vector that lies in the convex hull of the set of payoff vectors of the game Γ and which dominates the pure minimax payoff vector $(-\frac{1}{2}, -\frac{1}{2}, -\frac{1}{2})$ of the game Γ is approachable by means of subgame perfect Nash equilibrium strategies of finite repetitions of the game Γ [which is equivalent to being ex-ante approachable by means of subgame perfect (Ellsberg) equilibrium strategies of finite repetitions of G]. The payoff vector $(-\frac{1}{3}, -\frac{1}{3}, -\frac{1}{3})$ is therefore ex-ante approachable by means of subgame perfect equilibrium strategies of finite repetitions of the Ellsberg game. Note that, ex-post, in each period of finite repetitions of the game Γ, players receive payoffs as in the game G and it is not possible to implement/approach the payoff vector $(-\frac{1}{3}, -\frac{1}{3}, -\frac{1}{3})$ in the repeated game as the ex-post sum of payoffs of players 1 and 2 is always greater than or equal 0. More importantly, the payoff vector $(-\frac{1}{3}, -\frac{1}{3}, -\frac{1}{3})$ does not belong to the set of feasible payoff vectors of the game G. In addition, applying the existing limit perfect folk theorems to the Ellsberg game does not guarantee that any feasible payoff vector of the game G which dominates the mixed strategy effective maxmin payoff vector of the game G can be approached by subgame perfect Nash equilibrium strategy of finite repetitions of the Ellsberg game and whose ex-post payoff vector is closed enough to the target payoff vector.

4.4 Conclusion

This paper presented a model of finitely repeated games with complete information and perfect monitoring in which players can strategically make use of objective ambiguity. In addition to the classic pure and mixed actions, Ellsberg urns are available to players. An Ellsberg urn captures the quantity of information a player might want to know and share about her intentions. The main theorem provides a weak condition under which any feasible payoff vector that dominates the maxmin payoff vector of the stage-game is achievable via subgame perfect equilibria of the finitely repeated game with discounting. This new model explains how players can sustain collusive payoff vectors for some cases in which the classic models of finitely repeated games with pure and mixed actions fail to explain the emergence of cooperation.

4.5 Appendix 4: Proofs

4.5.1 Sketch of the proof of Theorem 6

In this section, Given a feasible payoff vector that dominates the mixed strategy effective maxmin payoff vector, I explain how to construct a subgame perfect equilibrium strategy σ of the finitely repeated game with discounting and whose ex-post payoff vectors is closed enough to the target payoff vector.

Let $y \in V^*$. The construction of σ involves few ingredients. The most important are the target path and the end-game-strategy. The target path is a finite sequence of pure action profiles of the stage game. It is obtained by applying our Lemma 24 (payoff continuation lemma for finitely repeated games) to the payoff vector y. The end-game-strategy corresponds to the very last phase of the game. It is a family of subgame perfect equilibria of the finitely repeated game. It allows to independently leverage the behavior of players in the finitely repeated game, regardless of whether some players are equivalent or not.

The strategy profile σ involves 5 phases. The first phase consists in some conjunction of the target path. If a player unilaterally deviates early during this phase, the strategy σ prescribes to start the second phase and thereafter to go to the third phase.

The second phase is a punishment phase where a potential deviator i is punished. There is no specific requirement for players of the block $N\backslash\mathcal{J}(i)$ while players of the block $\mathcal{J}(i)$ have to remain silent, that is completely ambiguous. At the end of this phase, we record in a boolean vector α, the set of players who were silent during the punishment phase. We prove that for large discount factor, an ambiguity averse player of the block $N\backslash\mathcal{J}(i)$ will find it strictly dominant to remain silent during the punishment phase.

The third phase serves as a compensation. Indeed, it might be the case that the punishment phase is more severe than required and players of the block $\mathcal{J}(i)$ may receive a negative ex-ante payoff in each period of the punishment phase. The fourth phase serves as a transition. In the fifth phase, players are credibly rewarded.

Note that, if no deviation from σ occurs in the repeated game, players will follow some loops of the target path and then move to the end-game-strategy. In Section 4.5.4 I show that for sufficiently long time horizon and large discount factor, the strategy profile σ is a subgame perfect equilibrium of the finitely repeated game and that the deterministic part of the resulting discounted average payoff will be close enough to y and the ambiguous part will goes to 0.

Now I proceed to the detailed proof of Theorem 6. To ease this proof, I introduce three lemmata.

4.5.2 The payoff continuation lemma for finitely repeated games

Lemma 24 *For any $\varepsilon > 0$, there exists $k > 0$ and $\underline{\delta} < 1$ such that for all $x \in V$, there exists a deterministic sequence of stage game actions $\{s^\tau\}_{\tau=1}^k$ whose discounted average payoff is within ε of x for all discount factor $\delta \geq \underline{\delta}$.*

Lemma 24 establishes that for any positive ε, one can construct 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 or equal $\underline{\delta}$. This lemma allow to approach any feasible payoff vector by deterministic paths of the finitely repeated game in presence of discount factor.

Proof. of Lemma 24. Let $\varepsilon > 0$ and $y = \sum_{l=1}^m \alpha^l u(a^l) \in V$ be a feasible payoff, where $a^l \in S$ for $l = 1, ..., m$. Assume that there exists m integers $q_1, q_2, ..., q_m$ such that for all $l = 1, ..., m$, $\alpha^l = \frac{q_l}{Q}$ where $Q = \sum_{l=1}^m q_l$. Consider the sequences $\{b^{y,p}\}_{p=1}^Q$ and $\{c^{y,\tau}\}_{\tau=1}^\infty$ defined as follows.

$b^{y,p} = a^l$ if and only if $\sum_{l'<l} q_{l'} < p \leq \sum_{l' \leq l} q_{l'}$

$c^{y,\tau} = b^{y,p}$ if and only if $\tau - p \equiv 0[Q]$.

To have a clear view of the construction of the sequences $\{b^{y,p}\}_{p=1}^Q$ and

$\{c^{y,\tau}\}_{\tau=1}^{\infty}$, consider this simple example where $m = 3$, $q_1 = 2$, $q_2 = 1$, $q_3 = 4$, $Q = 7$, and therefore $y = \frac{2}{7}u(a^1) + \frac{1}{7}u(a^2) + \frac{4}{7}u(a^2)$. Table 4.3 provides the value of $b^{y,p}$ for $p = 1, ..., 7$ while Table 4.4 provides the value of $c^{y,\tau}, \tau \geq 1$.

