Subgame Perfect Equilibria in Stopping Games

Ayala Mashiah-Yaakovi[*][†]

April 20, 2009

Abstract

Stopping games (without simultaneous stopping) are sequential games in which at every stage one of the players is chosen according to a stochastic process, and that player decides whether to continue the interaction or stop it, whereby the terminal payoff vector is obtained by another stochastic process.
We prove that if the payoff process is integrable, a δ-approximate subgame perfect ϵ-equilibrium exists; that is, there exists a strategy profile that is an ϵ-equilibrium in all subgames, except possibly in a set of subgames that occurs with probability smaller than δ (even after deviation by some of the players).

1 Introduction

Stopping games (without simultaneous stopping) are n-player sequential games in which, at every stage, one player is chosen according to a stochastic process, and that player decides, whether to continue the game or to stop it. Once the chosen player decides to stop, the players receive a terminal payoff that is determined by a second stochastic process.

[*]School of Mathematical Sciences, Tel Aviv University, Tel Aviv 69978, Israel. e-mail: ayalam@post.tau.ac.il .

[†]This work is part of the Ph.D. thesis of the author done under the supervision of Prof. Eilon Solan. I owe special thanks to Prof. Solan for his valuable guidance, and infinite patience.

Stopping games have been introduced by Dynkin (1969), who studied two-player zero-sum games with bounded payoffs. Dynkin proved the existence of the value and pure optimal strategies in that game.

Dynkin's result has been extended to the case in which players can stop simultaneously (see, e.g., Kifer (1971), Neveu (1975), Rosenberg, Solan and Vieille (2001)).

Multi player stopping games (without simultaneous stopping) are a subclass of sequential games with perfect information. By Mertens (1987) it follows that every such game has an ϵ-equilibrium. The ϵ-equilibrium strategies that were constructed by Mertens (1987) employ threats of punishment, which might be non-credible. Stopping games were used to model, e.g., exit from shrinking markets (Fine and Li (1989), Ghemawat and Nalebuff (1985)), duels (Karlin (1959)), and investments (Kifer (2000)). In such applications, it is not clear whether the players will implement a Nash equilibrium that involves non-credible threats. Recent work concentrates on the existence of subgame perfect equilibria. The question whether every sequential game with perfect information has a subgame perfect equilibrium is still open and the existence was proved only for some subclasses of these games.

Solan and Vieille (2003) studied multi-player stopping games, where the order by which players are chosen is deterministic, and the probability that the game terminates once the chosen player decides to stop may be strictly less than 1. They proved that this game has a subgame perfect ϵ-equilibrium in Markovian strategies. Furthermore, if the game is not degenerate, this ϵ-equilibrium is actually in pure strategies.

Solan (2005) studied an n-player game in which both the terminal payoff process and the process by which players are chosen are stationary. Solan proved the existence of either a stationary ϵ-equilibrium or a subgame perfect 0-equilibrium.

Mashiah-Yaakovi (2008) generalized Solan's (2005) result to the case where both of the processes that define the stopping game are Markovian and periodic, rather than stationary. Mashiah-Yaakovi proved the existence of either a periodic subgame perfect ϵ-equilibrium or a subgame perfect 0-equilibrium in pure strategies.

Maitra and Sudderth (2007) present sufficient conditions for the existence of subgame perfect equilibria in multi player stochastic games with Borel state space and compact metric action sets. Their conditions do not hold when the payoff is undiscounted.

Recently, Flesch et al. (2008) proved the existence of a subgame perfect ϵ-equilibrium in every n-player limiting average recursive games with perfect information in which the payoffs in all absorbing state are non-negative.[1]

In a general stopping game there might be a continuum of subgames. Furthermore, the expected continuation payoff at a given stage k is a measurable function (with respect to the σ-algebra generated by the play up to stage k), and therefore it is defined almost surely, and not for every history. Hence, the traditional concept of a subgame perfect equilibrium should be adapted.

In this paper we define a variant of the concept of subgame perfect equilibrium, a **δ-approximate subgame perfect ϵ-equilibrium**, which is appropriate to stopping games. A strategy profile σ is a δ-approximate subgame perfect ϵ-equilibrium if there is an event G that occurs with probability smaller than δ, such that for every stage $K \in \mathbf{N}$, and every event $F \in \mathcal{F}_k$ in the complement of G that occurs with positive probability, no player can gain more than ϵ by deviating in the game that starts at stage K, conditioned that the event F occurs. That is, σ induces an ϵ-equilibrium in every subgame, except perhaps a set of subgames that occur with probability smaller than δ.

We show that every stopping game (without simultaneous stopping) has a δ-approximate subgame perfect ϵ-equilibrium, under merely an integrability condition on the payoff process.

The structure of the proof is similar to that of Shmaya and Solan (2004), who proved that every two-player stopping game (with simultaneous stopping) has an ϵ-equilibrium in randomized stopping times. Their proof is

[1] A state s in a stochastic game is called absorbing, if once the game reaches that state, it stays there forever, whatever the players play; otherwise it is called non-absorbing. A stochastic game is called recursive if the stage payoff in every non-absorbing state is 0.

based on a stochastic variation of Ramsey Theorem, which allows to reduce the problem of existence of an ϵ-equilibrium in a general two-player stopping game to that of studying properties of ϵ-equilibria in a simple class of stochastic games with finite state space: the class of two-player absorbing games, that always have stationary ϵ-equilibria (Flesch, Thuijsman and Vrieze (1997)).

In the model that we study the approximating game is equivalent to a multi-player absorbing game, yet stationary equilibria may not exist in such games, and it was not known whether such games have subgame perfect ϵ-equilibria. The core of our proof is that indeed subgame perfect ϵ-equilibria exist in this class of games. Furthermore, under some sufficient conditions, one can bound from below the probability of termination in each period of the game, under these ϵ-equilibria, a property which is needed for using the technique of Shmaya and Solan.[2]

The paper is organized as follows: In Section 2 we present the model, some basic definitions, and the main result. In Section 3 we study the class of periodic stopping games whose filtration consists of finite σ-algebra, and we prove that games in this class have a subgame perfect ϵ-equilibrium. The proof of the main result appears in Section 4.

2 The model and the main results

2.1 The model

DEFINITION 2.1 *A* **stopping game** *is given by* $\Gamma = (I, \Omega, \mathcal{A}, \mathbf{P}, \mathcal{F}, (i_k)_{k=1}^{\infty}, (a_k)_{k=1}^{\infty}, a_\infty)$ *where:*

- $I = \{1, ..., n\}$ *is a non-empty finite set of players.*
- $(\Omega, \mathcal{A}, \mathbf{P})$ *is a probability space.*
- $\mathcal{F} = (\mathcal{F}_k)_{k=1}^{\infty}$ *is a filtration over* $(\Omega, \mathcal{A}, \mathbf{P})$*, representing the information available to the players at stage* k*.*

[2]Flesch et al. (2008) study recursive games with positive payoffs, whereas we study stopping games with general payoff process. It is therefore not clear whether their result can be used in our proof.

- *$(i_k)_{k=1}^{\infty}$ and $(a_k)_{k=1}^{\infty}$ are $\mathcal{F}$-adapted processes. $(i_k)_{k=1}^{\infty}$ is an I-valued process, which indicates the player who decides whether to stop the game or to continue at stage k. $(a_k)_{k=1}^{\infty}$ is a $\Re^n$-valued process, which indicates the terminal payoff if player i_k stops.*

- *$a_\infty \in L^1(\mathbf{P})$ is a payoff function, representing the payoff if no player ever stops at stage k.*

The game is played as follows: An element $\omega \in \Omega$ is chosen according to $\mathbf{P}$. At every stage $k \in \mathbf{N}$, the players learn the atom of $\mathcal{F}_k$ that contains ω, and player $i_k(\omega)$ decides whether to stop the game or to continue.[3] If player $i_k(\omega)$ decides to stop, the game terminates with terminal payoff vector $a_k(\omega)$. If player $i_k(\omega)$ decides to continue, the play continues to stage $k+1$. If the game never terminates, the payoff is $a_\infty(\omega)$.

2.2 Strategies and equilibria

To save notations, we assume that each player i chooses actions in stages in which $i_k = i$, even after the game terminates.

Definition 2.2 *A* **pure strategy** *for player $i \in I$ is a $\{0,1\}$-valued $\mathcal{F}$-adapted process $\sigma^i := (\sigma_k^i)_{k=1}^{\infty}$. $\sigma_k^i(\omega) = 1$ if player i stops the game at ω when chosen at stage k (provided the game did not terminate before that stage), while $\sigma_k^i(\omega) = 0$ if player i continues at ω when chosen at stage k.*

We denote by 0^i the strategy of player i in which he continues whenever he is chosen.

Definition 2.3 *A* **(behavior) strategy** *for player i is a $[0,1]$-valued $\mathcal{F}$-adapted process $\sigma^i = (\sigma_k^i)_{k=1}^{\infty}$. $\sigma_k^i(\omega)$ is the probability that player i stops at ω when chosen at stage k (provided the game did not terminate before that stage).*

A **profile** is a vector of strategies, one for each player. We denote by σ^{-i} the vector of strategies of all the players excluding player i.

Let θ be the termination stage, that is, the first stage k in which player i_k chooses to stop. In case the game never terminates we set $\theta = +\infty$. Note

[3] Formally, for every $A \in \mathcal{F}_k$ the players learn whether $\omega \in A$ or $\omega \notin A$.

that a_θ is the payoff in the game.

A profile σ is called **terminating** if $\mathbf{P}_\sigma(\theta < \infty) = 1$; Namely, under σ, the game terminates with probability 1.

A **play** is given by ω together with an infinite sequence of players' actions. Notice that the play is infinite even when θ is finite, since players choose actions even after the game terminates.

Each profile σ together with $\mathbf{P}$ induces a distribution $\mathbf{P}_\sigma$ over the set of plays. Let E_σ be the expectation operator that corresponds to $\mathbf{P}_\sigma$. The expected payoff vector under σ is

$$\gamma(\sigma) := E_\sigma[a_\theta].$$

For every $F \in \mathcal{A}$ such that $\mathbf{P}(F) > 0$, denote by $\gamma_{|F}(\sigma)$ the conditional expected payoff vector under σ given F occurs, that is

$$\gamma_{|F}(\sigma) = E_\sigma[a_\theta | F].$$

DEFINITION 2.4 *Let $\epsilon \geq 0$, and $F \in \mathcal{A}$ such that $\mathbf{P}(F) > 0$. A profile σ is an* ϵ**-equilibrium given** F *if for every player $i \in I$ and every strategy $\overline{\sigma}^i$ of player i,*

$$\gamma^i_{|F}(\sigma) \geq \gamma^i_{|F}\left(\sigma^{-i}, \overline{\sigma}^i\right) - \epsilon.$$

The vector $\gamma_{|F}(\sigma)$ is called an ϵ**-equilibrium payoff vector given** F.

In particular, σ is an ϵ-equilibrium if and only if it is an ϵ-equilibrium given Ω.

For every $K \in \mathbf{N}$ define the game that starts at stage K, by

$$\Gamma_{|K} := \left(I, \Omega, \mathcal{A}, \mathbf{P}, (\mathcal{F}_k)_{k=K}^\infty, (i_k)_{k=K}^\infty, (a_k)_{k=K}^\infty, a_\infty\right).$$

Every strategy σ^i of player i in Γ induces a strategy $\sigma^i_{|K}$ in $\Gamma_{|K}$, by ignoring the play in the first $K-1$ stages.

In finite games in extensive form, a profile is a subgame perfect equilibrium if it induces an equilibrium in every subgame. In a stopping game, there is a continuum of subgames. Therefore, it is natural to define the concept of subgame perfect equilibrium, such that for every $K \in \mathbf{N}$, and every event $F \in \mathcal{F}_K$ such that $\mathbf{P}(F) > 0$, the profile induces an equilibrium in the subgame which starts at stage K given F occurs. We will study a weaker concept of subgame perfect equilibrium, a **δ-approximate subgame perfect ϵ-equilibrium**, which requires that the induced strategy be an ϵ-equilibrium in every subgame that starts at stage K, given event F occurs, for every $K \in \mathbf{N}$, and every event $F \in \mathcal{F}_K$ such that $\mathbf{P}(F) > 0$, except possibly a set of subgames which occurs with probability smaller than δ.

Definition 2.5 *Let $\epsilon, \delta \geq 0$. A profile σ is a* **δ-approximate subgame perfect ϵ-equilibrium** *if there is an event $G \in \mathcal{A}$ with $\mathbf{P}(G) < \delta$, such that for every stage $K \in \mathbf{N}$, and every event $F \in \mathcal{F}_k$, such that $\mathbf{P}(F) > 0$ and $F \cap G = \emptyset$, $\sigma_{|K}$ is an ϵ-equilibrium in $\Gamma_{|K}$ given F.*

2.3 The main result

The main result of the paper is:

Theorem 2.6 *Every stopping game such that $\|a_\infty\|_\infty, \sup_{k\in\mathbf{N}} \|a_k\|_\infty \in L^1(\mathbf{P})$ has a δ-approximate subgame perfect ϵ-equilibrium, for every $\delta, \epsilon > 0$.*

It is not known whether the game has a subgame perfect ϵ-equilibrium. The rest of the paper is devoted to the proof of Theorem 2.6. We first provide a sketch of the proof.

To simplify presentation, suppose that each σ-algebra $\mathcal{F}_k$ is finite. For every positive integer k, every stopping time $\tau > k$, every atom F of $\mathcal{F}_k$, and every $\omega \in F$, define a periodic stopping game $\Gamma_{k,\tau}(\omega)$ that starts at stage k, and if no player stops before stage τ, then it restart at stage k with a new state $\omega' \in F$ that is chosen according to $\mathbf{P}$ conditioned on F. The game $\Gamma_{k,\tau}(\omega)$ is a finite stochastic game with perfect information. In Section 3 we prove that the game has a subgame perfect ξ-equilibrium with special properties. If $\mathcal{F}_k$ is general, the fact that they can be approximated by finite filtrations without affecting the strategic properties of $\Gamma_{k,\tau}$ is proven in Shmaya and Solan (2004, Section 6).

We now attach a color $c_{k,\tau}$ from a finite set C to each of the periodic stopping games $\Gamma_{k,\tau}(\omega)$; the color captures the properties of the subgame perfect ξ-equilibrium in $\Gamma_{k,\tau}(\omega)$. Using a stochastic variation of Ramsey's Theorem (Shmaya and Solan (2004), Theorem 4.3), we concatenate subgame perfect

ξ-equilibria in $(\Gamma_{k,\tau}(\omega))$ to construct a δ-approximate subgame perfect ϵ-equilibrium in Γ.

