

Quitting Games

Eilon Solan[†] and Nicolas Vieille[‡]

January 4, 2001

Abstract

Quitting games are n -player sequential games in which, at any stage, each player has the choice between continuing and quitting. The game ends as soon as at least one player chooses to quit; player i then receives a payoff r_S^i , which depends on the set S of players that did choose to quit. If the game never ends, the payoff to each player is 0.

The paper has four goals. (i) we prove the existence of a subgame perfect uniform ε -equilibrium, under some assumptions on the payoff structure, (ii) we study the structure of the ε -equilibrium strategies, (iii) we present a new method for dealing with n -player games, and (iv) we study an example of a four-player quitting game where the “simplest” equilibrium is cyclic with period 2.

We also discuss the relation to Dynkin’s stopping games, and provide a generalization of our result to these games.

Keywords: Quitting games, uniform equilibrium, Dynkin’s stopping games, n -player stochastic games.

[†] MEDS Department, Kellogg Graduate School of Management, Northwestern University, *and* the School of Mathematical Sciences, Tel Aviv University,

Tel Aviv 69978, Israel.

e-mail: eilons@post.tau.ac.il

Laboratoire d'Econométrie, Ecole Polytechnique *and* Département Finance
et Economie, HEC, 1, rue de la Libération, 78 Jouy-en-Josas, France.

Tel: +33-1-46-34-30-02

e-mail: vielle@poly.polytechnique.fr

The first author thanks the hospitality of the Center for Rationality and Interactive Decision Theory and of the Laboratoire d'Econométrie de l'Ecole Polytechnique, and the financial support of the Institute of Mathematics at the Hebrew University of Jerusalem. We have benefited from several conversations with Ehud Kalai.

Introduction

Quitting games are sequential games in which, at any stage, each player has the choice between continuing and quitting. The game ends as soon as at least one player chooses to quit; player i then receives a payoff r_S^i , which depends on the set S of players that did choose to quit. If the game never ends, the payoff to each player is 0.

In the present paper we study subgame perfect uniform ε -equilibria in these games. In the case of two players, stationary ε -equilibria do exist. A three-player example was devised by Flesch, Thuijsman and Vrieze (1997), where ε -equilibrium strategies are more complex - they have a cyclic structure, and the length of the cycle is at least 3. This gave the impetus to the study of the three-player case, solved by Solan (1999) (for a more general class of games).

The paper has four goals:

- We prove the existence of a subgame perfect uniform ε -equilibrium strategies under some assumptions on the payoff function.
- We study the structure of these ε -equilibrium strategies. We show that there always exist cyclic ε -equilibrium strategies; that is, the players repeat the same behavior over and over again until the game ends.
- We introduce a new method of analyzing n -player games, that does not use fixed point arguments, in contrast to the methods that appear in the literature.
- We study an instance of a four-player quitting game, and show that in this game, the “simplest” equilibrium is cyclic with a period 2: there

is neither a stationary ε -equilibrium nor an ε -equilibrium where the players play, up to ε , the same mixed action combination every stage.

When the game is symmetric, that is (i) the payoff depends on the number of players that quit, (ii) all the players who quit receive the same payoff and (iii) all the players who continue receive the same payoff, we get a stronger result. There is always a pure stationary equilibrium. However, a symmetric ε -equilibrium, i.e. one in which all players follow the same strategy, does not necessarily exist.

Quitting games are a variant of the popular attrition models, first introduced in evolutionary biology by Maynard-Smith (1974), also used in auction theory (Krishna and Morgan (1997)), entry deterrence models (we refer to Hammerstein and Selten (1995), Wilson (1992) and Fudenberg and Tirole (1991) for references), or in the analysis of strategic exit (see Ghemawat and Nalebuff (1985) or Li (1989)).

One difference between the literature dealing with the war of attrition and our work is that attrition models are usually continuous-time models, in which strategic interaction lasts as long as two players at least did not quit. The major departure point is that papers on attrition models have mostly dealt with incomplete information situations, and focused on the existence and the analysis of equilibria for a given discounting function, whereas we are interested in uniform ε -equilibria.

Quitting games are also deeply related to Dynkin's stopping games. The latter are two-player, zero-sum games, where the players choose stopping times τ_1 and τ_2 , and the payoff is given by $X_{\tau_1} 1_{\tau_1 < \tau_2} + Y_{\tau_2} 1_{\tau_1 > \tau_2} + W_{\tau_1} 1_{\tau_1 = \tau_2}$, where $(X_n), (Y_n), (W_n)$ are processes. Dynkin (1969) gave sufficient conditions on these processes for the game to have a value. Subsequently, some classes of two-player, non zero-sum games were analyzed (Morimoto (1986)),

Ohtsubo (1987)). Thus, we deal here with constant payoff processes, and allow for randomized stopping times. This enables us to deal with any number of players, and different sets of assumptions on payoffs.

Another interest in our result is from the point of view of stochastic games. Quitting games form a class of n -player stochastic games, where existence of equilibrium payoff in general is still an open problem. Our result is the first result for n -player games where current methods fail to work, and the method that we use may be useful in the general case as well.

We assume that the payoff function satisfies two conditions: (i) each player prefers to be the only quitter, rather than to have the game continue forever, and (ii) if any player decides to quit, he prefers to be the only quitter, rather to have more players quit with him. The first assumption is not too restrictive, but the second assumption is.

The main interest in studying the uniform (or undiscounted) equilibrium payoff is the robustness that it offers: a uniform ε -equilibrium is a 2ε -equilibrium in any discounted game with discount factor sufficiently close to 1, and in any finite game which is sufficiently long. Recall that usually a strategy that is good for one discount factor, may yield a very low payoff if the discount factor is slightly changed. Sorin (1986) showed that even if for every discount factor a game possesses the same discounted equilibrium payoff, this payoff needs not be a uniform equilibrium payoff.

Our method to prove the result is significantly different from methods used in the literature. Whereas usually one uses a fixed point argument to show existence in the discounted game, or takes a limit of discounted equilibria as the discount factor tends to 1 to show the existence of the uniform equilibria for 2 and 3 player games, we look for a cycle, or a periodic point, in the correspondence that assigns to each possible continuation payoff

the set of all Nash equilibria in the corresponding one-shot game.

The paper is arranged as follows. In section 1 we set up the model and state the main result. In section 2 we prove the main result and provide a generalization for stopping games. Finally, in section 3 we study an example that shows that our result is sharp.

1 The Model and the Main Results

A quitting game is a pair $(\mathbf{N}, (r_S)_{\emptyset \subset S \subseteq \mathbf{N}})$, where (i) $\mathbf{N} = \{1, \dots, N\}$ is a finite set of players, and (ii) for every $S \subset \mathbf{N}$, $r_S \in \mathbf{R}^{\mathbf{N}}$.

The game is a sequential game, that is played as follows. The set of stages is the set \mathbf{N} of positive integers. At every stage each player chooses an action, either continue or quit. Let S be the subset of the players who choose to quit. If $S = \mathbf{N}$, then the game terminates, and each player i receives the payoff r_S^i . If $S \neq \mathbf{N}$, the game continues to the next stage. If the game never terminates, each player gets 0.

We denote the two actions of player i by $\{c^i, q^i\}$. A strategy for player i is a function $\mathbf{x}^i = (x_n^i) : \mathbf{N} \rightarrow [0, 1]$, x_n^i being the probability that player i continues at stage n , provided the game has not terminated before. If $x_n^i = 1$ then at stage n player i plays the pure action c^i , or continue, while if $x_n^i = 0$ then at stage n player i plays the pure action q^i , or quit. In particular, \mathbf{c}^i is the strategy of player i where he always continues, and \mathbf{q}^i is the strategy where he always quits.

For every stage $n \in \mathbf{N}$, S_n is the set of players that quit at that stage, and a_n is the action combination that is played.

A profile is a vector of strategies, one for each player. A profile $\mathbf{x} = (x_n)_{n \in \mathbf{N}}$ induces a probability distribution $\mathbf{P}_{\mathbf{x}}$ over the set of plays. We

denote by $\mathbf{E}_{\mathbf{x}}$ the corresponding expectation operator. If the players abide by \mathbf{x} , their expected payoff in the game is given by

$$\gamma(\mathbf{x}) = \mathbf{E}_{\mathbf{x}}(r_{S_t} \mathbf{1}_{t < +\infty}),$$

where $t = \inf\{n, S_n = \cdot\}$ is the termination stage. For each n , we denote by $\mathbf{x}_n = (x_n, x_{n+1}, \dots)$ the profile induced by \mathbf{x} in the subgame starting at stage n .

We say that the profile \mathbf{x} is terminating if $\mathbf{P}_{\mathbf{x}}(t < +\infty) = 1$: that is, if the players follow \mathbf{x} then with probability 1 eventually some players quit. This is equivalent to

$$\prod_{n \in \mathbf{N}} \prod_{i \in \mathcal{N}} x_n^i = 0.$$

We say that \mathbf{x} is cyclic if there exists $n_0 \in \mathbf{N}$ such that $x_n = x_{n+n_0}$ for every $n \in \mathbf{N}$, and that \mathbf{x} is stationary if $x_n = x_1$ for each $n \in \mathbf{N}$.

As usual, \mathbf{x}^{-i} stands for $(\mathbf{x}^j)_{j \neq i}$. We shall abbreviate similarly whenever convenient.

DEFINITION 1.1 *A profile \mathbf{x} is an ε -equilibrium if for every player i and every strategy \mathbf{y}^i of player i ,*

$$\gamma^i(\mathbf{x}) - \gamma^i(\mathbf{x}^{-i}, \mathbf{y}^i) \geq -\varepsilon.$$

It is a subgame perfect ε -equilibrium if for every $n \in \mathbf{N}$ the profile \mathbf{x}_n is an ε -equilibrium. The corresponding payoff vector $\gamma(\mathbf{x})$ is an ε -equilibrium payoff.

Our main result is:

THEOREM 1.2 *Let $\varepsilon > 0$. Every quitting game that satisfies:*

A.1 $r_{\{i\}}^i = 1$ for every $i \in \mathcal{N}$.

A.2 $r_S^i \geq 1$ for every S such that $i \in S$.

has a cyclic subgame perfect ε -equilibrium.

Assumption **A.1** essentially claims that any player prefers his unilateral termination to indefinite continuation. Assumption **A.2** is somewhat restrictive and can be partially weakened (see Proposition 2.2). It claims that if some player i decides to quit, then he cannot profit if some other players also quit.¹

Some of the quitting games that arise from economic situations are symmetric in the sense that the payoff r_S^i to i only depends on the size of S and on whether or not i belongs to S . We get a strengthened version of Theorem 1.2 for such games.

THEOREM 1.3 *Every symmetric quitting game has a pure, stationary (0)-equilibrium.*

Such an equilibrium is of course subgame perfect. We prove this result in subsection 1.1. We also provide an example that shows that a symmetric quitting game needs not admit a symmetric ε -equilibrium; that is, a profile \mathbf{x} such that $\mathbf{x}^i = \mathbf{x}^j$ for every $i, j \in N$.

In the literature on stochastic games, a stronger concept of uniform equilibrium is often used. In section 2.6 we prove that in quitting games, any ε -equilibrium is a uniform ε -equilibrium.