$b^{y,p}$	$b^{y,1}$	$b^{y,2}$	$b^{y,3}$	$b^{y,4}$	$b^{y,5}$	$b^{y,6}$	$b^{y,7}$
a^l	a^1	a^1	a^2	a^3	a^3	a^3	a^3

Table 4.3: Values of $b^{y,\tau}, \tau \geq 1$

$c^{y,\tau}$	$c^{y,1}$	$c^{y,2}$	$c^{y,3}$	$c^{y,4}$	$c^{y,5}$	$c^{y,6}$	$c^{y,7}$	$c^{y,8}$	$c^{y,9}$	$c^{y,10}$	$c^{y,11}$	$c^{y,12}$	$c^{y,13}$	$c^{y,14}$	...
$b^{y,p}$	$b^{y,1}$	$b^{y,2}$	$b^{y,3}$	$b^{y,4}$	$b^{y,5}$	$b^{y,6}$	$b^{y,7}$	$b^{y,1}$	$b^{y,2}$	$b^{y,3}$	$b^{y,4}$	$b^{y,5}$	$b^{y,6}$	$b^{y,7}$	...
a^l	a^1	a^1	a^2	a^3	a^3	a^3	a^3	a^1	a^1	a^2	a^3	a^3	a^3	a^3	...

Table 4.4: Values of $c^{y,\tau}, \tau \geq 1$

We can observe that the undiscounted average payoff of the sequence $\{c^{y,\tau}\}_{\tau=1}^{\infty}$ is equal to $\frac{2}{7}u(a^1) + \frac{1}{7}u(a^2) + \frac{4}{7}u(a^2)$.

Going back to the general case, let $l \in \{1, ..., m\}$, $\Theta = AQ + B$ where $A > 0$ and $0 \leq B < Q$ and consider

$$N(l, c^y, \Theta) = \{\tau \mid c^{y,\tau} = a^l\}$$

and

$$\beta(l, c^y, \Theta) = \tfrac{1-\delta}{1-\delta^{\Theta}} \textstyle\sum_{\tau \in N(l,c^y,\Theta)} \delta^{\tau-1}.$$

We have

$$\tfrac{1-\delta}{1-\delta^{\Theta}} \textstyle\sum_{\tau \leq \Theta} \delta^{\tau-1} u(c^{y,\tau}) = \sum_{l=1}^{m} \beta(l, c^y, \Theta) u(a^l)$$

and

$$\beta(1, c^y, \Theta) = \tfrac{1-\delta}{1-\delta^{\Theta}} \left[\tfrac{1-\delta^{p_1}}{1-\delta} \tfrac{1-\delta^{AQ}}{1-\delta^{Q}} + \delta^{AQ} \tfrac{1-\delta^{p_1'}}{1-\delta} \right]$$

where $p_1' = \min\{B, p_1\}$;

$$\beta(2, c^y, \Theta) = \tfrac{1-\delta}{1-\delta^{\Theta}} \left[\delta^{p_1} \tfrac{1-\delta^{p_2}}{1-\delta} \tfrac{1-\delta^{AQ}}{1-\delta^{Q}} + \delta^{AQ+p_1} \tfrac{1-\delta^{p_2'}}{1-\delta} \right]$$

where $p_2' = \min\{\max\{0, B - p_1\}, p_2\}$;

.

.

.

$$\beta(m, c^y, \Theta) = \frac{1-\delta}{1-\delta^\Theta}\left[\delta^{p_1+...+p_{m-1}}\frac{1-\delta^{p_m}}{1-\delta}\frac{1-\delta^{AQ}}{1-\delta^Q} + \delta^{AQ+p_1+...+p_{m-1}}\frac{1-\delta^{p'_m}}{1-\delta}\right]$$

where $p'_m = \min\{\max\{0, B - p_1 - ... - p_{m-1}\}, p_m\}$.

As

$$\lim_{\delta\to 1}\beta(l, c^y, \Theta) = \frac{p_l}{AQ+B}\frac{AQ}{Q} = \frac{p_l}{Q+\frac{B}{A}}$$

and

$$\lim_{A\to+\infty}\frac{p_l}{Q+\frac{B}{A}} = \frac{p_l}{Q},$$

there exists $\underline{A}^y > 0$ such that, for all $A \geq \underline{A}^y$, there exists $\delta^{y,A} < 1$ such that for all $\delta > \delta^{y,A}$,

$$\left\|\frac{1-\delta}{1-\delta^\Theta}\sum_{\tau\leq\Theta}\delta^{\tau-1}u(c^{y,\tau}) - y\right\| < \frac{\varepsilon}{2}$$

for all B, $0 \leq B < Q$. Let $\{\widetilde{B}(y, \frac{\varepsilon}{2}), y \in Y\}$[3] be a finite open covering of the compact set V where Y is the set of convex sum of stage game payoff vectors with rational coefficients. Pose $\underline{A} = \max\{\underline{A}^y, y \in Y\}$, $k = Q(\underline{A}+1)$ and $\underline{\delta} = \max\{\delta^{y,\underline{A}+1}, y \in Y\}$. Let $x \in V$ and $y \in Y$ such that $x \in \widetilde{B}(y, \frac{\varepsilon}{2})$. Take $s^\tau = c^{y,\tau}$ for $\tau = 1, ..., k$. ■

The next two lemmata explain how to leverage the behavior of players in the very last phase of the game where essentially only stage game (recursive) equilibrium profile are played.

4.5.3 The end-game-strategy

Lemma 25 *Assume that the Ellsberg decomposition $\emptyset \subseteq N_1 \subseteq ... \subseteq N_h$ of the game G is complete. Then there exists $\phi_e > 0$, $T > 0$, $\underline{\delta} \in (0,1)$ and for all $i \in N$, there exists $\sigma^{i,1}$ and $\sigma^{i,2}$ two strategy profiles of the $T-$fold repeated game such that*

1. *$\sigma^{i,1}$ and $\sigma^{i,2}$ are subgame perfect equilibria of the finitely repeated game $G(\delta, T)$ for all $\delta \in (\underline{\delta}, 1)$;*
2. *$u_i^\delta(\sigma^{i,1}) > \phi_e + u_i^\delta(\sigma^{i,2})$.*

[3] $\widetilde{B}(y, \frac{\varepsilon}{2}) = \{z \in V \mid \|z - y\|_\infty < \frac{\varepsilon}{2}\}$