The structure of the proof is similar to that of Shmaya and Solan (2004), yet some details are quiet different. Shmaya and Solan studied two-player stopping games with simultaneous stopping; they show that the two-player approximating game is equivalent to a two-player absorbing game, and they used the fact that two-player absorbing games have stationary ξ-equilibria (cf. Flesch, Thuijsman and Vrieze (1997)). In the model that we study the approximating game is equivalent to a multi-player stochastic game, yet stationary equilibria may not exist in such games, and it was not known whether the induced multi-player stochastic games have subgame perfect equilibria. Another difficulty arises with the concatenation of the subgame perfect equilibria in the finite games: when a stationary equilibrium exists, one can bound from below the probability of termination in each period of the game, and use this property to properly concatenate strategies in different approximating games. When a stationary equilibrium does not exist, such a lower bound is not available. In section 3 we provide sufficient conditions for the existence of such a lower bound, and in Section 4 we show that these conditions can be assumed w.l.o.g.

3 Periodic stopping games with a finite filtration

Periodic stopping games in which the filtration is finite play an important role in the proof of Theorem 2.6. In the present section we prove that these games have subgame perfect ϵ-equilibria.

Suppose that the filtration is finite, that is, each $\mathcal{F}_k$ has finitely many atoms, and $\mathcal{F}_1$ has a single atom. Given a bounded stopping time τ, consider the game that restarts at time τ: if no player stops before time τ, then at time τ the game restarts, with a new state $\omega \in \Omega$ that is chosen according to P, independently of previous choices of the state. Such a game can be represented as a stochastic game in which each atom of each $\mathcal{F}_k$ correspond to a state. We call such a game "a stopping game on a finite tree".

Definition 3.1 *A* **stopping game on a finite tree** *is given by* $T =$

$\left(S, S_0, r, (C_s, p_s, i_s, a_s)_{s \in S_0}, a_\infty\right)$, *where*

- $\left(S, S_0, r, (C_s)_{s \in S_0}\right)$ *is a tree. S is a nonempty finite set of nodes, $S_0 \subset S$ is the set of nodes which are not leaves, $S_1 := S \setminus S_0$ is a nonempty set of leaves, $r \in S$ is the root, and for each $s \in S_0$, $C_s \subseteq S \setminus \{r\}$ is the nonempty set of children of s.*

- *For every $s \in S_0$, p_s is a probability distribution over C_s.*

- *For every $s \in S_0$, $i_s \in I$ is the player who can terminate the game at node s, and $a_s \in R$ is the terminal payoff at that node.*

- *$a_\infty \in \Re^n$ is a payoff vector.*[4]

A stopping game on a finite tree starts at the root r. If the current node is $s \in S_0$ and the game did not terminate before, player i_s decides whether to stop the game or to continue. If he stops, the game terminates and the terminal payoff vector is a_s. Otherwise, the game continues, and a new node $s' \in S$ is chosen according to p_s. If $s' \in S_0$, the process repeats itself with s' as the current node; if $s' \in S_1$, the process repeats itself with r as the current node. If the game never terminates, the payoff vector is a_∞.

A stopping game on a finite tree is a stochastic game with perfect information. A (behavior) strategy σ^i for player i is a function from the set of all finite histories that end at a decision node of player i, to $[0, 1]$. A strategy σ^i is **stationary** if the play after every given history is a function of the last decision node in the history.

Denote by H the space of all finite histories, and by H_∞ the set of all infinite histories. H_∞, equipped with the σ-algebra spanned by the cylinder sets, is a measurable space.

Let θ be the termination node, that is, the first node s in which player i_s chooses to stop. In case the game never terminates we set $\theta = +\infty$. Thus a_θ is the payoff vector in the game.

[4] For simplicity, in this model, unlike the general one, we assume that a_∞ is a payoff vector, rather than an integrable payoff function.

Each profile σ together with the distributions $(p_s)_{s\in S_0}$, induces a distribution $\mathbf{P}_\sigma$ over H_∞. Let E_σ be the expectation operator that corresponds to $\mathbf{P}_\sigma$. The expected payoff vector under σ is

$$\gamma(\sigma) := E_\sigma[a_\theta].$$

For every finite history $h \in H$, we denote by $T_{|h}$ the restriction of T to the subgame that start after history h occurs. Given a finite history $h = (s_1, b_1, s_2, b_2, ..., s_L) \in H$ (such that s_l is a node and b_l is a chosen action of player i_{s_l}), and strategy σ^i of player $i \in I$, we define the continuation strategy of player i given the history h occurs by

$$\sigma^i_{|h}(h') := \sigma^i(h, h'),$$

for every $h' = (s'_1, b'_1, s'_2, b'_2, ..., s'_M) \in H$ such that $s_L = s'_1$, $i_{s'_M} = i$, and $(s_1, b_1, ..., s_L, b'_1, s'_2, b'_2, ..., s'_M) \in H$. Note that $\sigma^i_{|h}$ is a strategy in $T_{|h}$.

Let $\gamma_{|h}(\sigma)$ be the expected payoff vector that corresponds to the profile $\sigma_{|h}$ in the subgame $T_{|h}$.

Definition 3.2 *Let $\epsilon \geq 0$. A profile σ is a* **subgame perfect ϵ-equilibrium** *if for every history $h \in H$, the profile $\sigma_{|h}$ is an ϵ-equilibrium in the subgame $T_{|h}$.*

Assuming no player ever stops, the collection $(p_s)_{s\in S_0}$ of probability distributions at the nodes induces a probability distribution over the set S_1 of leaves, and over the set of branches that connect the root to the leaves. For every set $E \subseteq S$, denote by π_E the probability that the chosen branch passes through E.

Let $S_{0,i} := \{s \in S_0 | i_s = i\}$ be the set of decision nodes of player i. In general, given a set of nodes E, we indicate by $E_i := E \cap S_{0,i}$ the subset of E which contains the decision nodes of player i in E.

Let $Z_i := \{s \in S_{0,i} \mid a^i_s = 0\}$ be the set of decision nodes of player i in which his terminal payoff is zero. Below we will assume w.l.o.g. that the maximal payoff that a player can achieve when all other players continue is 0, so that Z_i will be the set of all states in which player i can stop and obtain his optimal payoff.

The next theorem states that if π_{Z_i}, the probability that the game passes through Z_i, is close to 1, for every $i \in I$, and if every player can gain at most 0 by terminating the game by himself, then the game has a subgame perfect 2ϵ-equilibrium, for every $\epsilon > 0$. Furthermore, unless the equilibrium is stationary, there is a lower bound on the termination probability, even in case one player deviates, and this bound does not depend on the filtration or on the depth of the finite tree.

THEOREM 3.3 *Let $D \in \mathbf{N}$, $D > 2$, and let $\epsilon \in (0, \min\{\frac{1}{2D}, \frac{1}{n}\})$. Let $T = \left(S, S_0, r, (C_s, p_s, i_s, a_s)_{s \in S_0}, a_\infty\right)$ be a stopping game on a finite tree that satisfies the following conditions :*

Q.1 *for every $s \in S_0$, $a_s \in R := \left\{0, \pm\frac{1}{D}, \pm\frac{2}{D}, ..., \pm\frac{D}{D}\right\}^n$;*

Q.2 $\max_{s \in S_{0,i}} a_s^i = 0$, *for every* $i \in I$;

Q.3 $\pi_{Z_i} > 1 - \frac{\epsilon^3}{32}$, *for every $i \in I$.*

Then the game has a subgame perfect 2ϵ-equilibrium σ_*. Furthermore, unless the equilibrium is stationary, there is an integer $B = B(\epsilon, n)$ such that under σ_* the game terminates during every $3B$ periods with probability at least $\frac{\epsilon^6}{72B}$ even if one of the players deviates.

The proof of Theorem 3.3

The proof distinguishes between three cases. We start with identifying two cases where a stationary equilibrium exists (Section 3.1). We then study the periodic game excluding these cases (Section 3.2).

3.1 Stationary equilibria

We first discuss the case in which the vector a_∞ is non-negative. By Condition Q.2, the highest payoff that a player can obtain by stopping is 0. Therefore we obtain the following:

LEMMA 3.4 *If $a_\infty^i \geq 0$ for every $i \in I$, then $(0^i)_{i \in I}$ is a stationary 0-equilibrium.*

From now on we assume the following:
A.1. There is a player $j \in I$ such that $a^j_\infty < 0$.

For the second type of equilibrium we need some definitions.

Definition 3.5 *A player $i \in I$ is called* **dummy** *if $a^i_\infty \geq 0$, and $a^i_s \geq 0$ for every $s \in S_0$ such that $i_s \neq i$.*

A dummy player has no reason to terminate the game, since whatever happens his payoff is at least 0, his maximal payoff if he stops. Consider the game $\overline{T}$ in which all dummy players were recursively eliminated. Every subgame perfect ϵ-equilibrium in $\overline{T}$ can be extended to a subgame perfect ϵ-equilibrium in T, by instructing all dummy players to continue whenever chosen. Therefore, we assume the following w.l.o.g. :
A.2. There are no dummy players in T.

Assume that the profile σ is a stationary profile. Since σ is stationary, σ is a subgame perfect (ϵ-) equilibrium in T, if and only if for every $s \in S_0$, σ induces an (ϵ-) equilibrium in the subgame which starts at stage s. We denote by $\gamma_{|s}(\sigma)$ the expected payoff when the initial node is s rather than r when the players follow a stationary strategy σ.

Assume next that there is a player $i \in I$ who has a stationary strategy σ^i such that all the players (excluding player i) prefer that player i follows σ^i while all the other players continue whenever chosen, rather than they stop the game by themselves. Assume also that by following this strategy player i obtains 0, in every subgame (which is the maximal payoff he can receives by stopping). In this case we say that player i is a social welfare player. We show that under Assumptions A.1 and A.2, if player i is a social welfare player, then the game has a subgame perfect (ϵ-) equilibrium, in which player i follows σ_i, and all the other players (except, perhaps, for one) continue whenever chosen (see Lemma 3.7).

Definition 3.6 *A* **social welfare player** *is a non-dummy player $i \in I$ who has a pure stationary strategy σ^i of player i, such that*

$$\gamma^i_{|s'}\left(0^{-i}, \sigma^i\right) = 0 \quad \forall s' \in S_0 \quad \textit{and,} \tag{1}$$

$$\gamma^j_{|s'}\left(0^{-i}, \sigma^i\right) \geq \gamma^j_{|s'}\left(0^{-i,-j}, \sigma^i, \sigma^j\right) \quad \forall j \neq i,\ \sigma^j \neq 0^j \textit{ and } s' \in S_0, \tag{2}$$

where $0^{-i,-j}$ is the profile in which all the players, excluding players i and j, continue whenever they are chosen.[5]

Consider Definition 3.6. By Assumption A.1 and Condition Q.2, and since σ^i is stationary, such σ^i is necessarily terminating. Furthermore, by Eq. (1) player i can gain 0 by following σ^i, and therefore under the strategy σ^i, player i stops only at nodes $s \in Z_i$, that is, at nodes $s \in S_{0,i}$ where $a^i_s = 0$.

By Eq. (2), if player i is a social welfare player, then every player $j \neq i$ prefers that player i terminates the game according to the strategy σ^i rather than he himself does so. This leads us to next type of equilibrium:

LEMMA 3.7 *If there is a social welfare player, then there is a stationary subgame perfect ϵ-equilibrium, for every $\epsilon > 0$.*

Proof : Let $i \in I$ be a social welfare player. There are two cases.

1. $a^i_\infty < 0$: the profile "player i follows σ^i and all other players continue whenever chosen" is a stationary subgame perfect 0-equilibrium.

2. $a^i_s < 0$ for some $s \in S_0$: the profile "player i follows σ^i, player i_s stops with probability ϵ' sufficiently small, whenever the game reaches node s, and all other players continue whenever chosen" is a stationary subgame perfect ϵ-equilibrium.

Since player i is not dummy, at least one of these cases hold. ■

The result of Lemma 3.7 is tight, in the sense that the game needs not have a subgame perfect ϵ-equilibrium in pure strategies, nor a subgame perfect 0-equilibrium (cf. Solan and Vieille (2003), Example 3).

From now on we assume the following:
A.3. There are no social welfare players.

[5]The strategy σ^i of a social welfare player i could be a behavior strategy rather than pure. Later we prove that if there is no player that has a pure strategy which satisfies Eq. (1) and Eq. (2) then the game has a subgame perfect ϵ-equilibrium in pure strategies.

3.2 Non-stationary equilibria

Assumption A.1 ensures that the profile $(0^i)_{i\in I}$ (according to which all players always continue) is not a subgame perfect equilibrium. By Assumption A.3, every player $i \in I$ is not a social welfare player. Therefore, for every player $i \in I$ and every $s \in S_{0,i}$, there is a player $j \in I$ such that $a_s^j < 0$.

Under these Assumptions we prove the next theorem:

THEOREM 3.8 *If Conditions Q.1-Q.3 and Assumptions A.1-A.3 hold, then the game T has a subgame perfect 2ϵ-equilibrium in pure strategies.*

The rest of this section is devoted to the proof of Theorem 3.8. For every $v = (v_s)_{s\in S_1} \in \Re^{|S_1|\times n}$ define an auxiliary game $G(v)$, as a single round of the game T (that start at the root r), such that if no player stopped in this round, the game ends when it reaches a leaf $s \in S_1$, with final payoff v_s.

To prove Theorem 3.8, we construct a sequence of final payoffs $(v_m)_{m\in\mathbf{N}}$, and a sequence $(\mu_m)_{m\in\mathbf{N}}$ of subgame perfect ϵ-equilibria in $(G(v_m))_{m\in\mathbf{N}}$ with corresponding payoffs vectors $(u_m)_{m\in\mathbf{N}}$, in which for every $m > 1$, the final payoff in the leaves $v_m := (v_{m,s})_{s\in S_1}$ are subgame perfect ϵ-equilibrium payoffs in earlier games in the sequence (that is, for every leaf $s \in S_1$, $v_{m,s} = u_l$ for some $l < m$). We will then properly concatenate the profiles $(\mu_m)_{m\in\mathbf{N}}$ one after the other. Finally, we use a diagonal extraction argument to show that a limit of the concatenations of the profiles $(\mu_m)_{m\in\mathbf{N}}$ is a subgame perfect ϵ-equilibrium in T.

We start by explaining the main ideas of the construction of these sequences. The formal construction follows. We will simultaneously construct the three sequences by induction. Assume we already constructed the first m elements of each sequence as required.