The above result asserts the existence of cyclic ε -equilibria in a class of quitting games. For two- or three-player games, “better” results are available. For two-player games, stationary ε -equilibria exist. It is a consequence

¹Janos Flesch commented that the following weaker assumption can replace assumptions **A.1** and **A.2**: for every player i and every set S such that $i \in S$, $r_{\{i\}}^i \geq r_S^i$.

of more general results of Flesch, Thuijsman and Vrieze (1996) or Vieille (1992). For three-player games, it follows from Solan (1999) that either a stationary ε -equilibrium exists, or there exists a cyclic ε -equilibrium in which the probability of termination in any given stage is arbitrarily small (moreover, there is an ε -equilibrium payoff in $\text{co}\{r_{\{i\}}, i \in N\}$), or both. In section 3 we show on an example that this is no longer true for 4-player games. In that sense, our result is optimal.

1.1 Symmetric quitting games

1.1.1 Pure stationary 0-equilibria

The payoff function for a symmetric quitting game is characterized by numbers $(a_k)_{k=1,\dots,N}$ and $(b_k)_{k=1,\dots,N-1}$, where a_k (resp. b_k) is the payoff to a quitting (resp. non-quitting) player if the quitting coalition is of size k : $r_S^i = a_k$ if $|S| = k$ and $i \in S$ whereas $r_S^i = b_k$ if $|S| = k$ and $i \notin S$.

It is convenient to set $b_0 = 0$. Notice first that any stationary 0-equilibrium is a subgame perfect 0-equilibrium. We distinguish three cases.

Case 1: $a_1 = 0$.

The profile \mathbf{c} (every player continues in every stage) is a stationary equilibrium in which every player gets 0. Indeed, any (pure) deviation results in a payoff of a_1 to the deviating player.

Case 2: For some $k \in \{1, \dots, N-1\}$, $a_{k+1} = b_k$ and $a_k = b_{k-1}$.

With the profile $(\mathbf{q}^1, \dots, \mathbf{q}^k, \mathbf{c}^{k+1}, \mathbf{c}^N)$, players $i, i \leq k$, get a_k , and players $i, i > k$, get b_k . We claim that this profile is a stationary equilibrium. Indeed, for every pure deviation of player $i \leq k$, player i receives $b_{k-1} = a_k$. For every

pure deviation of player $i > k$, player i gets $a_{k+1} > b_k$.

Case 3: If neither case 1 nor case 2 hold, then $a_1 > 0 = b_0$ and, for every k ,

$$a_k > b_{k-1} \quad a_{k+1} > b_k.$$

This readily implies that $a_N > b_{N-1}$. Therefore, the profile \mathbf{q} (every player quits at every stage) is an equilibrium.

1.1.2 Non-existence of a symmetric equilibrium

We now show that a symmetric game does not necessarily admit a symmetric ε -equilibrium, even under assumptions **A.1** and **A.2**. That is, there needs not exist an ε -equilibrium \mathbf{x} such that $\mathbf{x}^i = \mathbf{x}^j$ for every $i, j \in N$.

Example 1: Consider the following two-player quitting game, where player 1 is the row player and player 2 is the column player:

		2	
		c^2	q^2
1	c^1	<i>continue</i>	2, 1
	q^1	1, 2	1, 1

Note that $a_1 = a_2 = 1$ and $b_1 = 2$, hence $(\mathbf{q}^1, \mathbf{c}^2)$ is a pure stationary equilibrium, as is $(\mathbf{q}^2, \mathbf{c}^1)$.

Assume that \mathbf{x} is a symmetric ε -equilibrium, where ε satisfies $\varepsilon + 2\bar{\varepsilon} < 1/2$. It is clear that $\mathbf{P}_{\mathbf{x}}(t < +\infty) = 1 - \varepsilon$. Indeed, player i could otherwise play \mathbf{x}^i until some distant stage n , and then quit. If the residual probability of termination after stage n is small enough, this strategy improves upon \mathbf{x}^i by more than ε .

Since \mathbf{x} is symmetric, $\gamma^1(\mathbf{x}) = \gamma^2(\mathbf{x})$. Moreover, $\varepsilon = \mathbf{P}_{\mathbf{x}}(t = +\infty) = \prod_{n \in \mathbf{N}} (x_n^1)^2$, hence

$$\mathbf{P}_{\mathbf{x}^1, \mathbf{c}^2}(t = +\infty) = \prod_{n \in \mathbf{N}} x_n^1 = \bar{\varepsilon}.$$

Since $r_S^1 + r_S^2 = 3$ for each S , one has $\gamma^i(\mathbf{x}) = 3/2$, for each player i . Finally, $\gamma^2(\mathbf{x}^1, \mathbf{c}^2) = 2\mathbf{P}_{\mathbf{x}^1, \mathbf{c}^2}(t < +\infty) = 2(1 - \bar{\varepsilon}) > 3/2 + \varepsilon$, where the last inequality follows by the choice of ε . However, this implies that \mathbf{x} is not an ε -equilibrium — a contradiction.

2 Existence Result

This section is devoted to the proof of Theorem 1.2. It is organized as follows. For every $w \in \mathbf{R}^N$, we define in section 2.1 an associated one-shot game $G(w)$, in which player i receives w^i if termination does not occur (in one stage). Thus, w should be interpreted as a continuation payoff. We define an ad hoc refinement of ε -equilibrium, which we call perfect ε -equilibrium.

In section 2.2, we show that there exists a profile \mathbf{x} such that, for every n , \mathbf{x}_n is terminating and x_n is a perfect ε -equilibrium in $G(\gamma(\mathbf{x}_{n+1}))$.

In section 2.5 we prove that existence of such a profile implies that either \mathbf{x} is a subgame perfect $\varepsilon^{\frac{1}{6}}$ -equilibrium, or that there exists a stationary $\varepsilon^{\frac{1}{6}}$ -equilibrium, or both.

2.1 The one-shot game

Fix a quitting game $G = (\mathbf{N}, (r_S)_{\emptyset \subset S \subset \mathbf{N}})$. Let $\rho = 2 \max\{|r_S^i| \mid i \in \mathbf{N}, S \subset \mathbf{N}\}$ be twice the maximal payoff in absolute values.

For every $w \in \mathbf{R}^N$ we define a one-shot game $G(w)$ as follows. Each player has two possible actions, continue and quit. Let S be the subset of

players that chose to quit. If $S = \emptyset$, the players receive the payoff w , and otherwise they receive the payoff r_S .

A profile in $G(w)$ is a vector $x \in [0, 1]^N$, x^i being the probability that player i chooses continue. In particular, c is the profile where everyone continues, and $c^{-i} = (c^j)_{j \neq i}$. With every profile x we associate the probability of termination:

$$p(x) = 1 - \prod_{i \in N} x^i$$

and the expected payoff in the one-shot game $G(w)$:

$$G(w), x = \left(\prod_{i \in N} x^i \right) w + \sum_{\emptyset \subset S \subset N} \left(\prod_{i \notin S} x^i \right) \left(\prod_{i \in S} (1 - x^i) \right) r_S.$$

With an abuse of notations, we denote by $\text{supp}(x^i)$ the actions that are played with positive probability under x^i . Thus, $\text{supp}(x^i) = \{q^i\}$, $\{c^i\}$, or $\{c^i, q^i\}$ when $x^i = 0, 1$ or $x^i \in (0, 1)$ respectively.

DEFINITION 2.1 *An action a^i of player i is an ε -best reply for x^{-i} if*

$$G(w), x^{-i}, a^i \geq \max_{b^i \in \{c^i, q^i\}} G(w), x^{-i}, b^i - \varepsilon.$$

A profile x in $G(w)$ is a perfect ε -equilibrium if for every player i , every action $a^i \in \text{supp}(x^i)$ is an ε -best reply for x^{-i} .

2.2 The proof

Fix an $\varepsilon > 0$ sufficiently small once and for all. Let $W \subset \mathbf{R}^N$ be a compact set. Define the correspondence $\psi : W \rightarrow W$ by: $\psi(w)$ is the subset of all vectors $G(w), x$ such that x is a perfect $\rho\varepsilon$ -equilibrium profile in $G(w)$ that satisfies $G(w), x \in W$ and $p(x) \leq \varepsilon$. Clearly, ψ is upper-semi-continuous. However, $\psi(w)$ might be empty.

Theorem 1.2 follows from the next three Propositions.

PROPOSITION 2.2 *Define*

$$W \stackrel{\text{def}}{=} \{w \in [-\rho, \rho]^N \mid \exists i \in N \text{ with } w^i = 1\}.$$

Assume that (i) **A.1** holds, and (ii) for every $w \in W$ there exists an equilibrium x in $G(w)$ such that either a) $x = c$ (everyone continues), or b) $x = c$ and for some $i \in N$, $x^i < 1$ and $G(w), x^i = 1$. Then ψ has non-empty values.

In other words, the lemma claims that if for any continuation payoff $w \in W$ there is an equilibrium in $G(w)$, such that one of the players who quit with positive probability receives at most 1, then ψ has non empty values

It is clear that if the game satisfies **A.1** and **A.2** then the conditions of Proposition 2.2 hold. However, there are games that do not satisfy **A.2** but satisfy the condition of Proposition 2.2.

PROPOSITION 2.3 *If there exists a compact set W such that ψ has non-empty values, then there exist a cyclic profile $\mathbf{x} = (x_n)_n$ in G such that, for every n :*

1. \mathbf{x}_n is terminating;
2. x_n is a perfect $(\rho + 2)\varepsilon$ -equilibrium of $G(\gamma(\mathbf{x}_{n+1}))$.

PROPOSITION 2.4 *Let $\mathbf{x} = (x_n)_n$ be a profile in G . Assume that the following properties hold for every n :*

1. \mathbf{x}_n is terminating;
2. x_n is a perfect ε -equilibrium of $G(\gamma(\mathbf{x}_{n+1}))$.

Then either \mathbf{x} is a subgame perfect $\varepsilon^{\frac{1}{6}}$ -equilibrium, or there is a stationary $\varepsilon^{\frac{1}{6}}$ -equilibrium.

The exponent $\frac{1}{6}$ has been chosen for computational ease. It is not the best possible.

The following three subsections are devoted to the proofs of these three

Define the periodic sequence

$$(w_l)_l = (w_1, \widehat{\psi}^{L-1}(w_1), \widehat{\psi}^{L-2}(w_1), \dots, \widehat{\psi}(w_1), w_1, \widehat{\psi}^{L-1}(w_1), \dots).$$

Observe that $w_l = \widehat{\psi}(w_{l+1})$ for each $l \in \mathbf{N}$. Let $l < L$, and denote by Ω_{l+1} the element of \mathcal{W} that contains w_{l+1} . By definition of $\widehat{\psi}$, there exists $\widehat{w}_{l+1} \in \Omega_{l+1}$, such that $w_l = \psi(\widehat{w}_{l+1})$, hence there exists a perfect $\rho\varepsilon$ -equilibrium profile x_l in $G(\widehat{w}_{l+1})$ such that $w_l = G(\widehat{w}_{l+1}, x_l)$ and $p(x_l) \geq \varepsilon$. We extend (x_1, \dots, x_{L-1}) to a periodic sequence $\mathbf{x} = (x_l)_l$.

We now prove that \mathbf{x} satisfies the conclusions of the proposition. Since $p(x_l) \geq \varepsilon$ for each l , the first conclusion is obvious. Our next step is to prove that $\|\gamma(\mathbf{x}_l) - w_l\| \leq \varepsilon$ for every l . We then show that x_l is a perfect $(\rho + 2)\varepsilon$ -equilibrium in $G(\gamma(x_{l+1}))$, and the second conclusion follows.