Lemma 26 *Suppose that the stage-game G has a complete Ellsberg decomposition. Then there exists $\phi > 0$ such that for all $p \geq 0$, there exists $r_p > 0, \underline{\delta} \in (0,1)$ and a family $\{\theta^p(\gamma) \mid \gamma \in \{0,1\}^n \cup \{(-1,\cdots,-1)\}\}$ of strategy profiles of the r_p-fold repeated game such that for all $\delta \in (\underline{\delta},1)$ and $\gamma \in \{0,1\}^n \cup \{(-1,\cdots,-1)\}$, $\theta^p(\gamma)$ is a subgame perfect equilibrium of the finitely repeated game $G(\delta, r_p)$. Furthermore, for all $\delta \in (\underline{\delta},1)$ $i \in N$ and $\gamma, \gamma' \in \{0,1\}^n$ we have*

$$u_i^\delta[\theta^p(1,\gamma_{-i})] - u_i^\delta[\theta^p(0,\gamma_{-i})] \geq \phi \tag{4.1}$$

$$u_i^\delta[\theta^p(\gamma)] - u_i^\delta[\theta^p(-1,\cdots,-1)] \geq \phi \tag{4.2}$$

$$|u_i^\delta[\theta^p(\gamma)] - u_i^\delta[\theta^p(\gamma_{\mathcal{J}(i)}, \gamma'_{\mathcal{J}\setminus\mathcal{J}(i)})]| < \frac{1}{2^p}. \tag{4.3}$$

The proofs of Lemmata 25 and 26 are respectively similar to the proofs of Lemmata 8 and 9 and are therefore omitted.

4.5.4 Proof of the Theorem 6

Proof of Theorem 6. Let G be a finite normal form game such that $V^* \neq \emptyset$. Let's shift the utility function of the game G to have the effective maxmin payoff of each player equal to 0 and so that within the same equivalence class, players, if many, have the same payoff function. This does not change the strategic behavior of players.

Part 1. (1⇒2). Assume that the Ellsberg decomposition of the game is complete. Let $\varepsilon > 0$ and $y \in V^*$. I wish to construct $\underline{\delta} < 1$ and $\underline{T} > 0$, and for all $\delta \geq \underline{\delta}$ and $T \geq \underline{T}$, a subgame perfect equilibrium strategy profile σ^T of $G(\delta, T)$ such that $\left\|u^\delta(h) - y\right\|_\infty < 3\varepsilon$ for all history h in the support of $P(\sigma^T)$.

Apply the payoff continuation lemma (see Lemma 24) to ε and obtain $k > 0$, $\underline{\delta}_0 < 1$, and a deterministic path

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

such that

$$d(y, u^\delta(\pi^y)) < \varepsilon$$

for all $\delta \in (\underline{\delta}_0, 1)$. For all $\delta \in (\underline{\delta}_0, 1)$, let

$$\overline{y} = \lim_{\delta \to 1} u^\delta(\pi^y).$$

Obtain ϕ, r_1 and θ^1 with $p = 1$ from the Lemma 26 and let

$$u^{1,r_1}[\theta^1(1, \cdots, 1)] = \lim_{\delta \to 1} u^\delta[\theta^1(1, \cdots, 1)].$$

Let $q_1 > 0$ and $q_2 > 0$ such that

$$0 < q_1 k u_i(\Delta S) + q_2 r_1 k u_i^{1,r_1}[\theta^1(1, \cdots, 1)] < \frac{q_1 k + q_2 r_1 k}{2} \overline{y}_i \text{ for all } i \in N$$

and

$$-2k\rho + \frac{q_1 k}{2} \overline{y}_i > 0 \text{ for all } i \in N$$

where

$\rho = \max_{a \in A} \|u(a)\|_\infty$.

Given q_1, q_2 and r_1, choose r such that

$$-2(q_1 k + q_2 r_1 k)\rho + r\phi > 0.$$

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

$$\frac{q_2 r_1 k}{2} \overline{y}_i - \frac{r}{2^p} > \overline{y}_i - \frac{r}{2^p} > 0 \text{ for all } i \in N.$$

Apply the Lemma 26 to p and obtain r_p and θ^p. Update $q_1 \leftarrow r_p q_1; q_2 \leftarrow r_p q_2$; $r \leftarrow r_p r$. The parameters ϕ, θ^1, q_1, q_2, r, r_1 and θ^p are such that

$$-2q_1 k\rho + r\phi > 0; \tag{4.4}$$

$$\overline{y}_i - \frac{r}{2^p} > 0; \tag{4.5}$$

$$2k\rho + q_1 k u_i(\Delta S) + q_2 r_1 k u_i^{1,r_1}[\theta^1(1, \cdots, 1)] + \frac{r}{2^p} - (q_1 k + q_2 r_1 k - k)\overline{y}_i < 0 \tag{4.6}$$

and

$$-2(q_1 k + q_2 r_1 k)\rho + r\phi > 0 \text{ for all } i \in N. \tag{4.7}$$

Let

$$\pi = (\underbrace{\pi^y, ..., \pi^y}_{C + q_1 + q_2 r_1 \text{ times}})$$

Set $\alpha = (1, \cdots, 1)$.

From now on, a deviation by a player from an ongoing path is called "early deviation" if it occurs during the first Ck periods of the game. In the other case, the deviation is called "late deviation". Consider the strategy profile σ of the finitely repeated game described by the following 5 phases.

$\mathbf{P}_0$ (Main path): At any time t, play the t th action profile of the path π. [If player i deviates early, start the Phase $\mathbf{P}_i$; if player i deviates late, start **LD**. Ignore any simultaneous deviation.] Go to Phase **EG**.

$\mathbf{P}_i$ (Punish player i): Reorder the profile of actions in each upcoming cycle of length k of the main path according to player i's preferences, starting from her best profile.
This phase last for $q_1 k$ periods and each player of the block $\mathcal{J}(i)$ has to remain silent (completely ambiguous). Each player of the block $N \backslash \mathcal{J}(i)$ can play whatever Ellsberg action she wants. [If any player $j \in \mathcal{J}(i)$ deviates early, start $\mathbf{P}_j$; if player $j \in \mathcal{J}(i)$ deviates late, start **LD**.]

At the end of this phase, for all $j \notin \mathcal{J}(i)$, set $\alpha_j = 0$ if there is at least one period of the punishment phase where player j was not silent (completely ambiguous) and set $\alpha_j = 1$ otherwise. Go to Phase **SPE**.

SPE Follow $q_2 k$ times the subgame perfect equilibrium $\theta^1(1, \cdots, 1)$. Go to Phase $\mathbf{P}_0$.

LD Each player can play whatever action she wants till period $(C+q_1+q_2 r_1)k$. At period $(C + q_1 + q_2 r_1)k$, set $\alpha = (-1, \cdots, -1)$ and go to Phase **EG**.

EG Follow $\frac{r}{r_p}$ times the subgame perfect equilibrium $\theta^p(\alpha)$.