The naive construction is to set $v_{m+1,s} = u_m$ for every leaf $s \in S_1$. That is, the final payoff in $G(v_{m+1})$ is the ϵ-equilibrium payoff in the previously constructed game. If u_m has a negative coordinate, that is, there is a player who prefers to stop rather than obtaining the final payoff u_m, then in any equilibrium in $G(v_{m+1})$ the game terminates (by the players) with positive probability. If moreover $u_m \ngeq -\overrightarrow{\frac{\epsilon^2}{2}} := \left(-\frac{\epsilon^2}{2}, ..., -\frac{\epsilon^2}{2}\right)$, that is the final payoff is below $-\frac{\epsilon^2}{2}$ for some player,[6] then we could bound from below the probability of termination in any equilibrium in $G(v_{m+1})$. If $u_m \not\geq -\overrightarrow{\frac{\epsilon^2}{2}}$ for every $m \in \mathbf{N}$, one can construct a subgame perfect equilibrium along the lines described in the previous paragraph.

Unfortunately, u_m may not be a non-positive vector, in which case all players will prefer to continue in $G(v_{m+1})$, and we will not be able to use this procedure to construct a terminating profile. Therefore, when $u_m \geq -\overrightarrow{\frac{\epsilon^2}{2}}$ our construction is more intricate. Roughly, we set $m' < m$ to be the maximal index such that $u_{m'} \not\geq -\overrightarrow{\frac{\epsilon^2}{2}}$. In $G(v_{m+1})$, for <u>some</u> nodes $s \in S_0$, we instruct the player i_s who controls s to stop at s, if he stopped at s in one of the strategies $\mu_{m'}, \mu_{m'+1}, ..., \mu_m$. If player i_s stops at s according to $\mu_{\hat{m}}$, where $m' \leq \hat{m} \leq m$, then it means that player i_s prefer to stop at s, when the play after s coincides with $\mu_{\hat{m}}$ and the final payoff in all leaves s' that can be reached from s, is $v_{\hat{m},s'}$. We then set $v_{m,s'} = v_{\hat{m},s'}$ in these leaves. We will show that in such a construction, infinitely many vectors among $(u_m)_{m \in \mathbf{N}}$ are $\not\geq -\overrightarrow{\frac{\epsilon^2}{2}}$, and therefore a subgame perfect ϵ-equilibrium can be constructed along the lines described above.

To state the formal construction we need additional preparation, which are done in Step 1. The formal construction of the sequences will be given in Steps 2-4.

<u>**Step 1: Perturbation of the payoffs in T**</u>

Assumption A.1 ensures that in any subgame perfect equilibrium the game terminates with probability 1. Assumptions A.1-A.3 ensure that if at most one player, say player i, uses a stationary strategy in T that is not 0^i, and all the other players continue, then there is some player j whose expected payoff is negative. In particular, there is at least one node $s \in S_{0,j}$ in which player j is better off by stopping than continuing. If it was possible to bound from above the negative payoff of player j uniformly (over all the trees), one could bound from below the probability that player j stops, and use this bound for constructing a subgame perfect ϵ-equilibrium in the game T. Unfortunately, such a uniform bound does not necessarily exist. In order to

[6] For every $b, a \in \Re^n$, denote $b \geq a$ if and only if $b^i \geq a^i$ for every $i \in I$; $b \not\geq a$ if and only if there is $i \in I$ such that $b^i < a^i$.

get some bound, we will perturb the payoffs in T as follows: we reduce the payoff a^i_∞ by ϵ^2 for every $i \in I$. Likewise, for every node s, we reduce the payoff of every player who does not control the node by ϵ^2. Formally, let $\widehat{T} := \left(S, S_0, r, (C_s, p_s, i_s, \hat{a}_s)_{s\in S_0}, \hat{a}_\infty\right)$ where $\hat{a}^j_\infty := a^j_\infty - \epsilon^2$ for every $j \in I$; $\hat{a}^j_s := a^j_s - \epsilon^2$ for every $s \in S_0$ and every $j \neq i_s$. Every subgame perfect ϵ-equilibrium in $\widehat{T}$ is a subgame perfect 2ϵ-equilibrium in T.

Assumptions A.1-A.3 hold for $\widehat{T}$. Hence there is a player $j \in I$ such that $\hat{a}^j_\infty < -\epsilon^2$. In addition, for every $i \in I$ and every stationary strategy $\mu^i \neq 0^i$ of player i such that i stops at a set $E \subseteq Z_i$ and otherwise he continues, there is a player $j \in I$ who loses at least ϵ^2 in $\widehat{T}$ under the profile $(\mu^i, 0^{-i})$.

We are now ready to simultaneously construct the sequences $(v_m)_{m\in\mathbf{N}}$, $(\mu_m)_{m\in\mathbf{N}}$, and $(u_m)_{m\in\mathbf{N}}$ in the perturbed game $\widehat{T}$.

Step 2: The initial value u_0

Set $u_0 := \hat{a}_\infty$. Then $u_0 \ngeq -\overrightarrow{\frac{\epsilon^2}{2}}$.

Assume we already constructed the first m elements of each sequence as required. We first discuss the case in which $u_m \ngeq -\overrightarrow{\frac{\epsilon^2}{2}}$.

Step 3: The case $u_m \ngeq -\overrightarrow{\frac{\epsilon^2}{2}}$

In this case there is a nonempty set of players I', such that $u^i_m < -\frac{\epsilon^2}{2}$ for every player $i \in I'$. We set $v_{m+1,s} := u_m$ for every $s \in S_1$. We also set $g_{m+1,s} := m$, for every $s \in S_1$. This is the index of the game whose ϵ-equilibrium payoff was determined as $v_{m+1,s}$.

By using an (approximate) backward induction process, we will construct a specific ϵ-equilibrium in $G(v_{m+1})$. Let μ_{m+1} be the profile obtained by using a backward induction in which at every node $s \in S_0$:

- If $i_s \in I'$, player i_s stops at s if and only if his payoff if he stops is at least as much as his expected payoff if he continues.
- If $i_s \notin I'$, player i_s stops at s if and only if his payoff if he stops exceeds his expected payoff if he continues by at least ϵ.

We say that players $i \in I'$ use the regular rule, and players $i \notin I'$ use the ϵ-rule

This backward induction process produces a profile μ_{m+1}, in which, in every subgame, no player can gain more than ϵ by deviating. That is, μ_{m+1} is a subgame perfect ϵ-equilibrium in $G(v_{m+1})$.

Note that, since the players in I' use the regular rule, and since (by Condition Q.2) every player has nodes in which he is allowed to stop and receive zero, the game will terminate with positive probability. Furthermore, all the players who are not in I' use the ϵ-rule, so they delay their stopping even if they lose, thereby increasing the probability that players in I' stop. As a result, and by Condition Q.3, the game $G(v_{m+1})$ terminates with some positive probability which is bounded from below (see Corollary 3.10). Moreover the probability that either the game $G(v_{m+1})$ or the game $G(v_{m+2})$ terminate is bounded from below, even if one of the players deviates (see Corollary 3.13).

We now formally state and prove some properties of the subgame perfect ϵ-equilibrium μ_{m+1} in case $u_m \ngeq -\overrightarrow{\frac{\epsilon^2}{2}}$.

Denote by $\phi(\mu_{m+1})$ the probability that under the profile μ_{m+1} the game $G(v_{m+1})$ terminates, and by $\phi_i(\mu_{m+1})$ the probability that $G(v_{m+1})$ is terminated by player i.

The following lemma presents a lower bound on the probability $\phi_i(\mu_{m+1})$, that the game $G(v_{m+1})$ is terminated by a player $i \in I'$. The lower bound depends on the probability $\phi_{-i}(\mu_{m+1})$ that the game is terminated by the other players. Furthermore, this bound is a decreasing function of $\phi_{-i}(\mu_{m+1})$.

LEMMA 3.9 *For every player $i \in I'$ (that is, $u_m^i < -\frac{\epsilon^2}{2}$), the probability that the game $G(v_{m+1})$ will be terminated by player i, satisfies*

$$\phi_i(\mu_{m+1}) \geq 1 - \frac{\epsilon^3}{32} - \frac{\phi_{-i}(\mu_{m+1})\left(1 + \frac{\epsilon^2}{2}\right)}{\frac{\epsilon^2}{2}}.$$

Proof Assume all the players follow μ_{m+1}. Recall that the set Z_i is the set of decision nodes of player i in which his terminal payoff is zero. For every node $s \in Z_i$, the game $G(v_{m+1})$ does not terminate in s with positive probability

only if one of the following two cases holds:
(i) the game is terminated before it reaches s by some player, which happens with probability at least $\pi_{\{s\}}$;[7]
(ii) the game does not terminate before it reaches s, and player i (who uses the regular rule) prefers to continue, since his continuation payoff at s is not negative.
Denote by r_s and q_s the probability that the subgame that starts at s is terminated by player i, and by all of the players except i, respectively. The subgame that starts at s is terminated either by player i, in which case i receives at most 0, or by the other players, in which case he receives at most 1, or it might ends at the leaves, in which case player i receives at most $u^i_m < -\frac{\epsilon^2}{2}$. As a result, the continuation payoff for player i at s is at most

$$r_s \cdot 0 + q_s \cdot 1 + (1 - r_s - q_s) \cdot (-\frac{\epsilon^2}{2}).$$

Hence, a necessary condition for player i to prefer continuing is:

$$q_s \cdot 1 + (1 - r_s - q_s) \cdot (-\frac{\epsilon^2}{2}) > 0,$$

which is equivalent to:

$$r_s > 1 - q_s \frac{1 + \frac{\epsilon^2}{2}}{\frac{\epsilon^2}{2}}. \tag{3}$$

In particular, if all players except i continue (that is, if $q_s = 0$), then player i stops at s with probability $r_s = 1$.

By Condition Q.3, the probability that the chosen branch passes through Z_i is at least $1 - \frac{\epsilon^3}{32}$. For every $s \in S$, denote by $Succ\,(s)$ the set of all descendants of s in the finite tree $\left(S, S_0, r, (C_s)_{s \in S_0}\right)$. Let $Z_i^F := \{s \in Z_i | \ \forall s' \in Z_i, \quad s \notin Succ(s')\}$ be the **frontier** of Z_i, that is the set of all the nodes in $s' \in Z_i$ such that there is no node in Z_i that appears before s'. We divide the set Z_i^F into three disjoint sets X, Y and W: X contains the nodes in which the game is terminated by one of the players except i before it reaches s (i.e., case (i) holds); Y contains the nodes in which the game does not terminate before it reaches s, but player i prefers to continue (i.e., case (ii) holds); W contains

[7]Recall that, for every set $E \subseteq S$, π_E is the probability that the chosen branch passes through E, given all the players continue whenever chosen.

the nodes $s \in Z_i^F$ in which the game is terminated by player i, either at the s, or at some node which appears before s.

If the chosen branch passes through X (with probability π_X), then the game is necessarily terminated by a player other than i; if the chosen branch passes through Y (with probability π_Y), then by Eq. (3), the conditional probability that the game is terminated by player i is higher than $1 - \frac{\sum_{s\in Y} \pi_{\{s\}} q_s}{\pi_Y} \frac{1+\frac{\epsilon^2}{2}}{\frac{\epsilon^2}{2}}$; if the chosen branch passes through W (with probability π_W), then the game is necessarily terminated by player i. The probability that the chosen branch passes through either X, Y, or W is equal to the probability that the chosen branch passes through Z_i, so that

$$\pi_X + \pi_Y + \pi_W = \pi_{Z_i} \geq 1 - \frac{\epsilon^3}{32}.$$

To summarize the probability that the game is terminated by player i, $\phi_i(\mu_{m+1})$, satisfies:

$$\begin{aligned}
\phi_i(\mu_{m+1}) &\geq \pi_Y \cdot \Big(1 - \frac{\sum_{s\in Y} \pi_{\{s\}} q_s}{\pi_Y} \frac{1+\frac{\epsilon^2}{2}}{\frac{\epsilon^2}{2}}\Big) + \pi_W \\
&= \pi_Y + \pi_W - \sum_{s\in W} \pi_{\{s\}} q_s \frac{1+\frac{\epsilon^2}{2}}{\frac{\epsilon^2}{2}} \\
&= \pi_{Z_i} - \pi_X - \sum_{s\in W} \pi_{\{s\}} q_s \frac{1+\frac{\epsilon^2}{2}}{\frac{\epsilon^2}{2}} \\
&\geq 1 - \frac{\epsilon^3}{32} - \pi_X - \sum_{s\in W} \pi_{\{s\}} q_s \frac{1+\frac{\epsilon^2}{2}}{\frac{\epsilon^2}{2}}.
\end{aligned}$$

One can verify that the probability $\phi_{-i}(\mu_{m+1})$ that the game is terminated by one player other than i, is at least $\pi_X + \sum_{s\in W} \pi_{\{s\}} q_s$, hence

$$\phi_i(\mu_{m+1}) \geq 1 - \frac{\epsilon^3}{32} - \pi_X - \sum_{s\in W} \pi_{\{s\}} q_s \frac{1+\frac{\epsilon^2}{2}}{\frac{\epsilon^2}{2}} \geq 1 - \frac{\epsilon^3}{32} - \frac{\phi_{-i}(\mu_{m+1}) \left(1+\frac{\epsilon^2}{2}\right)}{\frac{\epsilon^2}{2}}.$$

■

From Lemma 3.9, one can derive a lower bound on the probability that $G(v_{m+1})$ terminates, which only depends on ϵ.

COROLLARY 3.10 $\phi\left(\mu_{m+1}\right) > \frac{\epsilon^2}{4}$.

The next lemmas introduce bounds on the probability that the game $G\left(v_{m+1}\right)$ terminates when one player deviates. The following lemma claims that if the deviator i is not in I' (i.e., $u_m^i \geq -\frac{\epsilon^2}{2}$), then the game terminates with probability at least $\frac{\epsilon^4}{6}$, even if player i deviates.

LEMMA 3.11 *Let $i \in I$ such that $u_m^i \geq -\frac{\epsilon^2}{2}$. Then*

$$\phi\left(\mu_{m+1}^{-i}, 0^i\right) \geq \frac{\epsilon^4}{6}. \tag{4}$$

Proof If the probability that the game is terminated by player i satisfies $\phi_i\left(\mu_{m+1}\right) < \frac{\epsilon^2}{4n}$, then Eq. (4) follows from Corollary 3.10. Assume then that the probability that the game is terminated by player i satisfies $\phi_i\left(\mu_{m+1}\right) \geq \frac{\epsilon^2}{4n}$.

Denote by X the set of decision nodes of player i in which the game $G\left(v_{m+1}\right)$ is terminated by player i. In particular, $\pi_X = \phi_i\left(\mu_{m+1}\right)$.

Fix a node $s \in X$. Since $i \notin I'$, player i uses the ϵ-rule. Since player i stops at s, his continuation payoff at s is necessarily at most -ϵ.