Let $l^* \in \mathbf{N}$ maximize $\|\gamma(\mathbf{x}_l) - w_l\|$, and set $\alpha = \|\gamma(\mathbf{x}_{l^*}) - w_{l^*}\|$. Observe first that

$$\gamma(\mathbf{x}_{l^*}) = G(\gamma(\mathbf{x}_{l^*+1}), x_{l^*}).$$

Since $w_{l+1} - \widehat{w}_{l+1} \leq \varepsilon^2$, one has

$$\|w_{l^*} - G(w_{l^*+1}, x_{l^*})\| \leq \varepsilon^2.$$

Thus

$$\begin{aligned} \alpha &= \|\gamma(\mathbf{x}_{l^*}) - w_{l^*}\| \\ &\leq \|w_{l^*} - G(w_{l^*+1}, x_{l^*})\| + \|\gamma(\mathbf{x}_{l^*}) - G(w_{l^*+1}, x_{l^*})\| \\ &\leq \varepsilon^2 + \|G(\gamma(\mathbf{x}_{l^*+1}), x_{l^*}) - G(w_{l^*+1}, x_{l^*})\| \\ &\leq \varepsilon^2 + \|\gamma(\mathbf{x}_{l^*+1}) - w_{l^*+1}\| \times (1 - \varepsilon) \\ &\leq \varepsilon^2 + (1 - \varepsilon)\alpha, \end{aligned}$$

where the third inequality uses the fact that the probability under x_{l^*} that a player quits is at least ε . Thus $\alpha \leq \varepsilon$, as claimed.

Since x_l is a perfect $\rho\varepsilon$ -equilibrium in $G(\widehat{w}_{l+1})$ and

$$\widehat{w}_{l+1} - \gamma(\mathbf{x}_{l+1}) = \widehat{w}_{l+1} - w_{l+1} + w_{l+1} - \gamma(\mathbf{x}_{l+1}) = \varepsilon^2 + \varepsilon,$$

x_l is a perfect $(\rho + 2)\varepsilon$ -equilibrium in $G(\gamma(\mathbf{x}_{l+1}))$, which ends the proof. ■

We provide now a simpler derivation, which avoids the discretization of the above proof, but does not result in cyclic profiles. It is based on the following result, which has its own interest.

LEMMA 2.5 *Let $\phi : K \rightarrow K$ be an upper-semi-continuous correspondence with non-empty values defined over a compact space K . Then there exists a sequence $k_1, k_2, \dots \in K$ such that $k_i \in \phi(k_{i+1})$ for every $i \in \mathbf{N}$.*

Since the correspondence ϕ has non empty values, for every $k_1 \in K$ there is a sequence k_1, k_2, k_3, \dots such that $k_{i+1} \in \phi(k_i)$. This lemma claims that an inverse iterative sequence also exists. This lemma is used as a substitute for a fixed point theorem.

Proof: Define $K_0 = K$ and $K_i = \phi(K_{i-1}) = \bigcup_{k \in K_{i-1}} \phi(k)$ for every $i \in \mathbf{N}$. Since ϕ has non-empty values, K_i is non empty, and since ϕ is upper-semi-continuous and K compact, K_i is compact. Clearly $K_i \subset K_{i-1}$, hence $K_\infty = \bigcap_{i \in \mathbf{N}} K_i$ is non-empty.

Choose $k_1 \in K_\infty$. In particular, $k_1 \in K_i$ for every i , and therefore for every i there exists a sequence $k_1 = k_1^i, k_2^i, \dots, k_j^i$ such that $k_j^i \in \phi(k_{j+1}^i)$. By taking subsequences, we can assume that the limit $k_j^\infty = \lim_i k_j^i$ exists for every j . Since ϕ is upper-semi-continuous, $k_j^\infty \in \phi(k_{j+1}^\infty)$, as desired. ■

Proof of Proposition 2.3 that does not give cyclic equilibrium: Denote by (w_n) the sequence obtained by use of Lemma 2.5 w.r.t. ψ . For every $n \in \mathbf{N}$ choose $y_n \in [0, 1]^N$ such that y_n is a perfect $\rho\varepsilon$ -equilibrium profile in $G(w_{n+1})$ that satisfies

- $G(w_{n+1}, y_n) = w_n$.
- $p(y_n) \leq \varepsilon$.

Since $p(y_n) \leq \varepsilon$ for every n , it follows that for every $n \in \mathbf{N}$, y_n is terminating and that $\gamma(y_n) = w_n$. Indeed, define

$$Y_n = \begin{cases} r_{S_t} & n > t \\ w_n & \text{otherwise} \end{cases}$$

Then Y_n is a bounded martingale under y , which coincides with r_{S_t} for $n > t$. Since t is finite y -a.s., $\lim_n Y_n = r_{S_t}$. By dominated convergence, $w_0 = \mathbf{E}_y[r_{S_t}] = \gamma(y)$. A similar argument shows that $w_n = \gamma(y_n)$.

By Proposition 2.4, a subgame perfect $\varepsilon^{\frac{1}{\delta}}$ -equilibrium \mathbf{x} exists. However, the profile \mathbf{x} needs not be cyclic. Moreover, it needs not be true that a cyclic one can be obtained by repeating a finite segment of \mathbf{x} , as the following example shows.

Example 2: consider the two-player quitting game

	2, 1
1, 1	1, 1

and define a profile \mathbf{x} by

$$\begin{aligned} x_n^1 &= 1 - \eta \text{ for every } n \\ x_n^2 &= 1 - \frac{\varepsilon}{2^n} \end{aligned}$$

where $\eta > \varepsilon$. Since the probability under \mathbf{x}_n^2 that player 2 will ever quit is at most ε , player 1's payoff cannot exceed $1 + \varepsilon$. Hence \mathbf{x} is a subgame perfect ε -equilibrium. Let now $\bar{\mathbf{x}}$ be a strategy profile that is obtained by taking a finite set of stages and repeating over time the restriction of \mathbf{x} to this set. Since \mathbf{x}^1 is stationary, $\bar{\mathbf{x}}^1 = \mathbf{x}^1$. Since $\eta > \varepsilon$, $\gamma^1(\bar{\mathbf{x}}) < \frac{3}{2}$. Under $\bar{\mathbf{x}}^2$, player two will eventually quit, therefore $\gamma^1(\mathbf{c}^1, \bar{\mathbf{x}}^2) = 2$. Thus, $\bar{\mathbf{x}}$ is not a $\frac{1}{2}$ -equilibrium.

Nevertheless, it can be shown that a slightly more sophisticated procedure than the one which repeats a given segment of \mathbf{x} may be used to modify \mathbf{x} into a cyclic subgame perfect ε' -equilibrium (where ε' goes to zero with ε). We skip the details. ■

2.5 Proof of the main proposition

This section contains the proof of Proposition 2.4, which we state again for convenience.

PROPOSITION 2.6 *Let $\mathbf{x} = (x_n)_n$ be a profile in G . Assume that the following properties hold for every n :*

1. \mathbf{x}_n is terminating;
2. x_n is a perfect ε -equilibrium of $G(\gamma(\mathbf{x}_{n+1}))$.

Then either \mathbf{x} is a subgame perfect $\varepsilon^{\frac{1}{6}}$ -equilibrium, or there is a stationary $\varepsilon^{\frac{1}{6}}$ -equilibrium.

In general, \mathbf{x} needs not be an $\varepsilon^{\frac{1}{6}}$ -equilibrium. Indeed, consider the following game, where only the payoffs of player 1 appear:

	$1 - \epsilon^2$	ϵ^2
$1 - \epsilon$	<i>continue</i>	5
ϵ	1	1

Figure 1

The stationary profile depicted in Figure 1 yields to player 1 a payoff $\frac{1+5\epsilon(1-\epsilon)}{1+\epsilon(1-\epsilon)}$, and hence player 1 receives, up to 6ϵ , the payoff 1 by either quitting or continuing in the one-shot game $G\left(\frac{1+5\epsilon(1-\epsilon)}{1+\epsilon(1-\epsilon)}\right)$. However, if player 1 always continues, then his payoff is 5.

Observe that in this case, player 1 quits with probability ϵ , and player 2 quits with probability ϵ^2 . Therefore, if player 2 did continue with probability 1, the expected payoff for the players would not change by much, while the altered profile would be an equilibrium. This idea is fundamental to the proof.

The proof requires careful accounting of the probability of termination and the amount a player can gain by deviating. We start by explaining the basic difficulty and the basic insights of the proof. We continue with a presentation of the proof.

2.5.1 Overview

The assumption says that, by changing x_n^i , player i may increase his expected continuation payoff by at most ϵ above $\gamma^i(\mathbf{x}_n)$. By changing x_n^i repeatedly to c^i for N_0 iterations, the expected payoff increases by at most ϵN_0 .

By playing c^i instead of x_n^i at stage n , player i increases his expected continuation payoff by at most ε times the probability of playing q^i at stage n , since whenever this probability is positive, q^i is an ε -best reply in the one-shot game. Thus, the increase is at most $\varepsilon(1 - x_n^i)$. These profits may accumulate only if, when playing repeatedly c^i instead of x_n^i , the rate of termination is not too large, compared to $\varepsilon(1 - x_n^i)$. Roughly speaking, this implies that, over a large number of stages, the probability that players $N \setminus \{i\}$ will quit is small compared to the probability that player i quits under \mathbf{x}^i . It follows that with high probability the terminating coalition is $\{i\}$, hence the payoff $\gamma(\mathbf{x})$ is close to $r_{\{i\}}$. In particular, replacing \mathbf{x}^{-i} with \mathbf{c}^{-i} leaves the payoff vector essentially unchanged. It is not difficult to conclude that, for some $\alpha > 0$, $((1 - \alpha)\mathbf{c}^i + \alpha\mathbf{q}^i, \mathbf{c}^{-i})$ is a stationary β -equilibrium, where β goes to zero with ε .

The technical difficulty is of course that the maximal profit $\varepsilon(1 - x_n^i)$ may vary with n , and comparing the undefined rate of termination with it makes no sense. We use the following construction which is made precise later. We divide the stages into consecutive (disjoint) blocks of stages, on each of which the probability of termination under \mathbf{x} is of the order ε^a , where a is to be specified. Thus, ε^a is the rate of termination when time is measured in blocks. We divide these blocks into two categories, depending on whether termination would essentially be due to player i or not: the first category contains the blocks on which the probability of termination remains at least ε^b if player i switches to \mathbf{c}^i (where $b > a$), and the second category contains all other blocks. Thus, for any block of the second type, the payoff under \mathbf{x} , conditional on the event that termination occurs there, is close to $r_{\{i\}}$, provided ε is sufficiently small.

Assume that, at some point in time, there are more than $\frac{1}{\varepsilon^e}$ consecutive

blocks of the second type, with $e > a$. At the beginning of the first block in the sequence, the expected continuation payoff is close to $r_{\{i\}}$, and in these blocks, quitting is done mainly by player i . As described above, in this case there exists a stationary ε -equilibrium with corresponding payoff $r_{\{i\}}$, in which only player i quits with positive probability.

Assume now that no two consecutive blocks of the first type are more than $\frac{1}{\varepsilon^e}$ blocks apart. We prove that no pure strategy of i improves upon \mathbf{x}^i by more than some β . This is done by showing that if player i is allowed to deviate only in a single block, he cannot profit too much compared to the probability that the other players quit during this block. This bound is the core of the whole proof.