Now, I show that given any history, there is no profitable unilateral and single shot deviation if the discount factor is high enough.

c-1) **It is strictly dominant for any player $j \notin \mathcal{J}(i)$ to remain silent during a punishment phase $\mathbf{P}_i$.**

As players are ambiguity averse and aim to maximize their worst expected utility, they will individually find it strictly dominant to remain silent during any punishment phase. Indeed, if during a punishment phase (say $\mathbf{P}_i$) player $j \notin \mathcal{J}(i)$ is silent, she receives at least

1. $-\frac{1-\delta^{q_1 k}}{1-\delta}\rho$ during the punishment phase;

2. $\delta^{q_1 k}\frac{1-\delta^{q_2 r_1 k}}{1-\delta}u_j^{\delta}(\theta^1(1,\cdots,1))$ during the **SPE** phase;

3. some payoff $U_j(\delta)$ up to the period $(C+q_1+q_2 r_1)k$;

4. an ex-ante payoff $\delta^{c+(q_1+q_2r_1)k}\frac{1-\delta^r}{1-\delta}u_j^{\delta}[\theta^p(1,\alpha_j)]$ in the Phase **EG**.

In total, she gets

$$-\frac{1-\delta^{q_1 k}}{1-\delta}\rho+\delta^{q_1 k}\frac{1-\delta^{q_2 r_1 k}}{1-\delta}u_j^{\delta}(\theta^1(1,\cdots,1))+U_j(\delta)+\delta^{c+(q_1+q_2r_1)k}\frac{1-\delta^r}{1-\delta}u_j^{\delta}[\theta^p(1,\alpha_j)]$$

If player j is not silent during the Phase $\mathbf{P}_i$, she receives at most

1. $\frac{1-\delta^{q_1 k}}{1-\delta}\rho$ during the punishment phase;

2. $\delta^{q_1 k}\ \frac{1-\delta^{q_2 r_1 k}}{1-\delta}u_j^{\delta}(\theta^1(1,\cdots,1))$ during the **SPE** phase;

3. the same payoff $U_j(\delta)$ till period $(C+q_1+q_2r_1)k$;

4. an ex-ante payoff $\delta^{c+(q_1+q_2r_1)k}\frac{1-\delta^r}{1-\delta}\left[u_j^{\delta}[\theta^p(1,\alpha_j)]-\phi\right]$ in the Phase **EG**, see Lemma 26.

In total, she gets

$$\frac{1-\delta^{q_1 k}}{1-\delta}\rho+\delta^{q_1 k}\ \frac{1-\delta^{q_2 r_1 k}}{1-\delta}u_j^{\delta}(\theta^1(1,\cdots,1))+U_j(\delta)+\delta^{c+(q_1+q_2r_1)k}\frac{1-\delta^r}{1-\delta}\left[u_j^{\delta}[\theta^p(1,\alpha_j)]-\phi\right]$$

Thus, player j will find it strictly dominant to remain silent if

$$-2\frac{1-\delta^{q_1 k}}{1-\delta}\rho+\delta^{c+(q_1+q_2r_1)k}\frac{1-\delta^r}{1-\delta}\phi>0 \tag{4.8}$$

As δ goes to 1, the left hand of the latter inequality goes to $-2q_1k\rho+r\phi$ which is strictly positive, see Equation (4.4). Therefore, there exists $\underline{\delta}_1\in(\underline{\delta}_0,1)$ such that the Inequality (4.8) holds for all $\delta\in(\underline{\delta}_1,1)$.

Now assume that $\delta\in(\underline{\delta}_1,1)$ so that, it is strictly dominant for any player $j\notin\mathcal{J}(i)$ to be silent on punishment phases. We wish to prove that, for sufficiently large discount factor, σ is a subgame perfect equilibrium strategy of the finitely repeated game.

c-2) **No early deviation from Phase $\mathbf{P}_i$ by a player $j \in \mathcal{J}(i)$ is profitable**.

If after l_1k+l_2 rounds in the Phase $\mathbf{P}_i$ player $j \in \mathcal{J}(i)$ deviates, she receives:

1. at most 0 from the beginning of the Phase $\mathbf{P}_i$ till the deviation period;
2. an ex-ante payoff $\delta^{l_1k+l_2+1}U_j^1(\delta)$ during the Phases $\mathbf{P}_j$ and the new ***SPE*** phase;
3. some payoff $U_j^2(\delta)$ till the period $(C+q_1+q_2r_1)k$;
4. an ex-ante payoff $\delta^{c+(q_1+q_2r_1)k}\frac{1-\delta^r}{1-\delta}u_j^\delta[\theta^p(\alpha)]$ in the Phase **EG**.

If player j does not deviates, she receives at least:

1. the ex-ante payoff $U_j^1(\delta)+\delta^{q_1k+q_2r_1k}\frac{1-\delta^{l_1k+l_2+1}}{1-\delta}u_i^\delta(\pi^y)$ till the end of the new ***SPE*** phase;
2. the payoff $\widetilde{U}_i^2(\delta)$ till period $(C+q_1+q_2r_1)k$;[4]
3. the ex-ante payoff $\delta^{c+(q_1+q_2r_1)k}\frac{1-\delta^r}{1-\delta}\left[u_j^\delta[\theta^p(\alpha)]-\frac{1}{2^p}\right]$ in the Phase **EG**, see Lemma 26.

As $\overline{y}_i-\frac{r}{2^p}>0$ [see Equation (4.5)], there exists $\underline{\delta}_2 \in (\underline{\delta}_1, 1)$ such that for all $\delta \in (\underline{\delta}_2, 1)$, no early deviation from Phase $\mathbf{P}_i$ is profitable.

c-3) **No early deviation from Phase $\mathbf{P}_0$ is profitable**.

If during Phase $\mathbf{P}_0$, a player let's say i deviates early, the strategy profile σ prescribes to start the punishment phase $\mathbf{P}_i$ followed by the Phase **SPE**, to update the boolean vector α and to go back to the Phase $\mathbf{P}_0$. For sufficiently high discount factor, such a deviation is not profitable. Indeed, if player i deviates early during Phase $\mathbf{P}_0$, she receives at most

1. ρ in the deviation period;
2. $\delta\frac{1-\delta^{q_1k}}{1-\delta}u_i(\Delta S)+\delta^{q_1k+1}\frac{1-\delta^{q_2r_1k}}{1-\delta}u_i^\delta[\theta^1(1,\cdots,1)]$ in the punishment phase;[5]

[4] Recall that $U_i^2(\delta)$ is the discounted sum of a deterministic sequence of payoffs and $\widetilde{U}_i^2(\delta)$ is the discounted sum over a permutation of the same deterministic sequence of payoffs. Therefore, as δ goes to 1, both sums converge to the same limit.

[5] Recall that all players will be effective punishers.

