Center for
Mathematical Economics
Working Papers

585

July 2018

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

Ghislain H. Demeze-Jouatsa

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

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

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

Ghislain H. DEMEZE-JOUATSA [1] [2]

This version: July 31, 2018

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

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 birth of cooperation in a larger class of normal form games. This class includes some stage-games

[1]Bielefeld University. Email: demeze_jouatsa@uni-bielefeld.de

[2]I wish to acknowledge the DAAD for funding this research project and thank Frank Riedel, Christoph Kuzmics, Patrick Beissner, Igor Muraviev, Mira Frick, and Francesc Dilm for their insightful comments. I also thank the seminar participants at the University of Bielefeld, the University of Douala and University of Yaound I.

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 birth 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 an example of a game in which the classic model of finitely repeated games with pure and mixed strategies can not explain the birth of cooperation while the introduction of an infinitesimal level of ambiguity in the model allows for sustaining cooperation. Section 3 presents the model as well as some preliminary results. The main theorem of the paper is presented and discussed in Section 4. Section 5 provides the proofs.

2 The benefit of the ambiguity

Allowing players to play Ellsberg actions make the set of stage-game action profiles larger and less tractable. However, this can allow to explain the birth of cooperation when the classic model of finitely repeated games with pure and mixed actions fails to do so. In this section, we use a three-player normal form game to illustrate that adding an infinitesimal level of ambiguity to the model of finitely repeated game can allow to explain the birth of cooperation in finite time horizon.

Consider the three-player normal form game G whose payoff matrix is given by table 1 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: 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 birth 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.

3 The Model

3.1 The stage-game

3.1.1 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**[3] 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 the set of feasible payoff vectors that strictly dominates the effective maxmin payoff vector ν.

3.1.2 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, $u_i(P'_i, P_{-i}) \leq u_i(P)$.

[3]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.

3.1.3 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 1 *In the Ellsberg game, remaining silent is an optimal punishment strategy.*

Proof. of Lemma 1. 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 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 2: 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

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

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

Lemma 2 *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 2. From Lemma 1, 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. ■

3.1.4 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 1 *The sequence $N_0 = \emptyset \subseteq N_1 \subseteq ... \subseteq N_h$ is the Ellsberg decomposition of the game G.*

Definition 2 *The Ellsberg decomposition $N_0 = \emptyset \subseteq N_1 \subseteq ... \subseteq N_h$ is complete if $N_h = N$.*

3.2 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 $\sigma_i : \cup_{t=1}^{T} \mathcal{P}^{t-1} \to \mathcal{P}_i$ 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 3 *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 4 *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 Main result and discussion

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

Definition 3 *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 4 *The support of a strategy profile of the finitely repeated game is the set of possible play paths.*

Definition 5 *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 $P(\sigma)$.*[4]

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.

4.1 Statement of the main result

Theorem 1 *Let G be a finite normal form game such that $V^* \neq \emptyset$. The following are equivalent.*

1. G has a complete Ellsberg decomposition.

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

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.*

The most laborious part of the proof of Theorem 1 is to show that, under the statement 1) of Theorem 1, 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 6 and Lemma 7. 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 7 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 5. 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 6.1, 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.

4.2 Discussion

While both necessary and sufficient, Statement 1) of Theorem 1 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 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 2 for details).

For the game in Table 1, 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 birth 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 1 and Lemma 4 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 3.

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

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

lent 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.

5 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 birth of cooperation.

6 Appendix 4: Proofs

6.1 Sketch of the proof of Theorem 1

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 5 (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 6.4 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 1. To ease this proof, I introduce three lemmata.

6.2 The payoff continuation lemma for finitely repeated games

Lemma 5 *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 5 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 5. 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 provides the value of $b^{y,p}$ for $p = 1, ..., 7$ while Table 5 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: 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 5: 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 $p'_2 = \min\{\max\{0, B - p_1\}, p_2\}$;

.

.

.

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

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

As

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

and

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

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 $\delta > \delta^{y,A}$,

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

for all B, $0 \leq B < Q$. Let $\{\widetilde{B}(y, \frac{\varepsilon}{2}), y \in Y\}$[5] 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.

6.3 The end-game-strategy

Lemma 6 *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. *$u_i^\delta(\sigma^{i,1}) > \phi_e + u_i^\delta(\sigma^{i,2})$.*

Lemma 7 *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*

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

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

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

The proofs of Lemmata 6 and 7 are provided in a more general case in a parallel working paper. We therefore omit them.