2.5.2 The two types of blocks

Let \mathbf{y} be a strategy profile. For any two integers $n \leq m$ we set

$$p_{\mathbf{y}}[n, m] = \mathbf{P}_{\mathbf{y}}(t \leq m/t \leq n) \quad (= \mathbf{P}_{\mathbf{y}_n}(t \leq m - n)).$$

This is the probability that the game terminates between stages n and m (conditional on $t \leq n$).

We write simply $p_{\mathbf{y}}[n]$ for $p_{\mathbf{y}}[n, n]$, and $p_{\mathbf{y}}[B]$ for $p_{\mathbf{y}}[n, m]$, when $B = \{n, \dots, m\}$.

In the proof, we use repeatedly the following facts:

Fact 1: if u is a random variable on a probability space $(\Omega, \mathbf{A}, \mathbf{P})$, bounded by ρ , and $A \in \mathbf{A}$,

$$|\mathbf{E}(u)| \leq \rho \mathbf{P}(A^c) + \sup_A |u|.$$

Fact 2: for every profile \mathbf{y} , and every $n \in \mathbf{N}$,

$$\gamma(\mathbf{y}) = \mathbf{P}_{\mathbf{y}}(t < n) \mathbf{E}_{\mathbf{y}}[r_{S_t} | t < n] + \mathbf{P}_{\mathbf{y}}(t \geq n) \gamma(\mathbf{y}_n).$$

Fact 3: for each i and x , the vector $G(w), (x^{-i}, q^i)$ is independent of w and is equal to $\gamma(\mathbf{x}^{-i}, \mathbf{q}^i)$.

We fix $a, b, e \in]0, 1[$ so that: $e > a > \frac{1}{5}$, $b - e > \frac{1}{5}$, $1 - b - e > \frac{1}{5}$ (choose e between $\frac{1}{5}$ and $\frac{3}{10}$, then a and b). These are the powers of ε that were mentioned in the overview. Assume for convenience that $\frac{1}{\varepsilon^e}$ is an integer. We will also assume that $\varepsilon > 0$ is sufficiently small, so that (a finite number of) inequalities like $N_0 \varepsilon^2 < \varepsilon$, $(1 - \varepsilon^a)^{\frac{1}{\varepsilon^e}} < \varepsilon$ will hold. Observe that $(1 - \varepsilon^a)^{\frac{1}{\varepsilon^e}} = \exp(\frac{1}{\varepsilon^e} \ln(1 - \varepsilon^a)) < \exp(-\varepsilon^{a-e})$ for $\varepsilon < 1$. Since $e > a$, $\frac{1}{\varepsilon} \exp(-\varepsilon^{a-e})$ goes to zero when ε goes to zero. Hence the second inequality for small ε .

We now introduce the partition of \mathbf{N} into a sequence $(B_k)_{k \geq 1}$ of blocks. Set $n_1 = 1$ and

$$n_{k+1} = \inf \{n > n_k, p_{\mathbf{x}}[n_k, n - 1] < \varepsilon^a\}.$$

Set $B_k = \{n_k, \dots, n_{k+1} - 1\}$. Since \mathbf{x}_n is terminating for each n , every B_k is a finite set.

Observe that n_{k+1} is characterized by the inequalities

$$p_{\mathbf{x}}[n_k, n_{k+1} - 2] < \varepsilon^a \quad p_{\mathbf{x}}[n_k, n_{k+1} - 1].$$

In particular,

- in every stage n of B_k except possibly the last one (stage $n_{k+1} - 1$), $p_{\mathbf{x}}[n] < \varepsilon^a$, hence x_n is close to c (recall that c is the action combination where everyone continues).
- no such estimate is available for $x_{n_{k+1}-1}$, and $p_{\mathbf{x}}[n_{k+1} - 1]$ may be “large”.

DEFINITION 2.7 *Let $i \in N$. A block B_k is of type I for player i if $p_{\mathbf{x}^{-i}, \mathbf{c}^i}[B_k] \geq \varepsilon^b$, and of type II otherwise.*

In words, a block B_k is of type I for i if there is a non-negligible probability that some other player will quit in B_k .

We prove Proposition 2.6 by discussing the following two cases:

Case 1: If, for some player i , there exist at least $1/\varepsilon^e$ consecutive blocks that are of type II for i then for some $\alpha > 0$, $(\alpha \mathbf{q}^i + (1 - \alpha) \mathbf{c}^i, \mathbf{c}^{-i})$ is a stationary $\varepsilon^{\frac{1}{6}}$ -equilibrium.

Case 2: Otherwise, \mathbf{x} is a subgame perfect $\varepsilon^{\frac{1}{6}}$ -equilibrium.

These two cases are dealt with in the following two subsections.

2.5.3 A stationary equilibrium when blocks of type I are sparse

We place ourselves here under the assumptions of **Case 1**. The assumptions of Proposition 2.6 are still satisfied when the first stage is shifted. Thus, we may assume that the first block of type I has an index $l > 1/\varepsilon^e$.

We prove first that both $\gamma(\mathbf{x})$ and $\gamma(\mathbf{x}_2)$ are close to $r_{\{i\}}$. To prove that we have two intermediate steps: we first show that if the game terminates before stage n_l , then with high probability the quitting coalition is $\{i\}$. We then show that with high probability the game terminates before stage n_l .

Let $k < l$ be given. Observe that $p_{\mathbf{x}^{-i}, \mathbf{c}^i}[B_k]$ (resp. $p_{\mathbf{x}}[B_k]$) is the probability that at least some player $j = i$ (resp. at least some player j) plays q^j during B_k . Since the block B_k is of type II,

$$\mathbf{P}_{\mathbf{x}}(a_n^j = q^j \text{ for some } j = i, \text{ and some } n_k \leq n < n_{k+1}) < \varepsilon^b. \quad (2)$$

Since $p_{\mathbf{x}}[B_k] \geq \varepsilon^a$, one deduces

$$\mathbf{P}_{\mathbf{x}}\{j = i, j \leq S_t/t \mid B_k\} < \varepsilon^{b-a},$$

that is, $\mathbf{P}_{\mathbf{x}}\{S_t = \{i\}/t \mid B_k\} \leq 1 - \varepsilon^{b-a}$. By summing over k , one obtains

$$\mathbf{P}_{\mathbf{x}}\{S_t = \{i\}/t < n_l\} \leq 1 - \varepsilon^{b-a}.$$

In words, if the game terminates before stage n_l , then with high probability the quitting coalition is $\{i\}$.

On the other hand, the probability $p_{\mathbf{x}}[B_k]$ of termination in block B_k is at least ε^a , for each $k < l$. Since $l \geq \frac{1}{\varepsilon^e}$,

$$\mathbf{P}_{\mathbf{x}}(t \leq n_l) \geq (1 - \varepsilon^a)^{\frac{1}{\varepsilon^e}} \geq \varepsilon.$$

In words, with high probability the game terminates before stage n_l .

Thus, $\mathbf{P}_{\mathbf{x}}(S_t = \{i\}, t < n_l) \geq 1 - \varepsilon - \varepsilon^{b-a}$. By **Fact 1**, one gets

$$\gamma(\mathbf{x}) - r_{\{i\}} \leq \rho(\varepsilon + \varepsilon^{b-a}). \quad (3)$$

Observe that this proof is also valid in the case $l = \frac{1}{\varepsilon^e}$. Hence, written for the blocks 2 to l , it yields $\gamma(\mathbf{x}_{n_2}) - r_{\{i\}} \leq \rho(\varepsilon + \varepsilon^{b-a})$.

We now deal with $\gamma(\mathbf{x}_2)$. If $p_{\mathbf{x}}[1] < \varepsilon^a$, one has $\gamma(\mathbf{x}) - \gamma(\mathbf{x}_2) \leq \rho\varepsilon^a$ (**Fact 2**), hence

$$\gamma(\mathbf{x}_2) - r_{\{i\}} \leq \rho(\varepsilon + \varepsilon^{b-a} + \varepsilon^a). \quad (4)$$

If $p_{\mathbf{x}}[1] \geq \varepsilon^a$, the block B_1 is reduced to stage 1, and $n_2 = 2$, hence (4) also holds in that case. Thus, we showed that both $\gamma(\mathbf{x})$ and $\gamma(\mathbf{x}_2)$ are close to $r_{\{i\}}$. We now construct a stationary $\varepsilon^{\frac{1}{\delta}}$ -equilibrium in the game.

Denote by $\bar{\mathbf{x}}^i$ the stationary strategy of player i which quits with probability $\bar{x}^i = \max(\varepsilon, x_1^i)$ in every stage. Since $p_{\mathbf{x}}[1] = p_{\mathbf{x}}[B_1] < \varepsilon^b$, one has

$x_1 - (c^{-i}, \bar{x}^i) < \varepsilon^b$. Since x_1 is an ε -equilibrium in $G(\gamma(\mathbf{x}_2))$, (c^{-i}, \bar{x}^i) is an $(\varepsilon + \rho N \varepsilon^b)$ -equilibrium in $G(\gamma(\mathbf{x}_2))$. Observe that $\gamma(\mathbf{c}^{-i}, \bar{\mathbf{x}}^i) = r_{\{i\}}$. Thus, by (4), (c^{-i}, \bar{x}^i) is an η -equilibrium in $G(\gamma(\mathbf{c}^{-i}, \bar{\mathbf{x}}^i))$ with $\eta = \varepsilon + \rho(\varepsilon + N \varepsilon^b + \varepsilon^a + \varepsilon^{b-a}) < \varepsilon^{\frac{1}{6}}$.

We now check that $(\mathbf{c}^{-i}, \bar{\mathbf{x}}^i)$ is an η -equilibrium of the quitting game.

Any pure deviation by player i from $\bar{\mathbf{x}}^i$ yields a payoff which is either 0 (deviation to \mathbf{c}^i) or $\gamma^i(\mathbf{c}^{-i}, \bar{\mathbf{x}}^i) = 1$ (any other deviation). Since $r_{\{i\}}^i = 1$, player i cannot profit by deviating. Since (c^{-i}, \bar{x}^i) is stationary, any pure deviation by player $j = i$ from \mathbf{c}^j yields a payoff equal to $\gamma^j(\mathbf{c}^{-i,j}, \bar{\mathbf{x}}^i, \mathbf{q}^j)$, which, by **Fact 3**, is also equal to

$$G(\gamma(\mathbf{c}^{-i}, \bar{\mathbf{x}}^i), (c^{-i,j}, \bar{x}^i, q^j)^j) = \gamma^j(\mathbf{c}^{-i}, \bar{\mathbf{x}}^i) + \eta.$$

Thus, $(\mathbf{c}^{-i}, \bar{\mathbf{x}}^i)$ is a stationary η -equilibrium of the quitting game G .

2.5.4 The case where blocks of type I are regularly scattered

We place ourselves in **Case 2**: for every player i , no two consecutive blocks of type I for i are too distant. We prove that \mathbf{x} is an $\varepsilon^{\frac{1}{6}}$ -equilibrium. Let $i \in \mathbf{N}$ be fixed.

First step: estimates for local deviations For each $k \in \mathbf{N}$, we define an auxiliary game Γ_k , played during the block B_k : it starts in stage n_k , and ends after stage $n_{k+1} - 1$, with payoffs given by r_{S_t} if $t < n_{k+1}$, and by $\gamma(\mathbf{x}_{n_{k+1}})$ if termination did not occur. We prove that player i cannot profit too much by deviating from $(x_{n_k}^i, x_{n_{k+1}}^i, \dots, x_{n_{k+1}-1}^i)$ in the game Γ_k , against the profile $(x_{n_k}^{-i}, x_{n_{k+1}}^{-i}, \dots, x_{n_{k+1}-1}^{-i})$. In the original quitting game, this will imply that player i cannot profit too much by deviating only during B_k .