Let q_s be the probability that the subgame that starts at s terminates under $\left(\mu_{m+1}^{-i}, 0^i\right)$. In that case, player i's payoff is at least

$$q_s \cdot \left(-1 - \epsilon^2\right) + (1 - q_s)\left(-\frac{\epsilon^2}{2}\right).$$

Therefore, if player i follows μ_{m+1}^i in the subgame that starts at s, he definitely receives at least this amount. A necessary condition for player i to prefer stopping at s is,

$$q_s \cdot \left(-1 - \epsilon^2\right) + (1 - q_s)\left(-\frac{\epsilon^2}{2}\right) \leq -\epsilon,$$

which is equivalent to

$$q_s \geq \frac{\epsilon - \frac{\epsilon^2}{2}}{1 + \frac{\epsilon^2}{2}} \geq \frac{2}{3}\epsilon.$$

As a result, in case that player i deviates, the game terminates with probability $\phi\left(\mu_{m+1}^{-i}, 0^i\right)$, which is at least

$$\phi\left(\mu_{m+1}^{-i}, 0^i\right) \geq \sum_{s\in X} \pi_s \cdot q_s \geq \sum_{s\in X} \pi_s \cdot \frac{2}{3}\epsilon = \phi_i\left(\mu_{m+1}\right) \cdot \frac{2}{3}\epsilon. \tag{5}$$

Consequently, since $\phi_i\left(\mu_{m+1}\right) \geq \frac{\epsilon^2}{4n}$, and $\epsilon < \frac{1}{n}$, Eq. (4) follows from Eq. (5). ■

As opposed to the previous case, if the deviator is a member of I', the probability that the game $G\left(v_{m+1}\right)$ terminates when player i deviates is not necessarily bounded from below. Nevertheless, the following lemma asserts that if this probability is too small, then u_{m+1}, the ϵ-equilibrium in $G\left(v_{m+1}\right)$ satisfies $u_{m+1} \not\geq -\overrightarrow{\frac{\epsilon^2}{2}}$, and $u_{m+1}^i \geq -\frac{\epsilon^2}{2}$. As a result, player i is not a member of I' in the game $G\left(v_{m+2}\right)$. Hence, by Lemma 3.11, the probability that the game $G\left(v_{m+2}\right)$ terminates when player i deviates, is bounded from below (see Corollary 3.13).

Lemma 3.12 *Assume there is a player $i \in I$ who satisfies:*
(a) $u_m^i < -\frac{\epsilon^2}{2}$, and (b) $\phi\left(\mu_{m+1}^{-i}, 0^i\right) < \frac{\epsilon^6}{64}$. Then
(i) $\phi_i\left(\mu_{m+1}\right) \geq 1 - \frac{\epsilon^2}{32}$, (ii) $u_{m+1}^i \geq -\frac{\epsilon^2}{2}$, and (iii) there is a player $j \neq i$ such that $u_{m+1}^j < -\frac{\epsilon^2}{2}$.

Proof
By Lemma 3.9 and by Conditions (a) and (b),

$$\phi_i\left(\mu_{m+1}\right) \geq 1 - \frac{\epsilon^3}{32} - \frac{\frac{\epsilon^6}{64}\left(1 + \frac{\epsilon^2}{2}\right)}{\frac{\epsilon^2}{2}} \geq 1 - \frac{\epsilon^2}{32}, \tag{6}$$

and (i) follows.

Consider a pure strategy μ^i of player i in which he only stops whenever he is chosen at nodes $s \in Z_i$, that is, at each node s in which $\hat{a}_s^i = 0$. By Condition (b) and Condition Q.3, if player i uses the strategy μ^i, while all the other players follow μ_{m+1}^{-i}, the game is terminated by player i with probability at least $1 - \frac{\epsilon^3}{32} - \frac{\epsilon^6}{64}$, so that his expected payoff is at least $\left(1 - \frac{\epsilon^3}{32} - \frac{\epsilon^6}{64}\right) \cdot 0 + \left(\frac{\epsilon^3}{32} + \frac{\epsilon^6}{64}\right) \cdot \left(-1 - \epsilon^2\right) \geq -\frac{\epsilon^3}{16}$.

Since player i uses the regular rule, μ^i_{m+1} is a best response against μ^{-i}_{m+1} in $G(v_{m+1})$, and therefore $u^i_{m+1} \geq -\frac{\epsilon^3}{16}$, and (ii) follows.

Since $u^i_{m+1} \geq -\frac{\epsilon^3}{16}$, the probability q by which under μ^i_{m+1} player i terminates the game at a node in which he receives a negative payoff is low. Indeed, in $G(v_{m+1})$ one of the following cases occur: (1) player i terminates the game in nodes in which he receives at most $-\frac{1}{D} < -\epsilon$, with probability q; (2) player i terminates the game at nodes in which he receives 0; (3) the game is terminated by the other players with probability $\frac{\epsilon^6}{64}$ and player i receives at most 1; (4) the game ends at the leaves, in which case player i receives less than 0. Therefore, player i's payoff is at most

$$q(-\epsilon) + \frac{\epsilon^6}{64} \cdot 1 > u^i_{m+1} > -\frac{\epsilon^3}{16},$$

so that

$$q < \frac{\epsilon^2}{8}.$$

Consider the strategy $\overline{\mu}^i$ of player i in which he stops according to μ^i_{m+1}, unless his terminal payoff is negative, in which case he continues. Since player i is not a social welfare player, there is a player j such that his expected payoff is less than -ϵ^2, given player i played $\overline{\mu}^i$, the other players follow μ^{-i}_{m+1}, and the game terminates by player i. Hence, if the players (including player i) follow μ_{m+1}, then with probability at least $1 - \frac{\epsilon^2}{32} - \frac{\epsilon^2}{8} - \frac{\epsilon^6}{64}$, the game is terminated by player i as if he follows $\overline{\mu}^i$, and player j receives less than -ϵ^2; otherwise player j receives at most $1-\epsilon^2$. Therefore, the ϵ-equilibrium payoff of player j satisfies

$$\begin{aligned} u^j_{m+1} &< \left(1 - \frac{\epsilon^2}{32} - \frac{\epsilon^2}{8} - \frac{\epsilon^6}{64}\right)(-\epsilon^2) + \left(\frac{\epsilon^2}{32} + \frac{\epsilon^2}{8} + \frac{\epsilon^6}{64}\right)(1-\epsilon^2) \\ &= -\epsilon^2 + \frac{\epsilon^2}{32} + \frac{\epsilon^2}{8} + \frac{\epsilon^6}{64} < -\frac{\epsilon^2}{2}, \end{aligned}$$

and (iii) follows. ■

Summarizing the last two lemmas, if the probability that the game $G(v_{m+1})$ terminates when player i deviates is less than $\frac{\epsilon^6}{64}$, then the probability that the game $G(v_{m+2})$ is terminated when player i deviates, is bounded from below by $\frac{\epsilon^6}{64}$.

Corollary 3.13 *For every $i \in I$, if $\phi\left(\mu_{m+1}^{-i}, 0^i\right) < \frac{\epsilon^6}{64}$ then $\phi\left(\mu_{m+2}^{-i}, 0^i\right) \geq \frac{\epsilon^6}{64}$.*

Proof By Lemma 3.11, if $\phi\left(\mu_{m+1}^{-i}, 0^i\right) < \frac{\epsilon^6}{64}$, necessarily $u_m^i < -\frac{\epsilon^2}{2}$.
Therefore, by Lemma 3.12, $u_{m+1}^i \geq -\frac{\epsilon^2}{2}$ and $u_{m+1} \ngeq -\overrightarrow{\frac{\epsilon^2}{2}}$.
Using again Lemma 3.11, it follows that, $\phi\left(\mu_{m+2}^{-i}, 0^i\right) \geq \frac{\epsilon^4}{6} \geq \frac{\epsilon^6}{64}$. ■

Step 4: The case $u_m \geq -\overrightarrow{\frac{\epsilon^2}{2}}$

As mentioned before, if we set u_m to be the final payoff in every leaf in $G\left(v_{m+1}\right)$, then we might not be able to use the iterative process to construct a subgame perfect ϵ-equilibrium. To overcome this difficulty, we choose the final payoffs $\left(v_{m+1,s}\right)_{s \in S_1}$ from the set of the previous equilibrium payoffs, such that the following conditions hold: (C1) The game $G\left(v_{m+1}\right)$ will terminate with positive probability that is bounded from below; (C2) there will be an integer B, such that for every $3B$ games $G\left(v_{m+1}\right), G\left(v_{m+2}\right), ..., G\left(v_{m+3B}\right)$, and every player $i \in I$, one of these games terminates with positive probability which is bounded from below, even if player i deviates; and (C3) there will be infinitely many games in the sequence such that the equilibrium payoffs $\ngeq -\overrightarrow{\frac{\epsilon^2}{2}}$.

Let $m' < m$ be the maximal index such that $u_{m'} \ngeq -\overrightarrow{\frac{\epsilon^2}{2}}$. m' is well defined since $u_0 = \hat{a}_\infty \ngeq -\overrightarrow{\frac{\epsilon^2}{2}}$. By Corollary 3.13, the game $G\left(v_{m'+1}\right)$ terminates with some positive probability, which is bounded from below, given that all players, except possibly one, follow $\mu_{m'+1}$.

Assume we already defined v_{m+1}. Consider the following profile in the game $G\left(v_{m+1}\right)$: at every node $s \in S_1$, player i_s stops, if and only if he stops at s according to one of the profiles $\mu_{m'+1}, \mu_{m'+2} ..., \mu_m$. Denote this profile by μ. Let $E_{m+1} := \{s \in S_0 | \mu^{i_s}(s) = 1\}$ be the set of nodes in which the players stop according to μ. Denote, $E_{m+1}^F := \{s \in E_{m+1} \mid \forall s' \in E_{m+1}, \;\; s \notin Succ\left(s'\right)\}$ the frontier of E_{m+1}. This is the set of the "highest" nodes in E_{m+1}-there is no node in E_{m+1} which appears before them in the tree. In particular, if the players stop only in nodes which are in the frontier E_{m+1}^F, rather than in every node in E_{m+1}, the probability that the game terminates, does not change: $\pi_{E_{m+1}} = \pi_{E_{m+1}^F}$.

By Corollary 3.10, $\phi(\mu_{m'+1}) > \frac{\epsilon^2}{4}$, and therefore under μ the game $G(v_{m+1})$ terminates with probability $\pi_{E_{m+1}} > \frac{\epsilon^2}{4}$. In the construction below we will use this property.

Although, the profile μ is not necessarily an ϵ-equilibrium in the game $G(v_{m+1})$, we can define a game $G(v_{m+1})$ for which there is a subgame perfect ϵ-equilibrium, such that the players stop at the frontier E^F_{m+1}. Let s be a node in E^F_{m+1}. We can give player i_s an incentive to stop at s, by choosing appropriate payoffs in the leaves of the subgame that start at s, as follows. We choose some $m_s \in \{m', m'+1, ..., m\}$ such that player i_s stops at s according to $\mu^{i_s}_{m_s}$, i.e., $\mu^{i_s}_{m_s}(s) = 1$ (if there are more than one such m_s, choose one arbitrarily). We then set the payoff for every leaf s' that can be reached from s, to be equal to its payoff in the game $G(v_{m_s})$, i.e., $v_{m+1,s'} := v_{m_s,s'}$. Furthermore, during the backward induction in $G(v_{m+1})$, we instruct the players to follows μ_{m_s} in the subgame of $G(v_{m+1})$ that starts at s. In other words, the subgame of $G(v_{m+1})$ that starts at node s, coincides with the subgame $G(v_{m_s})$ that starts at node s, and the profile $\mu_{m+1|s}$ coincides with the profile $\mu_{m_s|s}$.

Unfortunately, this profile might not satisfy Condition (C3), and therefore we cannot use the iterative process to construct a subgame perfect ϵ-equilibrium. Hence, in the game $G(v_{m+1})$, we choose only a subset of nodes $A_{m+1} \subseteq E^F_{m+1}$ such that only for each node $s \in A_{m+1}$, we instruct the players to follow the profile $\mu_{m_s|s}$ in the subgame that starts at node s. In each node which is not in one of these subgames, we instruct the players to use the ϵ-rule.

We now explain how we choose the set A_{m+1}. We first divide the set E^F_{m+1} into n disjoint subsets: for every player $i \in I$, let $E^F_{m+1,i}$ be the subset of nodes $s \in E^F_{m+1}$ that are controlled by player i. Let i_0 be the index that maximizes the probability $\pi_{E^F_{m+1,i}}$ that the chosen branch passes through this set (given the players continue whenever chosen). We next divide the chosen subset E^F_{m+1,i_0} into two disjoint subsets: the first set is Z_{m+1,i_0} which includes each node $s \in E^F_{m+1,i_0}$ such that the terminal payoff $a^{i_0}_s$ to player i_0 is 0, and the second set is N_{m+1,i_0} which includes each node $s \in E^F_{m+1,i_0}$ such that the terminal payoff $a^{i_0}_s$ to player i_0 is negative. Finally, we determine either the set Z_{m+1,i_0} or the set N_{m+1,i_0} to be the set A_{m+1}, according to the following rule:

(I) If in the previous game $G(v_m)$, either the set A_m was not defined, or there was a set A_m but $A_m \neq Z_{m,i_0}$, then

- If in the current game, the probability to path through the set N_{m+1,i_0} is higher than the probability to path through the set Z_{m+1,i_0}, then we set $A_{m+1} := N_{m+1,i_0}$.
- Otherwise, we set $A_{m+1} := Z_{m+1,i_0}$.

(II) If in the previous game $G(v_m)$, there was a set A_m that satisfies $A_m = Z_{m,i_0}$ then:

- If in the current game, the probability to pass through the set N_{m+1,i_0} exceeds the probability to pass through the set Z_{m+1,i_0} by more than $\frac{\epsilon^4}{8n}$, then we set $A_{m+1} := N_{m+1,i_0}$.
- Else, we set $A_{m+1} := Z_{m+1,i_0}$.

The motivation of the choice of the set A_{m+1} is as follows. The set A_{m+1} is contained in a set E^F_{m+1,i_0}, so each node in A_{m+1} is controlled by player i_0. Assume that player i_0 is forced to stop in each node $s \in A_{m+1}$, and all the other players continue whenever chosen. Then, by Corollary 3.10, and by (I) and (II), the game terminates with probability at least $\frac{\epsilon^2-\epsilon^4}{8n}$. If, in addition, $A_{m+1} = Z_{m+1,i_0}$ then, since player i_0 is not a social welfare player, there is another player j who is damaged in case that the game is terminated by player i_0, and after bounded number of games player j will prefer to stop. By (II), this argument is also proper in case that the game is terminated by the players other than i_0 with a low probability. On the other hand, if $A_{m+1} = N_{m+1,i_0}$ then, by the properties of the expectation, after finitely many games player i_0 will prefer to stop only at nodes in which he obtains 0.