3. some payoff $U_i^2(\delta)$ till the period $(C+q_1+q_2r_1)k$;

4. an ex-ante payoff $\delta^{c+(q_1+q_2r_1)k}\frac{1-\delta^r}{1-\delta}u_j^\delta[\theta^p(\alpha)]$ in the Phase **EG**.

If player i does not deviate during the Phase $\mathbf{P}_0$, she receives at least

1. $-\frac{1-\delta^l}{1-\delta}\rho$ till the end of the ongoing $k-$cycle (for some $l \leq k$);

2. $\delta^l\frac{1-\delta^{q_1k+q_2r_1k-k}}{1-\delta}u_i^\delta(\pi^y) - \delta^{l+q_1k+q_2r_1k-k}\frac{1-\delta^{1+k-l}}{1-\delta}\rho$ corresponding to the Phase $\mathbf{P}_i$ and the Phase **SPE**;

3. the payoff $\widetilde{U}_i^2(\delta)$ till the period $(C+q_1+q_2r_1)k$;[6]

4. the ex-ante payoff $\delta^{c+(q_1+q_2r_1)k}\frac{1-\delta^r}{1-\delta}u_j^\delta[\theta^p(\alpha_{\mathcal{J}(i)},\alpha'_{-\mathcal{J}(i)})]$ in the Phase **EG**. From Lemma 26, the latter ex-ante payoff is greater than or equal to $\delta^{c+(q_1+q_2r_1)k}\frac{1-\delta^r}{1-\delta}\left(u_j^\delta[\theta^p(\alpha)] - \frac{1}{2^p}\right)$.

Therefore, as δ goes to 1, the limit of the profit from deviating is above bounded by

$$2k\rho + q_1ku_i(\Delta S) + q_2r_1ku_i^{1,r_1}[\theta^1(1,\cdots,1)] + \frac{r}{2^p} - (q_1k+q_2r_1k-k)\overline{y}_i$$

which is strictly negative, see Equation (4.6).

Therefore, there exists $\underline{\delta}_3 \in (\underline{\delta}_2, 1)$ such that for all $\delta \in (\underline{\delta}_3, 1)$, no early deviation from Phase $\mathbf{P}_0$ is profitable.

c-4) **No late deviation is profitable.**

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

1. $\frac{1-\delta^{q_1k+q_2r_1k}}{1-\delta}\rho$ till the beginning of the phase **EG**;

2. the ex-ante payoff $\delta^{q_1k+q_2r_1k}\frac{1-\delta^r}{1-\delta}u_j^\delta[\theta^p(-1,\cdots,-1)]$ in the Phase **EG**.

If player j does not deviates, she receives at least

1. $-\frac{1-\delta^{q_1k+q_2r_1k}}{1-\delta}\rho$ till the beginning of the phase **EG**;

[6] Note that $\lim_{\delta\to 1}\widetilde{U}_i^2(\delta) = \lim_{\delta\to 1}U_i^2(\delta)$.

2. the ex-ante payoff $\delta^{q_1 k + q_2 r_1 k} \frac{1-\delta^r}{1-\delta} u_j^\delta[\theta^p(\alpha)]$ in the Phase **EG**, where $\alpha \in \{0.1\}^n$. From Lemma 26, the latter ex-ante payoff is strictly greater than $\delta^{q_1 k + q_2 r_1 k} \frac{1-\delta^r}{1-\delta} \left(u_j^\delta[\theta^p(-1, \cdots, -1)] + \phi \right)$.

As $-2(q_1 k + q_2 r_1 k)\rho + r\phi > 0$ [see Equation (4.7)], there exists $\underline{\delta}_4 \in (\underline{\delta}_3, 1)$ such that for all $\delta \in (\underline{\delta}_4, 1)$, no late deviation is profitable.

Therefore, for all $\delta \in (\underline{\delta}_4, 1)$ and given any history h of the repeated game, no player has any incentive to deviate from $\sigma_{|h}$. That is σ is a subgame perfect equilibrium for all $C > 0$. Choose $\underline{C}$ high enough and $\underline{\delta} \in (\underline{\delta}_4, 1)$ such that

$$\frac{1-\delta^k}{1-\delta^{(\underline{C}+q_1+q_2 r_1)k+r}} \rho + \delta^{(\underline{C}+q_1+q_2 r_1)k} \frac{1-\delta^r}{1-\delta^{(\underline{C}+q_1+q_2 r_1)k+r}} \rho < \varepsilon$$

For all $T \geq \underline{T}$ and $\delta \in (\underline{\delta}, 1)$, let σ^T be the restriction of σ to the last T periods of the finitely repeated game $G(\delta, T)$. Let h be an history in the support of $P(\sigma^T)$. We have

$$\left\| u^\delta(h) - u^\delta(\pi) \right\|_\infty < 2\varepsilon$$

and therefore

$$\left\| u^\delta(h) - y \right\|_\infty < 3\varepsilon$$

for all $T \geq \underline{T}$ and $\delta \geq \underline{\delta}$.

Part 2. (2⇒3). Assume that any point of V^* is approachable by means of subgame perfect Ellsberg equilibrium of the finitely repeated game. As V^* is non empty, V^* has non empty relative interior and statement 3) of Theorem 6 holds.

Part 3. (3⇒1). Assume that the Ellsberg decomposition of the game G is incomplete. By induction on the time horizon, players of the block $N \backslash N_n$ receive their unique stage game Ellsberg equilibrium payoff in each period of the finitely repeated game. That is, any player of the block $N \backslash N_n$ has a unique subgame perfect equilibrium payoff in the finitely repeated game. This contradicts the statement 3) of Theorem 6.

■

Chapter 5

Infinitely repeated games with discounting. What changes if players are allowed to use imprecise devices.

Abstract: In this note, I present a model of infinitely repeated game with complete information and perfect monitoring and where players are allowed to employ pure actions, mixed actions as well as an additional device which captures the willingness of a player to exercise her right to remain silent. I show that any feasible payoff that dominates the maxmin (modulo some players have equivalent utility functions) payoff vector is sustainable by means of pure strategy subgame perfect Nash equilibria of the infinitely repeated game with discounting.

5.1 Introduction

In game of conflict, cooperation is often observed. Even in a world where binding agreements can not be written. Some examples are the prisoners' dilemma, the peace negotiation game and the Cournot duopoly. The theory of repeated game argues that, an agent involved in such conflicts may expect to have a long-term relationship and, may abandon her short term interest and behave nicely because she expects a some future reward or alternatively because she fears future retaliations. Famous results of this theory are known as Folk Theorems. A Folk Theorem saids to holds if the set of Nash equilibrium payoffs or alternatively the set of subgame perfect Nash equilibrium payoffs of the repeated game includes any feasible payoff vector that dominates the minimax payoff vector. The idea behind Folk Theorems is that, if the stage game is rich enough [see for instance Benoit and Krishna (1984), Smith (1995) and Fudenberg and Maskin (1991) for sufficient and necessary conditions], players can construct credible punishment and reward paths and therefore sustain collusive payoffs, using Nash equilibrium or subgame perfect Nash equilibrium strategies of the repeated game.