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

6.4 Proof of the Theorem 1

Proof of Theorem 1. 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 5) 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 7 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

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

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 7 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 $\phi,\ \theta^1,\ q_1,\ q_2,\ r,\ r_1$ and θ^p are such that

$$-2q_1 k\rho + r\phi > 0; \tag{4}$$

$$\overline{y}_i - \frac{r}{2^p} > 0; \tag{5}$$

$$2k\rho + q_1 k u_i(\Delta S) + q_2 r_1 k u_i^{1,r_1}[\theta^1(1,\cdots,1)] + \frac{r}{2^p} - (q_1 k + q_2 r_1 k - k)\overline{y}_i < 0 \tag{6}$$

and

$$-2(q_1 k + q_2 r_1 k)\rho + r\phi > 0 \text{ for all } i \in N. \tag{7}$$

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.

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 **P**$_i$; if player i deviates late, start **LD**. Ignore any simultaneous deviation.] Go to Phase **EG**.

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 **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_2k 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 $(C+q_1+q_2r_1)k$. At period $(C+q_1+q_2r_1)k$, 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. $-\frac{1-\delta^{q_1k}}{1-\delta}\rho$ during the punishment phase;
2. $\delta^{q_1k}\frac{1-\delta^{q_2r_1k}}{1-\delta}u_j^\delta(\theta^1(1,\cdots,1))$ during the **SPE** phase;
3. some payoff $U_j(\delta)$ up to the period $(C+q_1+q_2r_1)k$;
4. an ex-ante payoff $\delta^{c+(q_1+q_2r_1)k}\frac{1-\delta^r}{1-\delta}u_j^\delta[\theta^p(1,\alpha_j)]$ in the Phase **EG**.

In total, she gets

$$-\frac{1-\delta^{q_1k}}{1-\delta}\rho+\delta^{q_1k}\frac{1-\delta^{q_2r_1k}}{1-\delta}u_j^\delta(\theta^1(1,\cdots,1))+U_j(\delta)+\delta^{c+(q_1+q_2r_1)k}\frac{1-\delta^r}{1-\delta}u_j^\delta[\theta^p(1,\alpha_j)]$$

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

1. $\frac{1-\delta^{q_1k}}{1-\delta}\rho$ during the punishment phase;
2. $\delta^{q_1k}\ \frac{1-\delta^{q_2r_1k}}{1-\delta}u_j^\delta(\theta^1(1,\cdots,1))$ 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 7.

In total, she gets

$$\frac{1-\delta^{q_1k}}{1-\delta}\rho+\delta^{q_1k}\,\frac{1-\delta^{q_2r_1k}}{1-\delta}u_j^{\delta}(\theta^1(1,\cdots,1))+U_j(\delta)+\delta^{c+(q_1+q_2r_1)k}\frac{1-\delta^r}{1-\delta}\left[u_j^{\delta}[\theta^p(1,\alpha_j)]-\phi\right]$$

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

$$-2\frac{1-\delta^{q_1k}}{1-\delta}\rho+\delta^{c+(q_1+q_2r_1)k}\frac{1-\delta^r}{1-\delta}\phi>0 \tag{8}$$

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). Therefore, there exists $\underline{\delta}_1\in(\underline{\delta}_0,1)$ such that the Inequality (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 $(C+q_1+q_2r_1)k$;[6]
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 7.

As $\overline{y}_i-\frac{r}{2^p}>0$ [see Equation (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.

[6] 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.

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;[7]
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 $(C+q_1+q_2r_1)k$;[8]
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 7, 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 (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

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

[8] Note that $\lim_{\delta\to1}\widetilde{U}_i^2(\delta)=\lim_{\delta\to1}U_i^2(\delta)$.

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**;

2. the ex-ante payoff $\delta^{q_1k+q_2r_1k}\frac{1-\delta^r}{1-\delta}u_j^\delta[\theta^p(\alpha)]$ in the Phase **EG**, where $\alpha \in \{0.1\}^n$. From Lemma 7, the latter ex-ante payoff is strictly greater than $\delta^{q_1k+q_2r_1k}\frac{1-\delta^r}{1-\delta}\left(u_j^\delta[\theta^p(-1,\cdots,-1)]+\phi\right)$.

As $-2(q_1k+q_2r_1k)\rho + r\phi > 0$ [see Equation (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

$$\frac{1-\delta^k}{1-\delta^{(\underline{C}+q_1+q_2r_1)k+r}}\rho + \delta^{(\underline{C}+q_1+q_2r_1)k}\frac{1-\delta^r}{1-\delta^{(\underline{C}+q_1+q_2r_1)k+r}}\rho < \varepsilon$$

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 1 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 1.

■

References

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. 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.

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.

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.

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. *Econometrica*, 63(2): 425–430, 1995.

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