As was explained before, for every node $s \in A_{m+1}$, we choose some $m_s \in \{m', m'+1, ..., m\}$ such that $\mu^i_{m_s}(s) = 1$. For every leaf $s' \in S_1 \cap Succ(s)$-in the subgame that starts at s define $v_{m+1,s'} := v_{m_s,s'}$, and the index of the game that its ϵ-equilibrium payoff was determined as $v_{m+1,s}$, is $g_{m+1,s'} := g_{m_s,s'}$, the same index as in the game $G(v_{m_s})$. Furthermore, during the backward induction in $G(v_{m+1})$, we instruct the players to follows μ_{m_s} in the subgame of $G(v_{m+1})$ that starts at s.

For every leaf $s' \in S_1$ that cannot be reached by nodes in A_{m+1}, we set $v_{m+1}(s) = u_m$ and $g_{m+1,s'} := m$. In addition, we instruct the players to use the ϵ-rule during the backward induction in every node that cannot be reached by nodes in A_{m+1}. Let μ_{m+1} be the appropriate ϵ-equilibrium.

If all the players follow μ_{m+1}, then (since $u_m \geq -\overrightarrow{\frac{\epsilon^2}{2}}$ and the players use the ϵ-rule) every subgame that does not include nodes in A_{m+1} terminates only in the leaves. Yet, Condition (C1) holds in $G(v_{m+1})$, and the game terminates with probability at least $\frac{\epsilon^2-\epsilon^4}{8n}$ (by Corollary 3.10). It is therefore left for us to prove that Conditions (C2) and (C3) also hold. In the next lemmas we prove that the way we choose the sets (A_m) guarantees that these conditions hold. That is, although there need not exist a lower bound for the probability of termination when one player deviates, Condition (C2) holds: there is a bound B (which only depends on ϵ and n), such that if the ϵ-equilibria payoffs $u_m, u_{m+1}, ..., u_{m+3B} \geq -\overrightarrow{\frac{\epsilon^2}{2}}$, then for every player $i \in I$, there is at least one game $G(v_{\hat{m}})$, for $\hat{m} \in \{m+1, m+2, ..., m+3B\}$, such that the game is terminated by one of the players with a probability which is bounded from below, even if player i deviates (cf. Lemma 3.15). In addition, Condition (C3) holds: there are infinite many games in the sequence such that the equilibrium payoffs $\ngeq -\overrightarrow{\frac{\epsilon^2}{2}}$.(cf. Lemma 3.16).

We next formally present and prove the properties of the subgame perfect ϵ-equilibrium μ_{m+1} in case $u_m \geq -\overrightarrow{\frac{\epsilon^2}{2}}$, which were introduced in the previous paragraph.

The following lemma claims that there is a bound B that depends on ϵ and n, such that if there are $B+1$ consecutive games $G(v_{m+1}), G(v_{m+2}), ..., G(v_{m+B+1})$ that satisfy that (i) the ϵ-equilibrium payoffs in these games are high, and (ii) the players other than i stop with sufficiently small probability in these games, then in one of these games player i is forced to stop in nodes in which he obtains 0 by stopping.

A node s is called a "zero-node" if $a_s^{i_s} = 0$: the terminal payoff to the player who control the node is 0. It is called a "negative-node" if $a_s^{i_s} < 0$.

LEMMA 3.14 *There is an integer $B = B(\epsilon, n)$ such that, if*

(i) $u_m, u_{m+1}, ...u_{m+B} \geq -\overrightarrow{\frac{\epsilon^2}{2}}$; and

(ii) $\phi\left(\mu_{m+1}^{-i},0^{i}\right),\phi\left(\mu_{m+2}^{-i},0^{i}\right),...\phi\left(\mu_{m+B+1}^{-i},0^{i}\right)<\frac{\epsilon^{6}}{30n}$*; Then there is* $\hat{m}\in\{m,m+1,...,m+B\}$ *such that* $A_{\hat{m}+1}$ *includes zero nodes of player* i*.*

Proof Let B be a sufficiently large integer; below we determine how large it should be. Assume that (i) and (ii) hold. By (i) $\overrightarrow{u_{m''}\geq-\frac{\epsilon^{2}}{2}}$ for every $m''\in\{m+1,m+2,...,m+B+1\}$, and by (ii) player i is forced to stop due to the construction in all the games $G\left(v_{m+1}\right),G\left(v_{m+2}\right),...,G\left(v_{m+B+1}\right)$. In particular, every game is necessarily terminated by player i with probability at least $\frac{\epsilon^{2}-\epsilon^{4}}{8n}-\frac{\epsilon^{6}}{30n}$.

If A_{m+1} includes zero nodes, we are done. Otherwise, assume that for every $m''\in\{m,m+1,...,m+B-1\}$, $A_{m''+2}$ includes negative nodes. Therefore, the probability that player i terminates the game $G\left(v_{m''+1}\right)$ at nodes in which he receives 0, is less than the probability that player i terminates the game at nodes in which he receives a negative payoff. Consequently, if the game is terminated by player i, his expected payoff is less[8] than $-\frac{\epsilon}{2}$. Hence, if the game terminates, player i's expected payoff is at most

$$\frac{\frac{\epsilon^{2}-\epsilon^{4}}{8n}-\frac{\epsilon^{6}}{30n}}{\frac{\epsilon^{2}-\epsilon^{4}}{8n}}\cdot\left(-\frac{\epsilon}{2}\right)+\frac{\frac{\epsilon^{6}}{30n}}{\frac{\epsilon^{2}-\epsilon^{4}}{8n}}\cdot 1<-\frac{2\epsilon^{2}}{3}.$$

Therefore, the game $G\left(v_{m''+1}\right)$ terminates with probability at least $\frac{\epsilon^{2}-\epsilon^{4}}{8n}$ in which case player i receives at most $-\frac{2\epsilon^{2}}{3}$; otherwise, the game ends at a leaf with probability at most $1-\frac{\epsilon^{2}-\epsilon^{4}}{8n}$ in which case player i receives $u_{m''}^{i}>-\frac{\epsilon^{2}}{2}$, so that

$$u_{m''+1}^{i}<\frac{\epsilon^{2}-\epsilon^{4}}{8n}\cdot\left(-\frac{2\epsilon^{2}}{3}\right)+\left(1-\frac{\epsilon^{2}-\epsilon^{4}}{8n}\right)\cdot u_{m''}^{i}.$$

Consequently, and by the construction of the games, it follows that, player i's expected payoff in the game $G\left(v_{m''+1}\right)$ is at most

$$u_{m''+1}^{i}<\left(1-\frac{\epsilon^{2}-\epsilon^{4}}{8n}\right)^{m''+1-m}\cdot 1+\left(1-\left(1-\frac{\epsilon^{2}-\epsilon^{4}}{8n}\right)^{m''+1-m}\right)\cdot\left(-\frac{2\epsilon^{2}}{3}\right). \tag{7}$$

[8]Recall that, by Condition Q.1, for every $s\in S_{0}$, $a_{s}\in R=\left\{0,\pm\frac{1}{D},\pm\frac{2}{D},...,\pm\frac{D}{D}\right\}^{n}$, where $\frac{1}{D}>\epsilon$.

For a sufficiently large m'' the value of the right-hand side of Inequality (7) is less than $-\frac{\epsilon^2}{2}$. We set B to be the minimal $m''+1-m$ such that the value in the right side is less than $-\frac{\epsilon^2}{2}$, so the lemma follows. ■

Let

$$B = B(\epsilon, n) = \min\left\{ k \in \mathbf{N} \mid \left(1 - \frac{\epsilon^2 - \epsilon^4}{8n}\right)^k \cdot 1 + \left(1 - \left(1 - \frac{\epsilon^2 - \epsilon^4}{8n}\right)^k\right) \cdot \left(-\frac{2\epsilon^2}{3}\right) < -\frac{\epsilon^2}{2} \right\}. \tag{8}$$

B is an appropriate bound for the previous lemma.

The next lemma deals with the constructions in Steps 3 and 4. It asserts that, for every consecutive games $G(v_{m+1}), G(v_{m+2}), ..., G(v_{m+3B})$, and every player $i \in I$, there is at least one game $G(v_{m''})$, where $m'' \in \{m, ..., m+B\}$, in which the probability that the game is terminated, when player i deviates, is bounded from below.

Lemma 3.15 *For every $3B$ consecutive games $G(v_{m+1}), G(v_{m+2}), ..., G(v_{m+3B})$, and every player $i \in I$, there is $m'' \in \{m+1, m+2, ..., m+3B\}$ such that the probability that the game $G(v_{m''})$ terminates, when player i deviates, is at least $\phi\left(\mu^{-i}_{m''+1}, 0^i\right) \geq \frac{\epsilon^6}{64B}$.*

Proof In case that for some $m'' \in \{m, m+1, ...m+3B-2\}$, there is $u_{m''} \ngeq -\overrightarrow{\frac{\epsilon^2}{2}}$, then the argument follows Corollary 3.13.

In addition, if there are $\widetilde{m}, \overline{m} \in \{m, m+1, ..., m+3B-2\}$, in which $u_{\widetilde{m}}, u_{\overline{m}} \geq -\overrightarrow{\frac{\epsilon^2}{2}}$, and in $G(v_{\widetilde{m}+1})$ player i is forced to stop, while in $G(v_{\overline{m}+1})$ player $j \neq i$ is forced to stop, then we are done.

We next assume that, for every $m'' \in \{m+1, m+2, ..., m+3B-2\}$, $u_{m''} \geq -\overrightarrow{\frac{\epsilon^2}{2}}$, and player i is forced to stop in each game $G(v_{m''+1})$.

In particular, for every $m'' \in \{m+1, m+2, ..., m+3B-2\}$, and every player $j \neq i$,

$$\phi\left(\mu^{-j}_{m''+1}, 0^j\right) \geq \phi\left(\mu^{i}_{m''+1}, 0^{-i}\right) \geq \frac{\epsilon^2 - \epsilon^4}{8n} > \frac{\epsilon^3 - \epsilon^5}{8} > \frac{\epsilon^6}{64B}.$$

It is left to prove that $\phi\left(\mu_{m''+1}^{-i}, 0^i\right) \geq \frac{\epsilon^6}{64B}$ for some m''.

Case 1: Assume to the contrary that for every $m'' \in \{m+1, m+2, ..., m+3B-2\}$,

$$\phi\left(\mu_{m''+1}^{-i}, 0^i\right) < \frac{\epsilon^6}{64B}.$$

We first deal with the case that A_{m+1} includes negative nodes. By Lemma 3.14 there is a minimal index $\hat{m} \in \{m, m+1, ..., m+B\}$ such that $A_{\hat{m}+1}$ includes zero nodes of player i. Furthermore, the probability that the chosen branch passes the zero nodes in $E_{\hat{m}+1}$ is at least the probability that the chosen branch passes the negative nodes in $E_{\hat{m}+1}$.

We will now estimate the probability that the game $G\left(v_{\hat{m}+1}\right)$ is terminated by player i in nodes which do not belong to $A_{\hat{m}+1}$.

Let X be the set of nodes in which the game $G\left(v_{\hat{m}+1}\right)$ is terminated by player i, excluding nodes in $A_{\hat{m}+1}$. By the construction, for every $s \in X$, the node s cannot be reached by nodes in $A_{\hat{m}+1}$, and player i uses the ϵ-rule at s. Denote by q_s the probability that the subgame of $G\left(v_{\hat{m}+1}\right)$ that starts at s is terminated by the other players. Then, if player i continues in every nodes in the subgame that starts at s, except for nodes in A_{m+1}, his payoff is at least $q_s \cdot (-1-\epsilon^2) + (1-q_s) \cdot \left(-\frac{\epsilon^2}{2}\right)$. Hence, by following $\mu_{\hat{m}+1}^i$ his payoff is at least this amount. Yet, he prefers to stop at s, so that

$$q_s \cdot \left(-1-\epsilon^2\right) + (1-q_s) \cdot \left(-\frac{\epsilon^2}{2}\right) \leq -\epsilon,$$

and therefore

$$\frac{2\epsilon}{3} \leq q_s.$$

Thus, the probability π_X, that the game is terminated by player i at nodes which do not belong to $A_{\hat{m}+1}$ satisfies,

$$\pi_X = \sum_{s \in X} \pi_s = \frac{3}{2\epsilon} \cdot \sum_{s \in X} \frac{2\epsilon}{3} \pi_s \leq \frac{3}{2\epsilon} \cdot \sum_{s \in X} q_s \pi_s \leq \frac{3}{2\epsilon} \cdot \phi\left(\mu_{m+1}^{-i}, 0^i\right) < \frac{3\epsilon^5}{2 \cdot 64B}.$$

As a result of the upper bound on π_X, the set $A_{\hat{m}+2}$ also includes zero nodes, since the probability that the chosen branch passes the zero nodes in $E_{\hat{m}+1}$

might decrease by at the most π_X, while the probability that the chosen branch passes through negative nodes in $E_{\hat{m}+1}$ increase by at the most π_X. Hence, the probability that the chosen branch passes through negative nodes in $E_{\hat{m}+2}$ may exceeds the probability that the chosen branch passes through zero nodes but by no more than

$$\frac{2 \cdot 3\epsilon^5}{2 \cdot 64B} + \frac{\epsilon^6}{\cdot 64B} < \frac{\epsilon^4}{8n}.$$

By repeating this argument for the following B games in the sequence, it follows that for every $m'' \in \{\hat{m}+1, \hat{m}+2, ..., \hat{m}+B\}$, $A_{m''}$ includes zero nodes.

Let $H = \bigcap_{m''=\hat{m}+1}^{\hat{m}+B} A_{m''}$, be the set of nodes in which player i is forced to stop in each game $G(v_{\hat{m}+1}), G(v_{\hat{m}+2}), ..., G(v_{\hat{m}+3B})$. By the upper bound of π_X it follows that

$$\pi_H \geq \frac{\epsilon^2 - \epsilon^4}{8n} - B \cdot \left(\frac{3\epsilon^5}{2 \cdot 64B} + \frac{\epsilon^6}{64B} \right) = \frac{\epsilon^2 - \epsilon^4}{8n} - \frac{\epsilon^5}{64} \cdot \left(\frac{3}{2} + \epsilon \right).$$

Player i is not a social welfare player, so there is a player $j \neq i$ whose expected payoff given the game terminates at H is less than -ϵ^2.

Hence, for every $m'' \in \{\hat{m}, ..., \hat{m}+B\}$, the expected payoff of j in $G(v_{m''+1})$ under $\mu_{m''+1}$, given the game terminates, is at most

$$\frac{\frac{\epsilon^2-\epsilon^4}{8n} - \frac{\epsilon^5}{64} \cdot \left(\frac{3}{2} + \epsilon\right)}{\frac{\epsilon^2-\epsilon^4}{8n}} \cdot \left(-\epsilon^2\right) + \frac{\frac{\epsilon^5}{64} \cdot \left(\frac{3}{2} + \epsilon\right)}{\frac{\epsilon^2-\epsilon^4}{8n}} \cdot 1 \leq -\frac{2\epsilon^2}{3}.$$

Therefore, the expected payoff of j in $G(v_{\hat{m}+B})$ satisfies

$$u^j_{\hat{m}+B} \leq \left(1 - \frac{\epsilon^2 - \epsilon^4}{8n}\right)^B \cdot 1 + \left(1 - \left(1 - \frac{\epsilon^2 - \epsilon^4}{8n}\right)^B\right) \cdot \left(-\frac{2\epsilon^2}{3}\right) < -\frac{\epsilon^2}{2},$$

a contradiction.