For notational simplicity, we deal with the game Γ_1 . The corresponding estimates for Γ_k will be obtained by conditioning on the event $\{t = n_k\}$.

In the game Γ_1 , the payoff induced by a strategy profile \mathbf{y} is

$$g(\mathbf{y}) = \mathbf{E}_{\mathbf{y}}[r_{S_t} 1_{t < n_2} + \gamma(\mathbf{x}_{n_2}) 1_{t \geq n_2}].$$

Observe that $g(\mathbf{x}) = \gamma(\mathbf{x})$.

We shall compute a bound on $g^i(\mathbf{x}^{-i}, \mathbf{s}^i)$, where \mathbf{s}^i is any pure strategy of player i . We analyze in turn the case where $\mathbf{s}^i = \mathbf{q}_n^i$ (the strategy of i where he continues at every stage except stage n , in which he quits with probability 1), for some $n < n_2$, and $\mathbf{s}^i = \mathbf{c}^i$.

The goal of this first step is to prove that

$$g^i(\mathbf{x}^{-i}, \mathbf{q}_n^i) = g^i(\mathbf{x}) + 2\rho\varepsilon^a + \varepsilon, \text{ for each } n < n_2, \text{ and} \quad (5)$$

$$g^i(\mathbf{x}^{-i}, \mathbf{c}^i) = g^i(\mathbf{x}) + 7\rho\varepsilon^a N \mathbf{P}_{\mathbf{x}^{-i}, \mathbf{c}^i}(t < n_2 - 1) + 2\varepsilon. \quad (6)$$

Note that the bound on the amount player i can profit in (6) is much smaller than the corresponding amount in (5). In particular, the proof of (6) is more involved. We need a better estimate in (6) since once a player quits the game terminates, while as long as he continues, the error in our estimate increases.

Quitting in the auxiliary game. Fix $n < n_2$. We rewrite $g^i(\mathbf{x})$ and $g^i(\mathbf{x}^{-i}, \mathbf{q}_n^i)$ as

$$g^i(\mathbf{x}) = \mathbf{P}_{\mathbf{x}}(t < n) \times \mathbf{E}_{\mathbf{x}}[r_{S_t}^i / t < n] + \mathbf{P}_{\mathbf{x}}(t = n) \times \gamma^i(\mathbf{x}_n), \text{ and} \quad (7)$$

$$g^i(\mathbf{x}^{-i}, \mathbf{q}_n^i) = \mathbf{P}_{\mathbf{x}^{-i}, \mathbf{c}^i}(t < n) \times \mathbf{E}_{\mathbf{x}^{-i}, \mathbf{c}^i}[r_{S_t}^i / t < n] + \mathbf{P}_{\mathbf{x}^{-i}, \mathbf{c}^i}(t = n) \times G(\gamma(\mathbf{x}_{n+1})), (x_n^{-i}, q^i)^i. \quad (8)$$

The last equality is derived as follows: if $t = n$ (this occurs with probability $\mathbf{P}_{\mathbf{x}^{-i}, \mathbf{c}^i}(t = n)$), player i quits in stage n , hence gets $\gamma^i(\mathbf{x}_n^{-i}, \mathbf{q}^i)$, which yields (8), using **Fact 3**.

Since x_n is an ε -equilibrium in the game $G(\gamma(\mathbf{x}_{n+1}))$, one has

$$\begin{aligned} G(\gamma(\mathbf{x}_{n+1})), (x_n^{-i}, q^i) & \geq G(\gamma(\mathbf{x}_{n+1})), x_n^{-i} + \varepsilon \\ & = \gamma^i(\mathbf{x}_n) + \varepsilon. \end{aligned} \quad (9)$$

By definition of n_2 and since $n < n_2$,

$$\mathbf{P}_{\mathbf{x}^{-i}, \mathbf{c}^i}(t < n) - \mathbf{P}_{\mathbf{x}}(t < n) < \varepsilon^a. \quad (10)$$

From (7), (8), (9) and (10), one gets

$$g^i(\mathbf{x}^{-i}, \mathbf{q}_n^i) \geq g^i(\mathbf{x}) + 2\rho\varepsilon^a + \varepsilon,$$

which is (5).

Continuing in the auxiliary game. As in the previous paragraph, we compare $g^i(\mathbf{x}^{-i}, \mathbf{c}^i)$ and $g^i(\mathbf{x})$. For notational simplicity, set

$$\begin{aligned} \pi_1 &= \mathbf{P}_{\mathbf{x}}(t < n_2 - 1, i \in S_t), \\ u_1 &= \mathbf{E}_{\mathbf{x}}[r_{S_t}^i / t < n_2 - 1, i \in S_t], \\ \pi_2 &= \mathbf{P}_{\mathbf{x}}(t < n_2 - 1, i \notin S_t), \text{ and} \\ u_2 &= \mathbf{E}_{\mathbf{x}}[r_{S_t}^i / t < n_2 - 1, i \notin S_t], \end{aligned}$$

so that

$$g^i(\mathbf{x}) = \pi_1 u_1 + \pi_2 u_2 + (1 - \pi_1 - \pi_2) \gamma^i(\mathbf{x}_{n_2-1}). \quad (11)$$

Set also

$$\begin{aligned} \pi_2^* &= \mathbf{P}_{\mathbf{x}^{-i}, \mathbf{c}^i}(t < n_2 - 1) \text{ and} \\ u_2^* &= \mathbf{E}_{\mathbf{x}^{-i}, \mathbf{c}^i}[r_{S_t}^i / t < n_2 - 1], \end{aligned}$$

so that

$$g^i(\mathbf{x}^{-i}, \mathbf{c}^i) = \pi_2^* u_2^* + (1 - \pi_2^*) G(\gamma(\mathbf{x}_{n_2})), (x_{n_2-1}^{-i}, c^i)^i. \quad (12)$$

Since x_{n_2-1} is an ε -equilibrium in the game $G(\gamma(\mathbf{x}_{n_2}))$, one gets (as in (9))

$$G(\gamma(\mathbf{x}_{n_2})), (x_{n_2-1}^{-i}, c^i)^i \geq \gamma^i(\mathbf{x}_{n_2-1}) + \varepsilon.$$

Plugging this inequality into (11) and (12), one obtains

$$g^i(\mathbf{x}^{-i}, \mathbf{c}^i) \geq g^i(\mathbf{x}) + \pi_2^* u_2^* - \pi_2 u_2 + (\pi_2 - \pi_2^*) \gamma^i(\mathbf{x}_{n_2-1}) + \pi_1 (\gamma^i(\mathbf{x}_{n_2-1}) - u_1) + \varepsilon. \quad (13)$$

In the next three lemmas, we provide estimates on $|\pi_2 - \pi_2^*|$, on $|u_2 - u_2^*|$ and on $\gamma^i(\mathbf{x}_{n_2-1}) - u_1$ that immediately yield, using (13),

$$g^i(\mathbf{x}^{-i}, \mathbf{c}^i) \geq g^i(\mathbf{x}) + 7\rho\varepsilon^a N \pi_2^* + 2\varepsilon,$$

which is (6). This will end the first step.

LEMMA 2.8 *One has $|\pi_2 - \pi_2^*| \leq \varepsilon^a \pi_2^*$.*

Proof: Observe first that by the definition of n_2 ,

$$\pi_1, \pi_2, \pi_2^* \leq \mathbf{P}_{\mathbf{x}}\{t < n_2 - 1\} < \varepsilon^a. \quad (14)$$

For $n < n_2 - 1$, set $X_n = 1$ if at least some player $j = i$ quits in stage n , $X_n = 0$ otherwise, and set $S_1 = \inf\{n \leq n_2 - 2, X_n = 1\}$ (by convention $\inf \emptyset = +\infty$). Set $Y_n = 1$ if player i quits in stage n , $Y_n = 0$ otherwise, and set $S_2 = \inf\{n \leq n_2 - 2, Y_n = 1\}$. Finally, set $T = S_1$ if $S_1 < S_2$ and $T = S_2$ otherwise.

Thus,

$$\pi_2^* = \mathbf{P}_{\mathbf{x}}(S_1 \leq n_2 - 2) \text{ and } \pi_2 = \mathbf{P}_{\mathbf{x}}(T \leq n_2 - 2).$$

In particular, $\pi_2 = \pi_2^*$. By Lemma 3.1,

$$|\pi_2 - \pi_2^*| = \pi_2^* \times \mathbf{P}_{\mathbf{x}}(S_2 = n_2 - 2).$$

Clearly, $\mathbf{P}_{\mathbf{x}}(S_2 = n_2 - 2) = \mathbf{P}_{\mathbf{x}}(t < n_2 - 1) < \varepsilon^a$. The conclusion of Lemma 2.8 follows. \blacksquare

LEMMA 2.9 *One has $|u_2 - u_2^*| \leq 2\rho\varepsilon^a$.*

Proof: Denote by Π_2 the distribution of t under $\mathbf{P}_{\mathbf{x}}$, conditional on the event $\{t = n_2 - 2, i / S_t\}$. Denote by Π_2^* the distribution of t under $\mathbf{P}_{\mathbf{x}^{-i}, \mathbf{e}^i}$, conditional on the event $\{t = n_2 - 2\}$. Since u_2 (resp. u_2^*) is the expectation of $r_{S_t}^i$ under Π_2 (resp. under Π_2^*), one has

$$|u_2 - u_2^*| \leq \rho \|\Pi_2 - \Pi_2^*\|_1 \quad (15)$$

where $\|\Pi_2 - \Pi_2^*\|_1 = \sum_n |\Pi_2(n) - \Pi_2^*(n)|$ is the L^1 -distance between Π_2 and Π_2^* .

With the notations of the previous lemma,

$$\begin{aligned} \Pi_2(n) &= \mathbf{P}_{\mathbf{x}}(S_1 = n / T \mid n_2 - 2) \\ \Pi_2^*(n) &= \mathbf{P}_{\mathbf{x}}(S_1 = n / S_1 \mid n_2 - 2). \end{aligned}$$

By Lemma 3.1, one has

$$\|\Pi_2 - \Pi_2^*\|_1 = 2\mathbf{P}_{\mathbf{x}}(S_2 = n_2 - 2) < 2\varepsilon^a.$$

The conclusion of Lemma 2.9 then follows from (15). \blacksquare

LEMMA 2.10 *One has*

$$|\gamma^i(\mathbf{x}_{n_2-1}) - u_1| \leq 4\rho N \pi_2^* + \varepsilon \quad (16)$$

where N is the number of players.

Proof: Recall that u_1 is the payoff received in the termination stage, conditional on the event that i belongs to the quitting coalition. We proceed in three steps.

Step 1: We first estimate how close u_1 is to $r_{\{i\}}^i = 1$.

Let $n < n_2 - 1$. Conditional on $\{t = n, a_n^i = q^i\}$, the vector of actions a_n^{-i} is distributed according to x_n^{-i} . Since the overall probability that some player $j = i$ quits before stage $n_2 - 1$ is π_2^* , one has

$$x_n^{-i} - c^{-i} \leq \pi_2^*. \quad (17)$$

Hence

$$|\mathbf{E}_{\mathbf{x}}[r_{S_t}^i | a_t^i = q^i, t = n] - 1| \leq \rho N \pi_2^*.$$

By summation over n , this yields

$$|u_1 - 1| \leq \rho N \pi_2^*. \quad (18)$$

Step 2: We now compare $\gamma^i(\mathbf{x}_{\bar{n}+1})$ to $r_{\{i\}}^i = 1$, where \bar{n} is the last stage prior to $n_2 - 1$, for which $x_{\bar{n}}^i < 1$.