Repeated game models till now have assumed that, in each period, players predetermine their actions (possibly mixed) for the next period. This assumption is quite restrictive as, on punishment path for instance, the player being punished best responds to the punishment strategy setled by her fellow players. To be efficient, punishers might find optimal to be impredictable and therefore not predetermine their actions or the probability distributions their actions will be issued from, and conceal their intentions in imprecise (and possibly probabilistic) devices as Ellsberg urns (urns with unknown composition). Such behavior creates an objective ambiguity to the target player and she can not secure a payoff that is strictly greater than her maxmin payoff.

Indeed, if the punishers exercise their right to remain silent on a punishment path and if in addition the target player is ambiguity averse (Ellsberg (1961) illustrates that players do display aversion to ambiguity), then she will respond to the silence of her fellow players by playing a prudent strategy and will expect to receive her maxmin payoff in each period of the ongoing punishment phase.

For some games, the maxmin payoff vector is strictly dominated by the minimax payoff vector. In those cases, using the model presented in this paper, it is possible to sustain some feasible payoff vectors that do not dominate the minimax payoff vector by means of subgame perfect equilibrium strategies of the infinitely repeated game. An illustration is the famous two-player zero sum game rock-paper-scissors where player are allowed to choose only pure actions. If randomization devices are available so that players can choose mixed actions, then the three-player game whose payoff matrix is given by Table 5.1, and where player 1 chooses the row (a or b), player 2 chooses the column (c or d) and player 3 chooses the matrix (e or f) is an example. In fact, each player has a minimax payoff equal to 0 and a maxmin payoff equal to $-\frac{1}{2}$.

	c	d
a	0 0 0	1 −1 1
b	−1 1 1	0 0 −1

e

c	d
1 1 −1	−1 1 1
1 −1 0	0 0 0

f

Table 5.1: Payoff matrix of a three-player game where the mixed strategy maxmin payoff vector is strictly dominated by the mixed strategy minimax payoff vector.

In this paper, I analyze infinitely repeated games with complete information, perfect monitoring and discounting. I allow players to exercise their right to remain silent so that they can conceal their strategy in imprecise (and possibly probabilistic) devices as Ellsberg urn. The main finding is that, any payoff vector that dominate the effective maxmin payoff vector is sustainable by means of pure strategy subgame perfect equilibrium of the infinitely repeated game.

5.2 The stage game

I represent a compact normal form game G by $(N, S, u = (u_i)_{i\in N})$ where the set of players $N = \{1, ..., n\}$ is finite and the set $S = \prod_{i\in N} S_i$ of actions is compact[1]. Given a player $i \in N$, S_i denotes the set of pure actions of player i. The utility of player i given $s = (s_1, ..., s_n) \in S$ is measured by $u_i(s)$. The utility function u is assumed to be continue. For all players $i, j \in N$, player i is equivalent to player j if there exists two real numbers β_{ij} and $\alpha_{ij} > 0$ such that $u_i(s) = \alpha_{ij} u_j(s) + \beta_{ij}$ for all $s \in S$. Denotes by $\widetilde{i}$ the set of players that are equivalent to player i. Let $\mu_i = \max_{j\in\widetilde{i}} \max_{p_j\in S_j} \min_{s_{-j}\in S_{-j}} u_i(s_{-j}, s_j)$ be the **pure strategy effective maxmin payoff of player** i and $\mu = (\mu_1, ..., \mu_n)$ the pure strategy effective maxmin payoff vector of the game G. The payoff vector $x = (x_1, ..., x_n)$ is **feasible** if there exists a sequence $(\lambda_l)_{1\le l\le p}$ of positive real numbers and a sequence $(a^l)_{1\le l\le p}$ of profile of pure actions such that $\Sigma_{l=1}^{p} \lambda_l = 1$ and $x = \Sigma_{l=1}^{p} \lambda_l u(a^l)$. Let V be the set of all feasible payoff vectors and V^* be the set of feasible payoff vectors that strictly dominate μ.

[1]One can think of G as a finite normal for game or a mixed extension of a finite normal form game.

5.3 The infinitely repeated game

Let G be a compact normal form game which we will refer to as the stage game. Given $\delta < 1$, let $G(\delta)$ be the infinitely repeated game with discount factor δ where players can use only pure actions. In the game $G(\delta)$, in each point of time, each player observes the profile chosen in the previous period, receive her payoff to the realized action profile as in the stage game and makes her choice for the next period (player i might choose a single pure action from S_i or decide to remain silent and choose the whole set S_i). The choice of a player at a given period may depend on the observed history. Formally, a pure strategy of player i in the game $G(\delta)$ is a map $\sigma_i : \cup_{t\geq 1} S'^{t-1} \to S'_i$ where S'^0 is the empty set, $S'_i = S_i \cup \{S_i\}$ and $S' = \prod_{i\in N} S'_i$. Given any history $h^t = (h_1, .., h_{t-1}) \in S^{t-1}$, the strategy σ_i of player i recommends to play $\sigma_i(h^t) \in S'_i$ at period t, $t \geq 1$. In the repeated game $G(\delta)$, the discounted average payoff of player $i \in N$ given any play path $\pi = (s^1, ..., s^t, ...) \in S^\infty$ is

$$u_i^\delta(\pi) = (1-\delta)\Sigma_{t=1}^{\infty} \delta^{t-1} u_i(s^t).$$

As players may exercise their right to remain silent in some periods, a pure strategy profile $\sigma = (\sigma_1, ..., \sigma_n)$ might generate different play paths. Let $P(\sigma)$ be set of possible play paths that can be generated by σ. The profile σ is a pure strategy Nash equilibrium of $G(\delta)$ if for all player i, and given σ_{-i}, the pure strategy σ_i maximizes the minimal expected payoff $\min_{\pi\in P(\sigma)} u_i^\delta(\pi)$ of player i. The strategy profile σ is a **pure strategy subgame perfect Nash equilibrium** of $G(\delta)$ if for all $t \geq 1$ and history $h^t \in S^{t-1}$, the restriction $\sigma_{|h^t}$ of the strategy profile σ to the observed history h^t is a pure strategy Nash equilibrium of the game $G(\delta)$.