Case 2: Assume now that A_{m+1} includes zero nodes. If for every $m'' \in \{m+1, m+2, ..., m+B\}$ the set $A_{m''}$ includes zero nodes, the proof is similar to the previous case. If, on the other hand, there is a $m'' \in \{m+1, m+$

$2, ..., m+B-1\}$ such that the set $A_{m''}$ includes negative nodes, then we return to Case 1, and we are done. ■

The following lemma asserts that situation described in Step 3 occurs infinitely often.

LEMMA 3.16 *There are infinitely many games such that $u_m \ngeq -\overrightarrow{\frac{\epsilon^2}{2}}$.*

Proof Assume to the contrary that $u_m \geq -\overrightarrow{\frac{\epsilon^2}{2}}$ for every $m > m''$. Following Step 4 in the construction, $(E_{m+1})_{m>m''}$ is a non decreasing sequence in the following sense: the set of all descendants of E_{m+1} is contained in the set of all descendants of E_{m+2}, for every $m > m''$. Furthermore, there is an infinite subsequence of $(E_{m+1})_{m>m''}$ which is increasing. Indeed, in each game $G(v_{m+1})$, for $m > m''$, there is at least one player whose expected payoff, given the game is terminated at the set A_{m+1}, is less than $-\epsilon^2$. Therefore, if $E_{m+1} = E_{m+2} = ...$, then $A_{m+1} = A_{m+2} = ...$, thus there is necessarily a game $G(v_{\overline{m}})$, where $\overline{m} > m+1$, such that $u_{\overline{m}} \ngeq -\overrightarrow{\frac{\epsilon^2}{2}}$, a contradiction.

Since there is an infinite increasing subsequence of $(E_{m+1})_{m>m''}$, and since the set of the nodes in each game is a finite set, there is necessarily a game $G(v_{\overline{m}})$ such that $E_{\overline{m}+1}$ contains only the root r, so that $u_{\overline{m}} \ngeq -\overrightarrow{\frac{\epsilon^2}{2}}$, a contradiction. ■

Concatenating the games and the profiles:
We are going to define a new sequence of finite games $(G_m)_{m\in\mathbf{N}}$, by properly concatenating the games in $(G(v_m))_{m\in\mathbf{N}}$. Let $G_1 := G(v_1)$. Assume we already define $(G_m)_{m=1}^{m''}$. Let $G_{m''+1}$ be the game that starts with one round of $\widehat{T}$ and, if no player terminates, and the game reaches a leaf $s \in S_1$, continues[9] with $G_{g_{m''+1,s}}$. In the same way we can concatenate the profiles $(\mu_m)_{m\in\mathbf{N}}$, and get a sequence of profiles $(\sigma_m)_{m\in\mathbf{N}}$ that are subgame perfect ϵ-equilibria in $(G_m)_{m\in\mathbf{N}}$.

By Lemma 3.16 it follows that, as m goes to infinity, the minimal length of the path that connects the root to a leaf in G_m goes to infinity.

[9] Recall that, for every $s \in S_1$ and every $m \in \mathbf{N}$, $g_{m+1,s}$ is the index of the game that its ϵ-equilibrium payoff was determined as $v_{m+1,s}$.

By Lemma 3.15, for every player $i \in I$ the probability that the game G_m terminates under $(\sigma_m^{-i}, 0^i)$ before it reach a leaf goes to 1, as m goes to infinity. Note that, every game G_m is equivalent to a finite partial game of $\widehat{T}$ which start at the root.

The profile σ_m is an element in the space $[0,1]^{n_m}$ for some $n_m \in \mathbf{N}$. σ_m can be identified with a point in $[0,1]^{\mathbf{N}}$, by identifying the n_m initial coordinates with σ_m and the other coordinates with 0. The space $[0,1]^{\mathbf{N}}$ is compact, and therefore $(\sigma_m)_{m\in\mathbf{N}}$ has a subsequent that converges to a limit σ.

Using a limiting argument, it is standard to prove that σ is a subgame perfect ϵ-equilibrium (cf. Solan and Vieille (2003)).

From the proof it follows that the bound $B = B(\epsilon, n)$ defined in Eq. (8) satisfies that according to σ the probability that the game terminates at every $3B$ periods in the period stopping game is at least $\frac{\epsilon^6}{64B}$, even if one of the players deviates. ■

Remark 3.17 *The proof of an existence of a subgame perfect 2ϵ-equilibrium in Theorem 3.8 can be easily generalized to the case of a periodic stopping game with bounded payoffs, without Conditions Q.1-Q.3. However, in the absence of Condition Q.3, the upper bound B, as well as the lower bound of termination probability under deviation, depend on the minimal probability that the chosen branch reaches a leaf s, over the set of leaves.*

4 The proof of Theorem 2.6

4.1 Preliminaries

Let $\Gamma = (I, \Omega, \mathcal{A}, \mathbf{P}, \mathcal{F}, (i_k)_{k=1}^{\infty}, (a_k)_{k=1}^{\infty}, a_\infty)$ be a stopping game such that $\|a_\infty\|_\infty \in L^1(\mathbf{P})$, and $\sup_{k\in\mathbf{N}} \|a_k\|_\infty \in L^1(\mathbf{P})$. Fix $\delta, \epsilon > 0$ once and for all. Since every δ-approximate subgame perfect ϵ-equilibrium is δ'-approximate subgame perfect ϵ'-equilibrium, for $\delta < \delta'$ and $\epsilon < \epsilon'$, we can assume that $0 < \delta < \epsilon < \frac{1}{4}$.

Since the payoffs satisfy $\|a_\infty\|_\infty \in L^1(\mathbf{P})$, and $\sup_{k\in\mathbf{N}} \|a_k\|_\infty \in L^1(\mathbf{P})$, and since we want to prove the existence of a δ-approximate subgame perfect ϵ-equilibrium, we can assume w.l.o.g. that the absolute values of the payoffs are bounded by 1, and moreover the range of the payoff processes is included in the finite set $R = \left\{0, \pm\frac{1}{D}, \pm\frac{2}{D}, ..., \pm\frac{D}{D}\right\}^n$, where $\epsilon^2 < \frac{1}{D} < \epsilon$ (cf. Shmaya and Solan (2004)).

Fix a value $\xi \in \left(0, \frac{\epsilon^2}{2}\right)$. In the rest of this subsection we prove that there is no loss of generality assuming that:
A.4. There is a set of players $J \subseteq I$, and an increasing sequence $1 \leq K_1 < K_2 < ...$ of integers such that for every $m \in \mathbf{N}$ the following condition holds, with probability at least $1 - \frac{\xi^3\delta}{64}$:
(1) for every player $j \in J$, there is a stage $k \in \{K_m, K_m + 1, ..., K_{m+1} - 1\}$ such that player j receives 0 by terminating the game by himself at stage k, and
(2) the maximal payoff every player $j \in J$ can gain by terminating the game by himself during stages K_m up to K_{m+1} is 0, and
(3) every player i who does not belong to J, cannot terminates the game during stages K_m up to K_{m+1}.

The motivation for this assumption is as follows. In the proof of Theorem 2.6 we will use a reduction to periodic games $\Gamma_{K_m\tau}$, which are played between stages K_m and a bounded stopping time $\tau \geq K_{m+1}$. In the next section it will be argued that these games can be approximated by games with a finite filtration, that is, a collection of games on tree. Assumption A.4 guarantees that with probability $1 - \frac{\delta}{2}$ the approximating games satisfy Conditions Q.1-Q.3 of Theorem 3.3. Therefore these approximating games have subgame perfect ϵ-equilibria with some useful properties.

To see that assumption $A.4$ can be assumed w.l.o.g, let $\widetilde{\mathcal{F}}$ be a finite filtration of Ω that satisfies:

1. a_∞ is measurable with respect to $\widetilde{\mathcal{F}}$.
2. For every $F \in \widetilde{\mathcal{F}}$, and every $(i, a) \in I \times R$ either $(i_k(\omega), a_k(\omega)) = (i, a)$ infinitely often, for every $\omega \in F$, or (b) $(i_k(\omega), a_k(\omega)) = (i, a)$ only finitely many times, for every $\omega \in F$.

Each $F \in \widetilde{\mathcal{F}}$ is measurable, and therefore there is K sufficiently large such that $\widetilde{\mathcal{F}}$ can be approximated by sets in $\mathcal{F}_K$: there is a finite partition $\widehat{\mathcal{F}}$ which is $\mathcal{F}_K$-measurable such that for every $F \in \widetilde{\mathcal{F}}$ there is $\widehat{F} \in \widehat{\mathcal{F}}$ that satisfies

$$P(F \triangledown \widehat{F}) = P(F \setminus \widehat{F}) + P(\widehat{F} \setminus F) < \frac{\xi^3 \delta}{128(n+1)|R|}.$$

Likewise, by replacing a_∞ by a constant we add an additional noise of $\frac{\xi^3\delta}{128(n+1)|R|}$. Therefore, Theorem 2.6 holds in the subgame that starts at stage K. By a backward induction we proof that Theorem 2.6 is also holds in Γ.

4.2 Approximating partial games

We will use a stochastic variation of Ramsey's Theorem due to Shmaya and Solan (2004). We now introduce the concepts which are needed for the formulation of this theorem and for using it.

DEFINITION 4.1 *Let $(\Omega, \mathbf{P}, \mathcal{F})$ be a measurable space, and let $(\mathcal{F}_k)_k$ be a filtration. Let Y be any space. A function $f : (k, \tau) \mapsto Y$ that assigns to every positive integer k and every stopping time $\tau > k$ an element in Y, is* **$\mathcal{F}$-consistent** *if for every $k \in \mathbf{N}$, every $\mathcal{F}_k$-measurable set F, and every two bounded stopping times τ_1, τ_2, we have*

$$\tau_1 = \tau_2 > k \textit{ on } F \textit{ implies } f_{k,\tau_1} = f_{k,\tau_2} \textit{ on } F.$$

That is, if $\mathbf{P}\left(F \cap \{\tau_1 = \tau_2 > k\}\right) = \mathbf{P}\left(F\right)$ then $\mathbf{P}\left(F \cap \{f_{k,\tau_1} \neq f_{k,\tau_2}\}\right) = 0$.

For every positive integer k, every stopping time $\tau > k$, and every $\omega \in \Omega$, consider the periodic stopping game $\Gamma_{k,\tau}(\omega)$ that starts at stage k, and if no player stops before stage τ, then it restart at stage k with a new state $\omega' \in \Omega$ which is chosen according to $\mathbf{P}$, such that for every $F \in \mathcal{F}_k$ either both of the state ω, ω' belong to F or not one of them belong to F.

To this end one should approximate this game by a stopping game on a tree.

Denote by $\lambda_k = \xi^6/216B^2 2^{k+2}$ for each $k \geq 0$. Set $\Lambda_k = \sum_{k' \geq k} \lambda_{k'} = \xi^6/216B^2 2^{k+1}$, so that $\sum_{k\geq 0} \Lambda_k = \xi^6/216B^2$.

DEFINITION 4.2 *Let $\tau_1 \leq \tau_2$ be two bounded stopping times. A* **λ-approximation of Γ between τ_1 and τ_2** *is a pair $\left((\mathcal{G}_k), (q_{G,k})\right)$ such that for every $k \geq 0$*

1. *$\mathcal{G}_k$ is a $\mathcal{F}_k$-measurable finite partition*[10] *of $\{\tau_1 \le k \le \tau_2\}$,*

2. *i_k and a_k are $\mathcal{G}_k$-measurable,*

3. *τ_1 and τ_2 are measurable w.r.t. $\mathcal{G}$; that is, for every $k \ge 0$, $\{\tau_1 = k\}$ and $\{\tau_2 = k\}$ are unions of atoms in $\mathcal{G}_k$,*

4. *any atom G of $\mathcal{G}_k$ such that $k < \tau_2$ on G is a union of some atoms in $\mathcal{G}_{k+1}$,*

5. *for every atom G of $\mathcal{G}_k$, $q_{G,k}$ is a probability distribution over the atoms of $\mathcal{G}_{k+1}$ that are contained in G, and*

6. *$\sum_{G' \in \mathcal{G}_{k+1}} |\mathbf{P}(G' \mid \mathcal{F}_k)(\omega) - q_{G,k}(G')| < \lambda_k$, for every atom G of $\mathcal{G}_k$ and almost every $\omega \in G$.*

THEOREM 4.3 *Let Γ be a stopping game such that $\sup_{k \in \mathbf{N}} \|a_k\|_\infty \in L^1(\mathbf{P})$. Then there is a $\mathcal{F}$-consistent function that assigns to every $k \ge 0$ and every bounded stopping time τ, a λ-approximation of Γ between k and τ.*

Theorem 4.3 was proven by Shmaya and Solan (2004) for two player stopping games with simultaneous stopping. Their proof extends to multi-player games.

Every λ-approximation $(\mathcal{G}_k, (q_{G,k}))$ of Γ between τ_1 and τ_2, defines a finite collection of games on trees, a game $T_{\tau_1,\tau_2}(G)$ for each atom $G \in \mathcal{G}_{\tau_1}$, as follows:

- The root of the tree is G.
- The nodes are all the non-empty atoms F of $(\mathcal{G}_k)$ such that (a) $F \subseteq G$, and (b) if $F \in \mathcal{G}_k$, then $\tau_2 \ge k$ on F.
- The leaves are all the atoms $F \in \cup_{k \ge \tau_1} \mathcal{G}_k$ where there is equality in (b).
- The chosen players and the terminal payoffs are given by $(i_k)_{\tau_1 \le k \le \tau_2}$ and $(a_k)_{\tau_1 \le k \le \tau_2}$.

[10] We identify $\mathcal{G}_k$ with the finite σ-algebra generated by $\mathcal{G}_k$, and we denote $\mathcal{G} = (\mathcal{G}_k)_{k \ge 0}$.

- The children of each atom F in $\mathcal{G}_k$ are all atoms F' in $\mathcal{G}_{k+1}$ which are subsets of F.
- Transition from any node F in $\mathcal{G}_k$ is given by $q_{F,k}$.
- $a_\infty \in \Re^n$ is a payoff vector.