Since $x_{\bar{n}}$ is a perfect ε -equilibrium in the one-shot game $G(\gamma(\mathbf{x}_{\bar{n}+1}))$, q^i is a best reply to $x_{\bar{n}}^{-i}$, up to ε :

$$\gamma^i(\mathbf{x}_{\bar{n}}) \geq G(\gamma(\mathbf{x}_{\bar{n}+1}), (x_{\bar{n}}^{-i}, q^i))^i + \varepsilon. \quad (19)$$

By (17),

$$\begin{cases} \gamma^i(\mathbf{x}_{\bar{n}+1}) \geq \gamma^i(\mathbf{x}_{\bar{n}}) + \pi_2^* N \rho, \\ G(\gamma(\mathbf{x}_{\bar{n}+1}), (x_{\bar{n}}^{-i}, q^i))^i \geq G(\gamma(\mathbf{x}_{\bar{n}+1}), (c^{-i}, q^i))^i + \pi_2^* N \rho = 1 + \pi_2^* N \rho \end{cases}$$

which yields, using (19)

$$\gamma^i(\mathbf{x}_{\bar{n}+1}) \geq 1 + \varepsilon + 2\rho N \pi_2^*. \quad (20)$$

Step 3: We finally compare $\gamma^i(\mathbf{x}_{\bar{n}+1})$ to $\gamma^i(\mathbf{x}_{n_2-1})$.

By definition of \bar{n} , player i continues in every stage between $\bar{n} + 1$ and $n_2 - 2$ under \mathbf{x}^i . Thus,

$$p_{\mathbf{x}}[\bar{n} + 1, n_2 - 2] = p_{\mathbf{x}^{-i}, \mathbf{c}^i}[\bar{n} + 1, n_2 - 2] \quad \pi_2^*.$$

By **Fact 2** (applied with $\mathbf{y} = \mathbf{x}_{\bar{n}+1}$), this yields

$$|\gamma^i(\mathbf{x}_{\bar{n}+1}) - \gamma^i(\mathbf{x}_{n_2-1})| \leq \rho\pi_2^*. \quad (21)$$

The conclusion of Lemma 2.10 follows from (18), (20) and (21). \blacksquare

Second step: Global estimates and conclusion. In this paragraph, we compare $\gamma^i(\mathbf{x}^{-i}, \mathbf{s}^i)$ to $\gamma^i(\mathbf{x})$ where \mathbf{s}^i is any pure strategy of player i : either \mathbf{c}^i or \mathbf{q}_n^i . In the first step, we proved (see (5) and (6)), that

$$\begin{aligned} \mathbf{E}_{\mathbf{x}^{-i}, \mathbf{q}_n^i} [r_{S_t}^i] &\leq \gamma^i(\mathbf{x}) + 2\rho\varepsilon^a + \varepsilon, \text{ for every } n < n_2, \text{ and} \\ \mathbf{E}_{\mathbf{x}^{-i}, \mathbf{s}^i} [r_{S_t}^i 1_{t < n_2} + \gamma^i(\mathbf{x}_{n_2}) 1_{t \geq n_2}] &\leq \gamma^i(\mathbf{x}) + 7\rho\varepsilon^a N\pi_2^* + 2\varepsilon, \text{ for any other pure strategy } \mathbf{s}^i \end{aligned}$$

These estimates, which were obtained for the game played on the first block, have analogs for the game Γ_k played on B_k , obtained by conditioning upon $\{t \leq n_k\}$. These analogs are

$$\mathbf{E}_{\mathbf{x}^{-i}, \mathbf{q}_n^i} [r_{S_t}^i / t \leq n_k] \leq \gamma^i(\mathbf{x}_{n_k}) + 2\rho\varepsilon^a + \varepsilon, \text{ for every } n_k < n_{k+1}, \quad (22)$$

and

$$\mathbf{E}_{\mathbf{x}^{-i}, \mathbf{s}^i} [r_{S_t}^i 1_{t < n_{k+1}} + \gamma^i(\mathbf{x}_{n_{k+1}}) 1_{t \geq n_{k+1}} / t \leq n_k] \leq \gamma^i(\mathbf{x}_{n_k}) + 7\rho\varepsilon^a N\mathbf{P}_{\mathbf{x}^{-i}, \mathbf{c}^i}(t < n_{k+1} / t \leq n_k) + 2\varepsilon, \quad (23)$$

for any \mathbf{s}^i that coincides with \mathbf{c}^i up to n_{k+1} .

It is convenient to introduce the sequence $(X_k)_{k \in \mathbb{N}}$ of random variables defined by

$$X_k = \begin{cases} \gamma^i(\mathbf{x}_{n_k}) & \text{if } t = n_k \\ r_{S_t}^i & \text{if } t < n_k \end{cases}$$

Notice that $\mathbf{E}_{\mathbf{x}^{-i}, \mathbf{c}^i}[X_k]$ is the payoff to player i if player i follows \mathbf{c}^i up to stage n_k , then \mathbf{x}^i , while players $-i$ use \mathbf{x}^{-i} . We first compare $\gamma^i(\mathbf{x}^{-i}, \mathbf{s}^i)$ and $\mathbf{E}_{\mathbf{x}^{-i}, \mathbf{c}^i}[X_k]$.

LEMMA 2.11 *One has $\gamma^i(\mathbf{x}^{-i}, \mathbf{s}^i) \geq \sup_k \mathbf{E}_{\mathbf{x}^{-i}, \mathbf{c}^i}[X_k] + \varepsilon + 2\rho\varepsilon^a$.*

Proof: Let $\mathbf{s}^i = \mathbf{c}^i$. Since there are infinitely many blocks of type I for player i , the profile $(\mathbf{x}^{-i}, \mathbf{c}^i)$ is terminating. Thus, $\gamma^i(\mathbf{x}^{-i}, \mathbf{c}^i) = \lim_k \mathbf{E}_{\mathbf{x}^{-i}, \mathbf{c}^i}[X_k]$, hence the conclusion of the lemma holds in that case.

Let now $\mathbf{s}^i = \mathbf{q}_n^i$, and let B_k be the block containing n . The payoff $\gamma^i(\mathbf{x}^{-i}, \mathbf{q}_n^i)$ can be written

$$\gamma^i(\mathbf{x}^{-i}, \mathbf{q}_n^i) = \mathbf{E}_{\mathbf{x}^{-i}, \mathbf{c}^i}[X_k 1_{t < n_k}] + \mathbf{E}_{\mathbf{x}^{-i}, \mathbf{q}_n^i}[r_{S_t}^i 1_{t \geq n_k}]. \quad (24)$$

By definition, $\gamma^i(\mathbf{x}_{n_k}) = \mathbf{E}_{\mathbf{x}^{-i}, \mathbf{q}_n^i}[X_k / t = n_k]$. Using (22), the equality (24) implies

$$\gamma^i(\mathbf{x}^{-i}, \mathbf{q}_n^i) \geq \mathbf{E}_{\mathbf{x}^{-i}, \mathbf{c}^i}[X_k] + \varepsilon + 2\rho\varepsilon^a.$$

■

LEMMA 2.12 $\sup_k \mathbf{E}_{\mathbf{x}^{-i}, \mathbf{c}^i}[X_k] \geq \gamma^i(\mathbf{x}) + 2\varepsilon^{1-b-e} + 7\rho N\varepsilon^a$.

Proof: Using the definition of (X_k) , rewrite (23) as

$$\mathbf{E}_{\mathbf{x}^{-i}, \mathbf{c}^i}[X_{k+1} / t = n_{k+1}] \geq \mathbf{E}_{\mathbf{x}^{-i}, \mathbf{c}^i}[X_k / t = n_k] + 7\rho N\varepsilon^a \mathbf{P}_{\mathbf{x}^{-i}, \mathbf{c}^i}(t < n_{k+1} / t = n_k) + 2\varepsilon.$$

Note that $X_{k+1} = X_k$ if $t < n_k$. Therefore,

$$\mathbf{E}_{\mathbf{x}^{-i}, \mathbf{c}^i}[X_{k+1}] = \mathbf{E}_{\mathbf{x}^{-i}, \mathbf{c}^i}[X_k] + 7\rho N \varepsilon^a \mathbf{P}_{\mathbf{x}^{-i}, \mathbf{c}^i}(t \in B_k) + 2\varepsilon \mathbf{P}_{\mathbf{x}^{-i}, \mathbf{c}^i}(t \in n_k),$$

which yields by summation over k

$$\sup_k \mathbf{E}_{\mathbf{x}^{-i}, \mathbf{c}^i}[X_k] = X_1 + 2\varepsilon \sum_{k=1}^{\infty} \mathbf{P}_{\mathbf{x}^{-i}, \mathbf{c}^i}(t \in n_k) + 7\rho N \varepsilon^a. \quad (25)$$

We use now the fact that consecutive blocks of type I are never distant by more than $1/\varepsilon^e$ blocks. Thus, in the first p/ε^e blocks, there are at least $p-1$ blocks of type I. For such a block B_k , $\mathbf{P}_{\mathbf{x}^{-i}, \mathbf{c}^i}\{t \in n_{k+1} | t \in n_k\} = 1 - \varepsilon^b$. Therefore,

$$\mathbf{P}_{\mathbf{x}^{-i}, \mathbf{c}^i}(t \in n_{\frac{p}{\varepsilon^e}}) = (1 - \varepsilon^b)^{p-1}.$$

Since $\mathbf{P}_{\mathbf{x}^{-i}, \mathbf{c}^i}(t \in n_k) = \mathbf{P}_{\mathbf{x}^{-i}, \mathbf{c}^i}(t \in n_{\frac{p}{\varepsilon^e}})$ for $\frac{p}{\varepsilon^e} - k < \frac{p+1}{\varepsilon^e}$, one obtains

$$\varepsilon \sum_{k=1}^{\infty} \mathbf{P}_{\mathbf{x}^{-i}, \mathbf{c}^i}(t \in n_k) \leq 2\varepsilon^{1-b-e}. \quad (26)$$

The conclusion of the lemma follows from (25) and (26), since $X_1 = \gamma^i(\mathbf{x})$. ■

By Lemmas 2.11 and 2.12, \mathbf{x} is a $\varepsilon^{\frac{1}{6}}$ -equilibrium of the quitting game. This ends the proof of Proposition 2.4.

2.6 Equilibrium and uniformity

In the literature on stochastic games a stronger concept of uniform equilibrium is often used. In this section we prove that in quitting games, any ε -equilibrium is a uniform ε -equilibrium.

For every $n \in \mathbf{N}$ define

$$\gamma_n(\mathbf{x}) = \mathbf{E}_{\mathbf{x}} \left[1_{t \leq n} r S_t \frac{n-t}{n} \right].$$

This is the expected average payoff if the players receive a payoff equal to zero prior to the termination stage, and the termination payoff in every subsequent stage.

A profile \mathbf{x} is a uniform ε -equilibrium if it is an ε -equilibrium for each payoff function γ_n , provided n is sufficiently large. That is, there exists $n_0 \in \mathbf{N}$ such that for every $n \geq n_0$, every $i \in \mathbf{N}$ and every strategy \mathbf{y}^i of i ,

$$\gamma_n^i(\mathbf{x}) - \gamma_n^i(\mathbf{x}^{-i}, \mathbf{y}^i) \leq \varepsilon.$$

We now prove that the requirement of uniformity has no bite in the context of quitting games. Note that uniform ε -equilibrium is a stronger concept than ε -equilibrium.