As argued in the introduction, remaining silent on a punishment path can be more severe than employing minimax strategies. A target player, if she is ambiguity averse, best responds to the silence of her opponents by playing a prudent strategy, aiming to secure her maxmin payoff. Theorem 7 shows that, in any pure strategy subgame perfect equilibrium of the infinitely repeated game with discounting, each player receives at least her effective maxmin payoff.

Theorem 7 *For all $i \in N$ and $\delta \in (0, 1)$, player i's average equilibrium payoffs in $G(\delta)$ are not less than her effective maxmin payoff.*

Proof of Theorem 7. To start, normalize the payoffs of the stage game such that $u_i = u_j$ if $i \in \widetilde{j}$ and the effective maxmin payoff of each player equals 0. Let $\delta \in (0,1)$ and σ be a pure strategy subgame perfect Nash equilibrium of $G(\delta)$. Assume that there exists $i \in N$ such that $u_i^\delta(\sigma) < 0$. Let $j \in \widetilde{i}$ be a player whose maxmin payoff is equal to 0. Let σ'_j be the pure strategy of player j in the infinitely repeated game $G(\delta)$ that consists in playing her stage game prudent strategy in every period. In each point of time, player j receives a positive payoff at (σ_{-j}, σ'_j). Therefore, we have $u_i^\delta(\sigma_{-j}, \sigma'_j) \geq 0$. Furthermore, $u_j^\delta(\sigma) = u_i^\delta(\sigma) < 0$. The deviation σ'_j is thus profitable. A contradiction holds as σ is a subgame perfect equilibrium.

Theorem 8 shows that the set of pure strategy subgame perfect equilibrium payoff vectors of the infinitely repeated game includes any feasible payoff vector that strictly dominates the effective maxmin payoff vector.

Theorem 8 *For any $x \in V^*$, there exists $\underline{\delta} \in (0,1)$ such that for all $\delta \in (\underline{\delta}, 1)$, $G(\delta)$ has a subgame perfect Nash equilibrium strategy with deterministic play path and with average payoff x.*

The proof of this theorem is constructive and employs the payoff continuation lemma of Fudenberg and Maskin (1991). I recall it below.

Lemma 27 *(Lemma 2 of Fudenberg and Maskin (1991)) For any $\varepsilon > 0$, there exists $\underline{\delta} < 1$ such that for all $\delta \geq \underline{\delta}$ and every $v \in V^*$ with $v_i \geq \varepsilon$ for all i, there is a deterministic sequence of pure strategies whose average discounted payoffs are v, and whose continuation payoffs at each time are within ε of v.*

Proof of Theorem 8. Let $x \in V^*$. Apply the payoff asymetry lemma of Abreu et al. (1994) and obtain a payoff asymetry family $(y^i)_{i\in N}$ of elements of V^* such that for all $i \in N$, player i prefers x to y^i and any player $j \notin \widetilde{i}$ prefers y^i to y^j.

Now choose $\varepsilon > 0$ such that

- $2\varepsilon < y_i^i$,
- $B(x, \varepsilon) \subseteq V^*$,
- $B(y^i, \varepsilon) \subseteq V^*$ for all $i \in N$ and
- $\frac{2\rho}{2\rho + (y_i^j - y_i^i)} < \frac{y_i^i - 2\varepsilon}{y_i^i}$ for all $i, j \in N$ such that $i \notin \widetilde{j}$,

- $\widetilde{x}_i > \widetilde{y}^i_i$ for all $\widetilde{x} \in B(x, \varepsilon)$, $\widetilde{y}^i \in B(y^i, \varepsilon)$ and $i \in N$,

- $\widetilde{y}^j_i > \widetilde{y}^i_i$ for all $\widetilde{y}^i \in B(y^i, \varepsilon)$, $\widetilde{y}^j \in B(y^j, \varepsilon)$ and $i, j \in N$ such that $i \notin j$

where for all $x' \in V$ and $\varepsilon' > 0, B(x', \varepsilon') = \{x" \in V \mid \|x" - x'\| < \varepsilon'\}$ and $\|x'\| = max_{i \in N}|x'_i|$.

Apply Lemma 27 to ε and obtain a discount factor threshold δ_ε, a deterministic path $\pi = \{s(t)\}$ and for all $i \in N$ a deterministic path $\pi^i = \{s^i(t)\}$ such that

- the discounted average $u^\delta(\pi)$ of the path π is equal to x $\forall \delta \in (\delta_\varepsilon, 1)$,

- for all $i \in N$, the discounted average $u^\delta(\pi^i)$ of the path π^i is equal to y^i $\forall \delta \in (\delta_\varepsilon, 1)$,

- at each point of time, and for all $\delta \in (\delta_\varepsilon, 1)$, the continuation payoffs of the sequences $\pi, \pi^1, ..., \pi^n$ are respectively within ε of $x, y^1, ..., y^n$.

Choose $\delta_1 \in (\delta_\varepsilon, 1)$ and T_1 such that

$$\rho(1 - \delta_1) < \varepsilon$$

and

$$\tfrac{2\rho}{2\rho + (y^j_i - y^i_i)} < \delta_1^{T_1+2} < \delta_1^{T_1} < \tfrac{y^i_i - 2\varepsilon}{y^i_i}.$$

Those choices are possible as the set $\{\delta^t \mid \delta \in (\delta_\varepsilon, 1) \text{ and } t > 0\}$ is dense in $[0, 1]$. There exists $\underline{\delta} \in (\delta_1, 1)$, such that for all $\delta \in (\delta_2, 1)$, there exists T such that the following inequalities holds.

- $\rho(1 - \delta) + \delta^t y^i_i < -\rho(1 - \delta^t) + \delta^t y^j_i$ for all $t \leq T$ and for all $i, j \in N$ such that $i \notin \widetilde{j}$, (1)

- $\rho(1 - \delta) + \delta^{T+1} y^i_i < y^i_i - \varepsilon$ for all $i \in N$. (2)

Let $\delta \in (\underline{\delta}, 1)$. For any player $i \in N$, consider the strategy σ_i in the game $G(\delta)$ defined by the 4 following phases.

1) Main path (P_0): Play $a_i(1)$ in the first period and continue to follow the path $\{a_i(t)\}$ till a unique player deviates. If player j deviates, go to phase P_j.

2) Punish player j (P_j): Remain silent for T periods. If a unique player j' break the silence, then, set $j = j'$ and restart phase P_j. Go to phase R_j.

3) Reward players of the block $N\backslash\{j\}$ (R_j): Start and follow the path $\{a_i^j(t)\}$ till a unique player deviates. If player k deviates, set $j = k$ and restart phase P_j.