DEFINITION 4.4 *Let $\tau_1 \leq \tau_2$ be two bounded stopping times. A $(\mathcal{F}, \tau_1, \tau_2)$-*strategy *is a sequence $\sigma = (\sigma_k)$ of random variables such that for every $k \geq 0$, (i) $\sigma_k : \{\tau_1 \leq k < \tau_2\} \to [0, 1]$, and (ii) σ_k is $\mathcal{F}_k$-measurable.*

Thus, a $(\mathcal{F}, \tau_1, \tau_2)$-strategy prescribes the player what to play in Γ between stages τ_1 and τ_2 (excluded).

Let $\sigma(G)$ be a strategy profile in $T_{\tau_1,\tau_2}(G)$. The collection $(\sigma(G))_{G\in\mathcal{G}_{\tau_1}}$ defines a $(\mathcal{F}, \tau_1, \tau_2)$-strategy profile, by instructing the players to follow the first period of each profile $\sigma(G)$ for every $G \in \mathcal{G}_{\tau_1}$. Similarly, every $(\mathcal{F}, \tau_1, \tau_2)$-strategy profile σ such that for every $k \in \mathbf{N}$, σ_k is $\mathcal{G}_k$-measurable, defines a stationary profile in $T_{\tau_1,\tau_2} := (T_{\tau_1,\tau_2}(G)))_{G\in\mathcal{G}_{\tau_1}}$. We denote by σ the collection $(\sigma(G))_{G\in\mathcal{G}_{\tau_1}}$, as well as the $(\mathcal{F}, \tau_1, \tau_2)$-strategy profile which induces by σ.

Let $F \in \mathcal{A}$ such that $\mathbf{P}(F) > 0$. Denote by $\pi(\sigma; \mathcal{F}, \tau_1, \tau_2 \mid F)$ (respectively, $\pi(\sigma; \mathcal{G}, \tau_1, \tau_2 \mid F)$) the conditional probability given F, that under σ the finite subgame of Γ which is played between stages τ_1 and τ_2 (respectively, the game T_{τ_1,τ_2}) terminates before stage τ_2. Denote by $\rho(\sigma; \mathcal{F}, \tau_1, \tau_2 \mid F)$ (respectively, $\rho(\sigma; \mathcal{G}, \tau_1, \tau_2 \mid F)$) the conditional expected payoff given F and given that this finite subgame (respectively, the game T_{τ_1,τ_2}) terminates before stage τ_2.

The following lemma provides an estimate for the difference between the conditional expected payoff and the difference between the conditional expected probability of termination, when one changes the filtration. The proof is similar to that of Lemma 6.3 in Shmaya and Solan (2004).

LEMMA 4.5 *Let $(\mathcal{G}_k, (q_{G,k}))$ be a λ-approximation of Γ between τ_1 and τ_2. For every $F \in \mathcal{A}$, such that $\mathbf{P}(F) > 0$, and every $(\mathcal{G}, \tau_1, \tau_2)$-strategy profile σ,*

1. $|\pi(\sigma; \mathcal{F}, \tau_1, \tau_2 \mid F) - (\pi(\sigma; \mathcal{G}, \tau_1, \tau_2 \mid F))| < \Lambda_{\tau_1}$.

2. $|\rho(\sigma;\mathcal{F},\tau_1,\tau_2 \mid F) - (\rho(\sigma;\mathcal{G},\tau_1,\tau_2 \mid F))|_\infty < \Lambda_{\tau_1}$.

Let $v:\Omega \to \Re^n$ be a $\mathcal{G}_{\tau_2}$-measurable function. Denote by $\Gamma^v_{\tau_1,\tau_2}$ the partial game which starts at stage τ_1, and if not terminate before, it terminates at stage τ_2 with payoff v. Denote by $\gamma_v(\sigma;\mathcal{F},\tau_1,\tau_2 \mid F)$ the conditional expected payoff given F under σ, in this game. The following lemma states that if $(\mathcal{G},(q_{G,k}))$ is a λ-approximation of Γ between τ_1 and τ_2, and if the opponent plays a $(\mathcal{G},\tau_1,\tau_2)$-strategy, then the player does not lose much in $\Gamma^v_{\tau_1,\tau_2}$ by considering only $(\mathcal{G},\tau_1,\tau_2)$-strategies (rather than $(\mathcal{F},\tau_1,\tau_2)$-strategies). The proof is similar to that of Lemma 6.4 in Shmaya and Solan (2004).

Lemma 4.6 *Let $(\mathcal{G}_k,(q_{G,k}))$ be a λ-approximation of Γ between τ_1 and τ_2. For every $F \in \mathcal{A}$, such that $\mathbf{P}(F) > 0$, every $(\mathcal{G},\tau_1,\tau_2)$-strategy profile σ^{-i}, and every $(\mathcal{F},\tau_1,\tau_2)$-strategy $\overline{\sigma}^i$ of player i, there is a $(\mathcal{G},\tau_1,\tau_2)$-strategy σ^i of player i, such that*

$$|\gamma^i_v(\sigma^{-i},\overline{\sigma}^i;\mathcal{F},\tau_1,\tau_2 \mid F) - \gamma^i_v(\sigma;\mathcal{F},\tau_1,\tau_2 \mid F)| < \Lambda_{\tau_1}. \tag{9}$$

4.3 Coloring the periodic games

In this section we define a finite set of "colors", and attach for every triplet (k,τ,ω), a color $c_{k,\tau}(\omega)$ in that finite set. Choose $0 < \xi < \min\{\frac{\epsilon^2}{4},\frac{1}{2n}\}$ once and for all.

By Theorem 4.3, there is a $\mathcal{F}$-consistent function that assigns for every $k \geq 0$ and every bounded stopping time $\tau > k$, a λ-approximation of Γ between k and τ, $\left(\mathcal{G}^{k,\tau}_{k'},\left(q^{k,\tau}_{G,k'}\right)\right)$. For each atom $G \in \mathcal{G}^{k,\tau}_k$, and every $\omega \in G$, we identify $T_{k,\tau}(\omega) = T_{k,\tau}(G)$.

Fix a stage $k \in \mathbf{N}$, a bounded stopping time $\tau > k$, and $\omega \in \Omega$. Let $T_{k,\tau}(\omega)$ be the game on a tree that is defined by $\left(\mathcal{G}^{k,\tau}_{k'},\left(q^{k,\tau}_{G,k'}\right)\right)$.

From now on, suppose that $T_{k,\tau}(\omega)$ satisfies Conditions Q.1-Q.3. By Theorem 3.3, one of the following holds for $T_{k,\tau}(\omega)$: either (a) $(0^i)_{i\in I}$ is a stationary equilibrium in the game, or (b) the game has a social welfare player, so that for every $\xi > 0$ the game has a stationary subgame perfect ξ-equilibrium, or (c) Assumptions A.1-A.3 and Conditions Q.1-Q.3 hold for the game, so that by Theorem 3.8, the game has a subgame perfect ξ-equilibrium $\sigma_{k,\tau}(\omega)$ in pure strategies which is supported by at least two players.

Let $\sigma_{k,\tau}(\omega)$ be a subgame perfect ξ-equilibrium in $T_{k,\tau}(\omega)$ that satisfies the conditions of Theorem 3.3. Roughly, the color will indicate the type of the equilibrium $\sigma_{k,\tau}(\omega)$, namely, whether case (a), (b) or (c) above holds; if case (b) holds, the color will also indicate the identity of the social welfare player, and if case (c) holds, the color will also indicate the set of payoffs along the equilibrium path, as well as some information on the termination probabilities.

Because $T_{k,\tau}(\omega)$ is a periodic game, and $\sigma_{k,\tau}(\omega)$ is a subgame perfect ξ-equilibrium in that game, every continuation payoffs u under $\sigma_{k,\tau}(\omega)$ at the root of the tree at some period of the game is a subgame perfect ξ-equilibrium payoff. Let $U_{k,\tau}(\omega)$ be the set of all the payoff vectors that are a continuation payoff under $\sigma_{k,\tau}(\omega)$ at the root of the tree at some period of the game. For every ξ-equilibrium payoff $u \in U_{k,\tau}(\omega)$, denote by $\sigma_{k,\tau;u}(\omega)$ the ξ-equilibrium in $T_{k,\tau}(\omega)$ which correspond to u. $\sigma_{k,\tau;u}(\omega)$ is the reduction of $\sigma_{k,\tau}(\omega)$ to the appropriate subgame.

For every ξ-equilibrium payoff $u \in U_{k,\tau}(\omega)$, define $l_u \in L := \{1, 2, ..., 3B\} \cup \{\infty\}$ as the minimal number of periods of the game $T_{k,\tau}(\omega)$ under $\sigma_{k,\tau;u}(\omega)$, that are needed to ensure that even if one player deviates, the probability that the game terminates is at least $\frac{\xi^6}{72B}$. That is, if case (a) holds then $l_u = \infty$; if case (b) holds then $l_u \in \{1, \infty\}$; and if case (c) holds then $l_u \in \{1, 2, ..., 3B\}$.

Let

$$UL_{k,\tau}(\omega) = \left\{ (u, l_u) \middle| u \in U_{k,\tau}(\omega) \right\}.$$

Although the set of payoff R is finite, the set of equilibrium payoff is not necessarily finite. We therefore have to approximates the set $UL_{k,\tau}(\omega)$ by a finite set.

Choose $Q \in \mathbf{N}$ that satisfies $Q > 6$ and

$$\left(1 - \frac{\xi^6}{144B}\right)^{\frac{Q}{3B}-1} < \frac{\epsilon}{4}. \tag{10}$$

Let $Y = \left\{0, \pm\frac{1}{W}, \pm\frac{2}{W}, ..., \pm\frac{W}{W}\right\}^n$, where $W > \frac{4Q}{\epsilon}$. Let $\mathcal{YL} = 2^{Y \times L}$ be the set of all the subsets of $Y \times L$.

For every periodic game $T_{k,\tau}(\omega)$, let $YL_{k,\tau}(\omega) \in \mathcal{YL}$ be a discretization of the set $UL_{k,\tau}(\omega)$. The set $YL_{k,\tau}(\omega)$ is obtained by replacing each pair $(u, l_u) \in UL_{k,\tau}(\omega)$ by a pair $(y, l_u) \in Y \times L$ such that $\|u - y\|_\infty \leq \frac{1}{2W}$. Note that the number of elements in $YL_{k,\tau;u}(\omega)$ is uniformly bounded: $|YL(s,\tau)| \leq |Y| \times |L| = (2W)^n \times (3B+1)$.

Denote by $I^P_{k,\tau}(\omega)$ the set of all players $i \in I$ such that $\sigma^i_{k,\tau}(\omega)$ is a pure strategy different from 0^i:

$$I^P_{k,\tau}(\omega) := \left\{ i \in I \ ; \ \sigma^i_{k,\tau}(\omega) \text{ is a pure strategy } \neq 0^i \right\}.$$

Denote by $I^M_{k,\tau}(\omega)$ the set of all players such that $\sigma^i_{k,\tau}(\omega)$ is a non-pure strategy (which is necessarily different from 0^i),

$$I^M_{k,\tau}(\omega) := \left\{ i \in I \ ; \ \sigma^i_{k,\tau}(\omega) \text{ is a non-pure strategy} \right\}.$$

The set $I^P_{k,\tau}(\omega) \bigcup I^M_{k,\tau}(\omega)$ is the set of all the players who stop with positive probability according to $\sigma_{k,\tau}(\omega)$.

Let $\mathcal{I} = 2^I$ be the set of all the subset of I.

We are now ready to attach a color for every triplet (k, τ, ω):
For every positive integer $k \in \mathbf{N}$, every bounded stopping time $\tau > k$, and every $\omega \in \Omega$ such that $T_{k,\tau}(\omega)$ satisfies Conditions Q.1-Q.3, attach the color $c_{k,\tau}(\omega) := \left(I^P, I^M, YL\right)_{k,\tau}(\omega)$ from the finite set $\mathcal{I} \times \mathcal{I} \times \mathcal{YL}$.

Observe that if case (a) holds, then $I^P_{k,\tau}(\omega) = I^M_{k,\tau}(\omega) = \emptyset$; if case (b) holds, then $|I^P_{k,\tau}(\omega)| = 1$ and $|I^M_{k,\tau}(\omega)| \leq 1$; if case (c) holds, then $|I^M_{k,\tau}(\omega)| \geq 2$ and $I^P_{k,\tau}(\omega) = \emptyset$. In particular, $|I^M_{k,\tau}(\omega)| \leq 1$, and if $|I^M_{k,\tau}(\omega)| = 1$ then necessarily $|I^P_{k,\tau}(\omega)| = 1$.

Finally, if the game on a tree $T_{k,\tau}(\omega)$ does not satisfy at least one of the Condition Q.1-Q.3, then we attach to it the color $c_{k,\tau}(\omega) = (\emptyset, \emptyset, \emptyset)$.

4.4 Using a stochastic variation of Ramsey's Theorem

For every positive integer k, every stopping time $\tau > k$, and every $\omega \in \Omega$, we attached an element in a finite set. By a stochastic variation of Ramsey's Theorem (Theorem 4.3 in Shmaya and Solan (2004)), there exists an increasing sequence of bounded stopping times $\tau_1 < \tau_2 < ...$ such that $\tau_t(\omega) \in \{K_1, K_2, ...\}$ for every $t \in \mathbf{N}$, and every $\omega \in \Omega$, and $c_{\tau_1,\tau_2} = c_{\tau_2,\tau_3} = ...$, with high probability. That is,

$$\mathbf{P}\left(\left(I^P, I^M, YL\right)_{\tau_1,\tau_2} = \left(I^P, I^M, YL\right)_{\tau_2,\tau_3} = ...\right) > 1 - \frac{\delta}{2}. \tag{11}$$

To simplify notations, we exchange the index $\{\tau_t, \tau_{t+1}\}$ by t; for example, we set $I_t^P(\omega) := I^P_{\tau_t,\tau_{t+1}}$. Denote by $\mathcal{G}^t$ the λ-approximation of Γ between τ_t and τ_{t+1}.

We define a profile σ_* in Γ as follows:

1. The definition for the subgame that starts at τ_1.

   (i) We first choose a subgame perfect ξ-equilibrium payoff in the game $T_1(\omega)$, for every $\omega \in \Omega$ such that the game $T_1(\omega)$ satisfies Conditions Q.1-Q.3.
   For every $\omega \in \Omega$, if the set $YL_1(\omega) \neq \emptyset$, choose $(y_1(\omega), l_1(\omega)) \in YL(\omega)$, such that **the choice is $\mathcal{G}^1$-consistent**. That is, for every ω, ω' in the same atom in $\mathcal{G}^1$, $(y_1(\omega), l_1(\omega)) = (y_1(\omega'), l_1(\omega'))$.
   Let $t := 1$.
   Let X be a set of elements $\omega \in \Omega$. Set at first $X := \Omega$. During the process we will remove from X elements that, according to σ_*, the players are allowed to choose arbitrary actions in the subgame that starts at stage τ_t for some $t \in \mathbf{N}$.