PROPOSITION 2.13 *If \mathbf{x} is an ε -equilibrium then it is a uniform ε' -equilibrium, provided $\varepsilon' > \varepsilon$.*

Proof: Since the details are standard, we only sketch the proof. Let \mathbf{x} be an ε -equilibrium, $i \in \mathbf{N}$ and $\varepsilon' > \varepsilon$. We shall prove that, for n large, player i can not profit more than ε' in the game with payoff function γ_n . Note that for every profile \mathbf{y} , $\gamma_n^i(\mathbf{y}) \leq \gamma^i(\mathbf{y})$.

Assume first that $(\mathbf{x}^{-i}, \mathbf{c}^i)$ is terminating. There exists N_0 such that for every $n \geq N_0$ and every strategy \mathbf{y}^i ,

$$|\gamma(\mathbf{x}^{-i}, \mathbf{y}^i) - \gamma_n(\mathbf{x}^{-i}, \mathbf{y}^i)| < (\varepsilon' - \varepsilon)/2.$$

By the ε -equilibrium property, one deduces $\gamma_n^i(\mathbf{x}^{-i}, \mathbf{y}^i) \leq \gamma_n^i(\mathbf{x}) + \varepsilon'$ for every $n \geq N_0$.

Assume now that $\mathbf{P}_{\mathbf{x}^{-i}, \mathbf{c}^i}(t < +\infty) < 1$. Choose N_0 such that

$$\mathbf{P}_{\mathbf{x}^{-i}, \mathbf{c}^i}(t < +\infty | t \geq N_0) \leq (\varepsilon' - \varepsilon)/3N\rho. \quad (27)$$

In particular, for each $n \geq N_0$, x_n^{-i} is close to c^{-i} . Let \mathbf{s}^i be any pure strategy of player i . We now estimate $\gamma_n^i(\mathbf{x}^{-i}, \mathbf{s}^i)$ for n sufficiently large.

- Assume $\mathbf{s}^i = \mathbf{c}^i$.

Since $\gamma^i(\mathbf{x}^{-i}, \mathbf{c}^i) \geq \gamma^i(\mathbf{x}) + \varepsilon$, one has $\gamma_n^i(\mathbf{x}^{-i}, \mathbf{c}^i) \geq \gamma_n^i(\mathbf{x}) + \varepsilon'$, for n large enough.

- Assume $\mathbf{s}^i = \mathbf{q}_{n_0}^i$, where $n_0 \geq N_0$.

In that case, $\|\gamma_n(\mathbf{x}^{-i}, \mathbf{q}_{n_0}^i) - \gamma(\mathbf{x}^{-i}, \mathbf{q}_{n_0}^i)\|$ converges to zero, uniformly in n_0 . Hence $\gamma_n^i(\mathbf{x}^{-i}, \mathbf{q}_{n_0}^i) \geq \gamma_n^i(\mathbf{x}) + (\varepsilon' - \varepsilon)/3$, for $n \geq N_1$, where $N_1 \geq 3N_0\rho/(\varepsilon' - \varepsilon)$.

- Finally, assume $\mathbf{s}^i = \mathbf{q}_{n_0}^i$, where $n_0 \geq N_0$.

By (27), for each $n \geq N_0$ the expected payoff of player i , conditional on $\{t = n_0\}$ satisfies

$$\|\mathbf{E}_{\mathbf{x}^{-i}, \mathbf{q}_{n_0}^i} [r_{S_t}^i | t = n_0] - 1\| < (\varepsilon' - \varepsilon)/3.$$

By the choice of N_0 , $\gamma_n^i(\mathbf{x}^{-i}, \mathbf{q}_{n_0}^i) \geq \gamma_n^i(\mathbf{x}^{-i}, \mathbf{q}_{N_0}^i) + (\varepsilon' - \varepsilon)/3$, for every $n \geq N_1$. The result follows from the second case with $n_0 = N_0$. ■

2.7 General payoff processes

We give here a slight extension of Theorem 1.2, within the framework of non zero-sum Dynkin games. Let $\mathbf{r} = (r_n)_{n \geq 1}$ be a process over $(\Omega, \mathbf{A}, \mathbf{P})$, where $r_n = (r_{n,S})_{\emptyset \neq S \subseteq \mathcal{N}}$ is a vector of \mathbf{R}^N -valued variables. The quitting game $\Gamma(\mathbf{r})$ is played as above: $r_{n,S}$ is the payoff vector if termination occurs at stage n , and the quitting coalition is S .

Set $H_n = \sigma(r_p, p \dots n)$. A strategy of player i is a process $\mathbf{x}^i = (x_n^i)_{n \geq 1}$ adapted to $(H_n)_{n \geq 1}$, with values in $[0, 1]$. Provided expectations are well-defined, the extension of the above definitions of γ and ε -equilibrium is straightforward.

THEOREM 2.14 *Assume that the sequence (r_n) converges to r_∞ , \mathbf{P} -a.s. Assume r_∞ satisfies assumptions **A.1** and **A.2** of Theorem 1.2, \mathbf{P} -a.s. Assume $\rho = \mathbf{E}[\sup_{n \geq 1} |r_n|] < +\infty$. Then, for every $\varepsilon > 0$, $\Gamma(\mathbf{r})$ admits an ε -equilibrium.*

Proof: The idea is the following. Take N_0 large enough, so that r_{N_0} and r_∞ are close. At stage N_0 , players start using an η -equilibrium of the quitting game with constant payoffs $\Gamma(r_{N_0}, r_{N_0}, \dots)$. Behavior in the first $N_0 - 1$ stages is then defined by backwards induction.

We address first the measurability issue. Set $m = N(2^N - 1)$. Denote by $X = [0, 1]^{N \times N}$ the space of strategy profiles in a quitting game with constant payoffs. Denote by $\Delta \subset \mathbf{R}^m$ the set of vectors $r = (r_S)_{\emptyset \neq S \subseteq N}$ which satisfy **A.1** and **A.2**. For $\eta > 0$, and $r \in \mathbf{R}^m$, denote by Δ_η the η -neighborhood of Δ , and by $E_\eta(r) \subset X$ the set of η -equilibria of the quitting game with constant payoffs $\Gamma(r, r, \dots)$. By Theorem 1.2, $E_\eta(r) \neq \emptyset$, for every $r \in \Delta$, $\eta > 0$.

It is clear that whenever $\|r - r'\| < \eta$, $E_\eta(r) \subset E_{2\eta}(r')$. Therefore, there is a measurable step function $\sigma_\eta : \mathbf{R}^m \rightarrow X$, with $\sigma_\eta(r) \in E_{2\eta}(r)$, for every $r \in \Delta_\eta$.

Choose $0 < \eta < \frac{1}{4+\rho}$ and N_0 such that

$$\mathbf{P}(\exists n \leq N_0, |r_n - r_\infty| > \eta) < \eta. \quad (28)$$

We now construct a profile \mathbf{x} . Set $\mathbf{x}_{N_0} = (x_{N_0}, x_{N_0+1}, \dots) = \sigma_\eta(r_{N_0})$. By construction, x_n is H_{N_0} -measurable, for $n \geq N_0$. We define x_{N_0-1}, \dots, x_1 inductively: for $k \geq N_0 - 1$, define x_k to be a H_k -measurable equilibrium in $G(\mathbf{E}[\gamma(\mathbf{x}_{k+1})/H_k])$.

Such a choice for x_k exists. Indeed, let R be a H_k -measurable correspondence, defined as the set of all mixed Nash equilibria in $G(\mathbf{E}[\gamma(\mathbf{x}_{k+1})/H_k])$. Then R has non-empty and closed values. By Kuratowski and Ryll-Nardzewski (1965) R has a H_k -measurable selection.

By (28), $\mathbf{P}(\mathbf{x}_{N_0} \in E_{3\eta}(r_{N_0}, r_{N_0+1}, \dots)) \geq 1 - \eta$: the probability of \mathbf{x}_{N_0} being a 3η -equilibrium is at least the probability that r_n will remain 2η -close to r_{N_0} , which is at least $1 - \eta$ by the choice of N_0 .

This easily yields that \mathbf{x} is a $(3\eta + \eta\rho)$ -equilibrium of $\Gamma(\mathbf{r})$. ■

3 An Example

In this section we study the following four player quitting game, which satisfies conditions **A.1** and **A.2**.

	2	4	2									
1	<table style="border-collapse: collapse; width: 100%; height: 100%;"> <tr> <td style="border: none; padding-right: 5px;"><i>continue</i></td> <td style="border: none; padding-left: 5px;">4, 1, 0, 0</td> </tr> <tr> <td style="border: none; padding-right: 5px;">1, 4, 0, 0</td> <td style="border: none; padding-left: 5px;">1, 1, 1, 1</td> </tr> </table>	<i>continue</i>	4, 1, 0, 0	1, 4, 0, 0	1, 1, 1, 1	1	<table style="border-collapse: collapse; width: 100%; height: 100%;"> <tr> <td style="border: none; padding-right: 5px;">0, 0, 4, 1</td> <td style="border: none; padding-left: 5px;">1, 1, 0, 1</td> </tr> <tr> <td style="border: none; padding-right: 5px;">1, 0, 1, 1</td> <td style="border: none; padding-left: 5px;">0, 1, 0, 0</td> </tr> </table>	0, 0, 4, 1	1, 1, 0, 1	1, 0, 1, 1	0, 1, 0, 0	
<i>continue</i>	4, 1, 0, 0											
1, 4, 0, 0	1, 1, 1, 1											
0, 0, 4, 1	1, 1, 0, 1											
1, 0, 1, 1	0, 1, 0, 0											
3	1	<table style="border-collapse: collapse; width: 100%; height: 100%;"> <tr> <td style="border: none; padding-right: 5px;">0, 0, 1, 4</td> <td style="border: none; padding-left: 5px;">0, 1, 1, 1</td> </tr> <tr> <td style="border: none; padding-right: 5px;">1, 1, 1, 0</td> <td style="border: none; padding-left: 5px;">1, 0, 0, 0</td> </tr> </table>	0, 0, 1, 4	0, 1, 1, 1	1, 1, 1, 0	1, 0, 0, 0	1	<table style="border-collapse: collapse; width: 100%; height: 100%;"> <tr> <td style="border: none; padding-right: 5px;">1, 1, 1, 1</td> <td style="border: none; padding-left: 5px;">0, 0, 1, 0</td> </tr> <tr> <td style="border: none; padding-right: 5px;">0, 0, 0, 1</td> <td style="border: none; padding-left: 5px;">-1, -1, -1, -1</td> </tr> </table>	1, 1, 1, 1	0, 0, 1, 0	0, 0, 0, 1	-1, -1, -1, -1
0, 0, 1, 4	0, 1, 1, 1											
1, 1, 1, 0	1, 0, 0, 0											
1, 1, 1, 1	0, 0, 1, 0											
0, 0, 0, 1	-1, -1, -1, -1											

Figure 2

In this game player 1 chooses a row (top row = continue), player 2 chooses

a column (left column = continue), player 3 chooses either the top two matrices or the bottom two matrices, (top two matrices = continue) and player 4 chooses either the left two matrices or the right two matrices (left two matrices = continue).