I now show that the strategy profile $\sigma = (\sigma_1, ..., \sigma_n)$ is a pure strategy subgame perfect equilibrium of $G(\delta)$. Precisely, I show that, no pure strategy one shot deviation from any phase leads to a higher continuation payoff.

a) If player i deviates during the phase P_0, as she is ambiguity averse, she receives at most $\rho(1-\delta) + \delta^T y_i^i$ as continuation payoff. If she does not deviate, she gets at least $x_i - \varepsilon$. From (2), such deviation is not profitable.

b) If player i deviates during the phase P_j where $i \notin \widetilde{j}$, she receives at most $\rho(1-\delta) + \delta[0.(1-\delta^T) + \delta^T y_i^i]$ as continuation payoff. If she does not deviate, she gets at least $-\rho(1-\delta^t) + \delta^t y_i^j$ for some $t \leq T$. From (1), such deviation is not profitable.

c) If player i deviates during the phase P_j where $i \in \widetilde{j}$, she receives at most $0.(1-\delta) + \delta[\underline{x}.(1-\delta^T) + \delta^T y_i^i]$ as continuation payoff, where $\underline{x} = u_i(\Delta S)$. If she does not deviate, she gets at least $\underline{x}(1-\delta^t) + \delta^t y_i^i$ for some $t \leq T$. As $t \leq T$ and $\underline{x} \leq 0$, we have that

$$\delta[\underline{x}.(1-\delta^T) + \delta^T y_i^i] < \underline{x}(1-\delta^t) + \delta^t y_i^i$$

and the deviation is not profitable.

d) If player i deviates during the phase R_j, as she is ambiguity averse, she receives at most $\rho(1-\delta) + \delta^T y_i^i$ as continuation payoff. If she does not deviate, she receives at least $y_i^i - \varepsilon$. From (2), such deviation is not profitable.

The strategy profile σ is therefore a pure strategy subgame perfect equilibrium of $G(\delta)$ and the associated payoff is x.

5.4 Conclusion

Theorem 8 shows that, if players are allowed to conceal their actions in imprecise devices, then the set of pure strategy subgame perfect equilibrium payoff vectors of the infinitely repeated game with discounting will includes the set of feasible payoff vectors that dominate the effective maxmin payoff vector of the stage-game, a set which is larger (and even strictly larger for a non degenerated set of games with at least three players) than the set of equilibrium payoff vectors predicted by the classic model (see Wen (1994) for a general result with the classic model).

Bibliography

Dilip Abreu, David Pearce, and Ennio Stacchetti. Toward a theory of discounted repeated games with imperfect monitoring. Econometrica: Journal of the Econometric Society, pages 1041–1063, 1990.

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

Robert J Aumann and Sylvain Sorin. Cooperation and bounded recall. Games and Economic Behavior, 1(1):5–39, 1989.

Robert J Aumann, Michael Maschler, and Richard E Stearns. Repeated games with incomplete information. MIT press, 1995.

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.

Daniel Ellsberg. Risk, ambiguity, and the savage axioms. The quarterly journal of economics, 75(4): 643–669, 1961.

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

Drew Fudenberg and Eric Maskin. On the dispensability of public randomization in discounted repeated games. Journal of Economic Theory, 53(2): 428–438, 1991.

Itzhak Gilboa and David Schmeidler. Maxmin expected utility with non-unique prior. Journal of mathematical economics, 18(2): 141–153, 1989.

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.

Joseph Greenberg. The right to remain silent. Theory and Decision, 48(2): 193–204, 2000.

David M Kreps and Robert Wilson. Reputation and imperfect information. Journal of economic theory, 27(2):253–279, 1982.

David M Kreps, Paul Milgrom, John Roberts, and Robert Wilson. Rational cooperation in the finitely repeated prisoners' dilemma. Journal of Economic theory, 27(2):245–252, 1982.

Jamie Brown Kruse, Stephen Rassenti, Stanley S Reynolds, and Vernon L Smith. Bertrand-edgeworth competition in experimental markets. Econometrica: Journal of the Econometric Society, pages 343–371, 1994.

George J Mailath, Steven A Matthews, and Tadashi Sekiguchi. Private strategies in finitely repeated games with imperfect public monitoring. Contributions in Theoretical Economics, 2(1), 2002.

John Nash. Non-cooperative games. Annals of mathematics, pages 286–295, 1951.

Abraham Neyman. Bounded complexity justifies cooperation in the finitely repeated prisoners' dilemma. Economics letters, 19(3):227–229, 1985.

Frank Riedel. Uncertain acts in games. Homo Oeconomicus, 34(4): 275–292, 2017.

Frank Riedel and Linda Sass. Ellsberg games. Theory and Decision, 76(4): 469–509, 2014.

Tadashi Sekiguchi. Existence of nontrivial equilibria in repeated games with imperfect private monitoring. Games and Economic Behavior, 40(2): 299–321, 2002.

Reinhard Selten. Reexamination of the perfectness concept for equilibrium points in extensive games. International journal of game theory, 4(1): 25–55, 1975.

Hugh Sibly and John Tisdell. Cooperation and turn taking in finitely-repeated prisoners' dilemmas: An experimental analysis. Journal of Economic Psychology, 2017.

Lones Smith. Necessary and sufficient conditions for the perfect finite horizon folk theorem. MIT Working Paper, 93(6), 1993.

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.

Short cv

Education

2012 DU (equivalent to a M.A.) in Applied Statistics, University of Toulouse, France, FOAD

2010 Master of Sciences in Mathematics, University of Yaound I, Cameroon

2009 M.A. in Education, Higher Teachers Training College, University of Yaound I, Cameroon

Research Fields

Game theory, Strategic ambiguity, and Political economy

Invited Research Visits

University of Ottawa (Canada), with Roland Pongou and Victoria Barham, August-September 2018

Ecole Polytechnique de Paris(France), with Olivier Gossner, Febuary 2018

Karl Franzens Universitt Graz (Austria), with Christoph Kuzmics, December 2017

University of Yaound I (Cameroon), with Moyouou Issofa, March-April 2017

University of Ottawa (Canada), with Roland Pongou, September-October 2016

Ecole Polytechnique de Paris (France), with Olivier Gossner, September 2015

Honors and Awards

DAAD Scholarship within the GSSP Program (Sept 2013 - June 2017).

BGTS Mobility Grants (2016 and 2018)

Excellence Award of the President of the Republic, Cameroon (2009 and 2010)

Selected Working Papers

A complete folk theorem for finitely repeated games (2018)

A note on Necessary and sufficient condition for the perfect finite horizon folk theorem; [Econometrica, 63 (2): 425-430, 1995]" (2018)

Finitely repeated games and strategic ambiguity (2018)