Note that there are the following symmetries in the payoff function: for every 4-tuple of actions (a, b, c, d) we have:

$$\begin{aligned} v^1(a, b, c, d) &= v^2(b, a, d, c), \\ v^1(a, b, c, d) &= v^4(c, d, b, a) \quad \text{and} \\ v^2(a, b, c, d) &= v^3(c, d, b, a), \end{aligned}$$

where $v^i(a, b, c, d)$ is the payoff to i if the action combination is (a, b, c, d) ($v^i(c^1, c^2, c^3, c^4) = 0$).

One can show that the game does not admit any stationary ε -equilibrium, nor an ε -equilibrium \mathbf{x} such that $\|x_n - c\| < \varepsilon$ for every n . Since the details are technical we omit them. The interested reader may consult Solan and Vieille (2000).

We now prove that the game possesses a cyclic equilibrium, where the length of the cycle is 2. At odd stages players 2 and 4 play c^2 and c^4 respectively, and players 1 and 3 continue with probability x and z respectively, both strictly less than 1. At even stages players 1 and 3 play c^1 and c^3 respectively, and players 2 and 4 continue with probability x and z respectively.

Thus, in some sense, the “simplest” equilibrium in this game is periodic with period 2.

Formally, we study now profiles \mathbf{y} that satisfy:

$$y_n = \begin{cases} (x, 1, z, 1) & n \text{ odd} \\ (1, x, 1, z) & n \text{ even} \end{cases}$$

where $x, z \in]0, 1[$ are independent of n .

The one-shot game played by players 1 and 3 at odd stages is

		3	
		z	$1 - z$
1	x	γ_c^1, γ_c^3	$0, 1$
	$1 - x$	$1, 0$	$1, 1$

Figure 3: The game of players 1 and 3 at odd stages

In this game player 1 is the row player, player 3 is the column player, and γ_c^i is the continuation payoff of player $i = 1, 3$. The payoffs received by players 2 and 4 if termination occurs in an odd stage are given by the matrix below, in which the first coordinate of each entry is player 2's payoff, and the second coordinate is player 4's payoff.

	$0, 4$	(29)
$4, 0$	$1, 0$	

The one-shot game played by players 2 and 4 at even stages is

		4	
		z	$1 - z$
2	x	γ_c^2, γ_c^4	$0, 1$
	$1 - x$	$1, 0$	$1, 1$

Figure 4: The game of players 2 and 4 at even stages

where player 2 is the row player, player 4 is the column player, and the payoffs that are received by players 1 and 3 if termination occurs are given by matrix (29). The two situations are identical (up to the continuation payoffs).

We now find necessary conditions on (x, z) . First, (x, z) is a fully mixed equilibrium of the matrix game in Figure (3), so that

$$x\gamma_c^3 = 1 \text{ and } z\gamma_c^1 = 1,$$

and both players 1 and 3 receive 1 in this equilibrium.

By the symmetry of the profile, the continuation payoffs (resp. initial payoffs) of players 2 and 4 must coincide with the initial payoffs (resp. continuation payoffs) of players 1 and 3. That is, (γ_c^1, γ_c^3) is the payoff received in the matrix game (29), when the empty entry is filled with $(1, 1)$ and the row and column players play according to x and z respectively, so that

$$\begin{cases} \gamma_c^1 = xz + 4z(1 - x) + (1 - x)(1 - z) \\ \gamma_c^3 = xz + 4z(1 - x) \end{cases}$$

Set $g = \gamma_c^1$ and $h = \gamma_c^3$. Since $x = \frac{1}{h}$ and $z = \frac{1}{g}$, one gets

$$\begin{cases} g^2h = 1 + 4(g - 1) + (g - 1)(h - 1) \\ gh^2 = 1 + 4(g - 1) \end{cases}$$

which is equivalent to

$$\begin{cases} g = \frac{3}{4-h^2} \\ h \text{ root of } (h - 1)(h^4 + 3h^3 - 2h^2 - 9h + 4) = 0 \end{cases} \quad (30)$$

Conversely, let (g, h) be a solution to (30) with $g, h > 1$, and define a cyclic profile by $x = \frac{1}{h}$, $z = \frac{1}{g}$. Given the above properties, in order to prove that it is an equilibrium, we need only prove that neither player 2 nor 4 can find it profitable to quit in the first stage. This is clear, since players 2 and 4 would receive at most 1 by quitting, whereas they get strictly more than 1 under the cyclic profile.

Thus, the existence of such a cyclic equilibrium is equivalent to the existence of a solution (g, h) to system (30) with $g, h > 1$. If $1 < h < 2$ then $1 < 3/(4 - h^2)$. Hence we need to assert the existence of a root in $]1, 2[$ of the polynomial

$$Q(X) = X^4 + 3X^3 - 2X^2 - 9X + 4.$$

Such a root exists since $Q(1) < 0 < Q(2)$.

Appendix A

The following lemma has been used in the proof of Lemma 2.6. It is independent of any other result in the paper.

LEMMA 3.1 *Let $\bar{N} \in \mathbf{N}$, and $(X_0, \dots, X_{\bar{N}}, Y_0, \dots, Y_{\bar{N}})$ be independent $\{0, 1\}$ -valued random variables. Let $S_1 = \inf\{n \leq \bar{N}, X_n = 1\}$, $S_2 = \inf\{n \leq \bar{N}, Y_n = 1\}$ be the first successes of the two sequences. Set $T = S_1$ if $S_1 < S_2$ and $T = +\infty$ otherwise. Assume that $\mathbf{P}(T \leq \bar{N}) > 0$. Then*

1. $\mathbf{P}(S_1 \leq \bar{N}) - \mathbf{P}(T \leq \bar{N}) = \mathbf{P}(S_1 \leq \bar{N})\mathbf{P}(S_2 \leq \bar{N})$;
2. $\sum_n \mathbf{P}(S_1 = n | S_1 \leq \bar{N}) - \mathbf{P}(S_1 = n | T \leq \bar{N}) = 2\mathbf{P}(S_2 \leq \bar{N})$.

Proof: For each n , $\{T = n\} \subseteq \{S_1 = n\}$, and $\{S_1 = n\} \setminus \{T = n\} = \{S_1 = n\} \cap \{S_2 \leq n\}$. Thus

$$\mathbf{P}(S_1 = n) - \mathbf{P}(T = n) = \mathbf{P}(S_1 = n)\mathbf{P}(S_2 \leq n) - \mathbf{P}(S_1 = n)\mathbf{P}(S_2 \leq \bar{N}). \quad (31)$$

The first claim follows by summation over n .

On the other hand,

$$\begin{aligned} |\mathbf{P}(S_1 = n | S_1 \leq \bar{N}) - \mathbf{P}(S_1 = n | T \leq \bar{N})| &= \left| \frac{\mathbf{P}(S_1 = n)}{\mathbf{P}(S_1 \leq \bar{N})} - \frac{\mathbf{P}(S_1 = n < S_2)}{\mathbf{P}(T \leq \bar{N})} \right| \\ &= \left| \frac{\mathbf{P}(S_1 = n)\mathbf{P}(S_2 \leq n)}{\mathbf{P}(S_1 \leq \bar{N})} + \mathbf{P}(S_1 = n < S_2) \left(\frac{1}{\mathbf{P}(S_1 \leq \bar{N})} - \frac{1}{\mathbf{P}(T \leq \bar{N})} \right) \right|. \end{aligned}$$

Observe that

$$\frac{\mathbf{P}(S_1 = n)\mathbf{P}(S_2 \leq n)}{\mathbf{P}(S_1 \leq \bar{N})} = \frac{\mathbf{P}(S_1 = n)}{\mathbf{P}(S_1 \leq \bar{N})}\mathbf{P}(S_2 \leq \bar{N}). \quad (32)$$

On the other hand,

$$\begin{aligned}
& \left| \mathbf{P}(S_1 = n < S_2) \left(\frac{1}{\mathbf{P}(S_1 \leq \bar{N})} - \frac{1}{\mathbf{P}(T \leq \bar{N})} \right) \right| \\
&= \frac{\mathbf{P}(S_1 = n < S_2)}{\mathbf{P}(S_1 \leq \bar{N})\mathbf{P}(T \leq \bar{N})} |\mathbf{P}(T \leq \bar{N}) - \mathbf{P}(S_1 \leq \bar{N})| \\
&= \frac{\mathbf{P}(S_1 = n < S_2)}{\mathbf{P}(S_1 \leq \bar{N})\mathbf{P}(T \leq \bar{N})} \times \mathbf{P}(S_2 \leq \bar{N})\mathbf{P}(S_1 \leq \bar{N}) = \frac{\mathbf{P}(T = n)}{\mathbf{P}(T \leq \bar{N})} \mathbf{P}(S_2 \leq \bar{N}),
\end{aligned} \tag{33}$$

using the first part of the lemma. The result follows from (32) and (33). ■

References

- [1] E.G. Dynkin. Game variant of a problem on optimal stopping. *Soviet Math. Dokl.*, 10:270-274, 1969.
- [2] J. Flesch, F. Thuijsman, and O.J. Vrieze. Recursive repeated games with absorbing states. *Math. Oper. Res.*, 21:1016-1022, 1996.
- [3] J. Flesch, F. Thuijsman, and O.J. Vrieze. Cyclic Markov equilibria in stochastic games. *Inter. J. Game Theory*, 26:303-314, 1997.
- [4] D. Fudenberg and J. Tirole. *Game Theory*. The MIT Press, 1991
- [5] P. Ghemawat and B. Nalebuff. Exit. *RAND J. Econ.*, 16(2):184–194, 1985.
- [6] P. Hammerstein and R. Selten. Game theory and evolutionary biology. In R. J. Aumann and S. Hart, editors, *Handbook of Game Theory with Economic Applications*, volume 2, pages 929–993. Elsevier Science Publishers B.V., 1995.

- [7] V. Krishna and J. Morgan. An Analysis of the War of Attrition and the All-Pay Auction, *J. Econ. Theory*, 72(2):343-362.
- [8] K. Kuratowski and C. Ryll-Nardzewski. A general theorem on selectors. *Bull. Acad. Polon. Sci. Ser. Sci. Math. Astronom. Phys.*, 13:397-403, 1965.
- [9] L. Li. Equilibrium exit in stochastically declining industries. *Games Econ. Behavior*, 1(1):40–59, 1989.
- [10] J. Maynard-Smith. The Theory of Games and the Evolution of Animal Conflicts. *Journal of Theoretical Biology*, 47:209-221, 1974.
- [11] H. Morimoto. Non zero-sum discrete parameter stochastic games with stopping times. *Probab. Th. Rel. Fields*, 72:155-160, 1986.
- [12] Y. Ohtsubo. A non zero-sum extension of Dynkin’s stopping problem. *Math. Oper. Res.*, 12(2):277-296, 1987.
- [13] E. Solan. Three-person Absorbing Games games, *Math. Oper. Res.*, 24(3):669-698, 1999.
- [14] Solan E. and Vieille N. (2000), Quitting Games, Cahiers du CEREMADE; <http://www.math.tau.ac.il/~eilons/>
- [15] S. Sorin. Asymptotic Properties of a Non-zerosum Stochastic Games, *Int. J. Game Theory*, 15, 101-107, 1986.
- [16] N. Vieille. Contributions à la théorie des jeux répétés. *Thèse de doctorat, Université Paris 6*, 1992.

- [17] R. Wilson. Strategic models of entry deterrence. In R. J. Aumann and S. Hart, editors, *Handbook of Game Theory with Economic Applications*, volume 1, pages 305-354. Elsevier Science Publishers B.V., 1992.