   (ii) We next define a $(\mathcal{F}, \tau_t, \tau_{t+1})$-strategies, one for each player.
   For every $\omega \in X$:

      i. if $YL_t(\omega) = \emptyset$ (i.e., at least one of the Condition Q.1-Q.3 does not hold in $T_t(\omega)$), instruct each player to choose arbitrary actions from now on. Remove ω from X.

      ii. Likewise, if $\left(I^P, I^M, YL\right)_t(\omega) \neq \left(I^P, I^M, YL\right)_{t-1}(\omega)$, instruct each player to choose arbitrary actions from now on, and remove ω from X.

      iii. Else, choose $u_t(\omega) \in U_t(\omega)$, a ξ-equilibrium payoff in $T_t(\omega)$, which is $\mathcal{G}^t$-consistent, such that $\|u_t(\omega) - y_t(\omega)\|_\infty \leq \frac{1}{2W}$, and $(u_t(\omega), l_t(\omega)) \in UL_t(\omega)$.

At stages τ_t given ω until $\tau_{t+1} - 1$, the players should follow $\sigma_{t;u_t(\omega)}(\omega)$, the ξ-equilibrium in $T_t(\omega)$ which correspond to the payoff $u_t(\omega)$.

(iii) We now find, for every continuation payoff at τ_{t+1} in every $T_t(\omega)$, an approximate payoff in $YL_t(\omega)$:
For every $\omega \in X$, let $\bar{u}_{t+1}(\omega)$ be the continuation payoff at τ_{t+1} given ω in $T_t(\omega)$ under $\sigma_{t;u_t(\omega)}(\omega)$, and $l_{t+1}(\omega) = l_{\bar{u}_{t+1}(\omega)}$.
If $YL_{t+1}(\omega) \neq \emptyset$, and $YL_t(\omega) = YL_{t+1}(\omega)$, then choose $y_{t+1}(\omega) \in Y_{t+1}(\omega)$ such that $\|y_{t+1}(\omega) - \bar{u}_{t+1}(\omega)\|_\infty \leq \frac{1}{2Y}$, and the choice is consistent with $\mathcal{G}^{t+1}$.
Go back to (ii) with $t+1$.

2. Use backward induction to define a subgame perfect equilibrium in the finite game that terminate at stage τ_1 with terminal payoff $y_1(\omega)$ for every $\omega \in \Omega$ such that the set $YL_1(\omega) \neq \emptyset$, and $y_1(\omega) = \overrightarrow{0}$ for every $\omega \in \Omega$ such that the set $YL_1(\omega) = \emptyset$.

Notice, the profile σ_* is a well defined profile in Γ, since every selection along the process is consistent with $\mathcal{G}^{t+1}$, for every $t \in \mathbf{N}$. That is, for every $t \in \mathbf{N}$ we choose a $(\mathcal{G}, \tau_t, \tau_{t+1})$-strategy profile, which, as we already mentioned, defines a $(\mathcal{F}, \tau_t, \tau_{t+1})$-strategy profile, which prescribes the players what to play in Γ between stages τ_t and τ_{t+1} (excluded).

4.5 The profile σ_* is a δ-approximate subgame perfect ϵ-equilibrium

We now prove that the profile σ_* is a δ-approximate subgame perfect ϵ-equilibrium. Namely, there is an event $G \in \mathcal{A}$ with $\mathbf{P}(G) < \delta$, such that for every stage $K \in \mathbf{N}$, and every event $F \in \mathcal{F}_k$, such that $\mathbf{P}(F) > 0$ and $F \cap G = \emptyset$, $\sigma_{|K}$ is an ϵ-equilibrium in $\Gamma_{|K}$ given F.

At first we define the event G. Let G_1 be an event that includes all the elements $\omega \in \Omega$, in which there are at least two finite games, $T_t(\omega)$ and $T_{t+1}(\omega)$ (for $t \in \mathbf{N}$), which do not have the same color, i.e., $\left(I^P, I^M, YL\right)_t(\omega) \neq \left(I^P, I^M, YL\right)_{t+1}(\omega)$. On G_1 the profile σ_* need not be an ϵ-equilibrium, but by Eq. (11), $\mathbf{P}(G_1) < \frac{\delta}{2}$.

Let G_2 include all the elements $\omega \in \Omega$, in which the game $T_t(\omega)$ does not satisfies at least one of the Condition Q.1-Q.3, for every $t \in \mathbf{N}$. On G_2 the profile σ_* need not be an ϵ-equilibrium. By Assumption A.4, and since the stopping times $(\tau_t)_{t\in\mathbf{N}}$ were limited to the set of stages $\{K_1, K_2, ...\}$, it follows that $\mathbf{P}(G_2) < \frac{\delta}{2}$.

We set $G := G_1 \cup G_2$. Then $\mathbf{P}(G) < \mathbf{P}(G_1) + \mathbf{P}(G_2) < \delta$.

For every $\omega \in \Omega$ that does not belong to G, the same color is attached to all the games $T_t(\omega)$, for $t \in \mathbf{N}$. That is, for every $\omega \in \Omega \backslash G$, $\left(I^P, I^M, YL\right)_t(\omega) = \left(I^P, I^M, YL\right)_{t+1}(\omega)$, for every $t \in \mathbf{N}$. In addition, Conditions Q.1-Q.3 are satisfied in all the games $T_t(\omega)$, for every $t \in \mathbf{N}$.

We next prove that for every stage $K \in \mathbf{N}$, and every event $F \in \mathcal{F}_K$ such that $F \cap G = \emptyset$ and $\mathbf{P}(F) > 0$, $\sigma_{*|K}$ is an ϵ-equilibrium in $\Gamma_{|K}$ given F.

Fix $K \in \mathbf{N}$ and $F \in \mathcal{F}_K$ such that $F \cap G = \emptyset$ and $\mathbf{P}(F) > 0$. We define a partition of the set F to a finite number of subsets, according to cases (a),(b), and (c) in Section 4.3. For each subset, we verify that the upper bound over the amount every player can gain by deviating is at most ϵ. Hence, we conclude that no player can gain more than ϵ, by deviating in $\Gamma_{|K}$ given F.

1) $\underline{(0^i)_{i\in I} \text{ is a stationary equilibrium:}}$
Let F_1 includes all the elements $\omega \in F$, such that $I_t^P(\omega) = I_t^M(\omega) = \emptyset$, for every $t \in \mathbf{N}$. Thus, for every $\omega \in F_1$, each game $T_t(\omega)$ has a stationary equilibrium, such that all the players continue whenever chosen. In particular, for every $i \in I$, $a_\infty^i \geq 0$, and every player expects to receive at least 0. However, by Condition Q.2, each player can gain at most 0 by terminating the game by himself, in every subgame which starts at K, so no deviation is worthwhile.

Note that, if $a_\infty^i \geq 0$, for every $i \in I$, then $F_1 = F$, which implies that no player can gain by deviating in $\Gamma_{|K}$ given F, and we are done.

If that not the case, i.e., there is at east one player $i \in I$ such that $a_\infty^i < 0$, then $F_1 = \emptyset$. Assume from now on, that there is at least one player $i \in I$ such that $a_\infty^i < 0$.

2) $\underline{j' \text{ is a social welfare player and } a_\infty^{j'} < 0}$:
For every $j' \in I$, let $F_2^{j'}$ be the set that includes all the elements $\omega \in F$, such that the game $T_t(\omega)$ satisfies Conditions Q.1-Q.3, $I_t^P(\omega) = \{j'\}$, and $I_t^M(\omega) = \emptyset$, for every $t \in \mathbf{N}$. Player j' is a social welfare player in each game $T_t(\omega)$; he has a pure stationary strategy $\sigma_t^i(\omega)$ that ensures every player $i \in I$ (including j') receives a non-negative payoff in $T_t(\omega)$, given the game terminates between stages τ_t and $\tau_{t+1} - 1$. In particular, the profile in $T_t(\omega)$, according to which player j' follows this strategy, and all the other continue whenever chosen, is a stationary equilibrium in $T_t(\omega)$ with equilibrium payoff at least 0, for each player, for every $t \in \mathbf{N}$.

Observe the game Γ, given $\omega \in F_2^{j'}$ is chosen. $\sigma_*(\omega)$ instruct all the players, except for j', to continue whenever chosen, while j' has a terminating pure strategy, which is derived from the strategies $\left(\sigma_t^{j'}(\omega)\right)_{t\in\mathbf{N}}$ of the games $(T_t(\omega))_{t\in\mathbf{N}}$. By Lemma 4.5, every player expect to receive at least $0 - \Lambda_t$ if the game Γ terminates between τ_t and $\tau_{t+1} - 1$, thus the expected payoff for each player in Γ, is at least $-\Lambda_0 = -\xi^6/432B^2 > -\frac{\epsilon}{8}$. However, if player $i \neq j'$ deviates and terminates the game, he receive at most 0 (by Q.2), while if player j' deviate and continue, the game never terminates, and he receive a negative payoff. Therefore, every player can gain at most $\frac{\epsilon}{8}$. Note that this argument is valid in every subgame.

3) $\underline{j' \text{ is a social welfare player and } a_\infty^{j'} \geq 0}$:
For every two players $j', j'' \in I$, let $F_3^{j',j''}$ be the set that includes all the elements $\omega \in F$, such that the game $T_t(\omega)$ satisfies Conditions Q.1-Q.3, $I_t^P(\omega) = \{j'\}$, $I_t^M(\omega) = \{j''\}$, and $Y_t(\omega) = \{y(\omega)\}$, for every $t \in \mathbf{N}$. These sets are similar to the previous case. Player j' is again a social welfare player in every game $T_t(\omega)$, who has a pure stationary strategy $\sigma_t^{j'}$ that ensures every player $i \in I$ (including j') receives a non-negative payoff in $T_t(\omega)$ given the game terminates between stages τ_t and $\tau_{t+1} - 1$, for every $t \in \mathbf{N}$. In addition, player j'' has a mixed strategy $\sigma_t^{j''}$, which threatens player j', by stopping the game with some sufficient small probability, at stages where the terminal payoff for player j' is negative, for every $t \in \mathbf{N}$. Since $y(\omega)$ is the approximated ξ-equilibrium payoff in the game $T_t(\omega)$, given the game terminates between stages τ_t and $\tau_{t+1} - 1$, then every player receive at least $y(\omega) - \frac{1}{2W}$ under this ξ-equilibrium, and at most $y(\omega) + \frac{1}{2W} + \xi$ by deviating, given the game $T_t(\omega)$ terminates between stages τ_t and $\tau_{t+1} - 1$, for every

$t \in \mathbf{N}$.

Observe the game Γ, given $\omega \in F_3^{j',j''}$ is chosen. $\sigma_*(\omega)$ instruct all the players, except for j', j'', continue whenever chosen, while j' has a terminating pure strategy, and j'' has a terminating mixed strategy, which are derived from the strategies in the games $T_t(\omega)$.

By Lemma 4.5, the expected payoff of player i in Γ is at least $y^i - \frac{1}{2W}\Lambda_t$, given $\omega \in F_5^{j',j''}$ is chosen, the players follow $\sigma_*(\omega)$, and the game terminated between stages τ_t and $\tau_{t+1} - 1$.

On the other hand, assume player i deviates in Γ, and his expected payoff is v_t, given the game terminates between stages τ_t and $\tau_{t+1} - 1$. By Lemma 4.6, there is a $(\mathcal{G}, \tau_1, \tau_2)$-strategy of player i in which he can lose at most Λ_t by considering this strategy, instead of the original, so its expected payoff is at least $v_t - \Lambda_t$, in Γ, given the game terminates between stages τ_t and $\tau_{t+1} - 1$. By Lemma 4.5, if player i deviates to this $(\mathcal{G}, \tau_1, \tau_2)$-strategy in $T_t(\omega)$, his expected payoff is at least $v_t - 2\Lambda_t$ given the game terminates between stages τ_t and $\tau_{t+1} - 1$, on the other hand, as we already mentioned, this amount must be at most $y^i + \xi + \frac{1}{2W}$, so, $v_t \le y^i + \frac{1}{2W} + \xi + 2\Lambda_t$.

Therefore, every player can gain at most $\xi + 3\Lambda_t + \frac{1}{W}$ by deviating in Γ given $\omega \in F_3^{j',j''}$ is chosen, and the game terminates between stages τ_t and $\tau_{t+1} - 1$. As a result, every player can gain at most $\xi + 3\Lambda_0 + \frac{1}{W} < \frac{\epsilon}{8}$ by deviating in Γ given $\omega \in F_3^{j',j''}$ is chosen. These arguments are valid in every subgame, which start at stage K.

4) <u>There is no social welfare player:</u>
For every subset of players $J \subseteq I$, $|J| \ge 2$, and every $\overline{YL} \in \mathcal{YL}$, let $F_4^{J,\overline{YL}}$ be the set that includes all the elements $\omega \in F$, such that the game $T_t(\omega)$ satisfies Conditions Q.1-Q.3, $I_t^P(\omega) = J$, $I_t^M(\omega) = \emptyset$, and $YL_t(\omega) = \overline{YL}$. Meaning, every player $i \in J$ use a pure strategy which is different from 0^i in $T_t(\omega)$, and the set of equilibrium payoff is $\overline{Y}$, for every $t \in \mathbf{N}$.

In order to prove that no deviation in $\Gamma_{|K}$ can be profitable, we further partition the set $F_4^{J,\overline{YL}}$. For every $t' \in \mathbf{N}$, let

$$F_{4,t'}^{J,\overline{YL}} = \left\{\omega \in F_4^{J,\overline{YL}} \mid \tau_{t'}(\omega) \leq K < \tau_{t'+1}(\omega)\right\}$$

The collection $\left(F_{4,t'}^{J,\overline{YL}}\right)_{t'=1}^{K}$ is a finite partition of $F_4^{J,\overline{YL}}$. Fix $t' \in \{1, 2, ..., K\}$ such that $P(F_{4,t'}^{J,\overline{YL}}) > 0$.

We define a new game on a tree T using the games $(T_t(\omega))_{t=t'}^{t'+Q}$ on $F_{4,t'}^{J,\overline{YL}}$ as follows: it starts at stage K in $(T_{t'}(\omega))_{\omega \in F_{4,t'}^{J,\overline{YL}}}$, and for every $\omega \in F_{4,t'}^{J,\overline{YL}}$, if the game $T_{t'}(\omega)$ reaches a leaf, T continues with $T_{t'+1}(\omega)$, and so on.

σ_* induces a strategy in T, in which every player can gain at most $\xi + 2Q\frac{1}{2W}$, given T is terminated before $\tau_{t'+Q}$. Indeed, every player can gain at most ξ by deviating in the game $T_t(\omega)$, given the game is terminated between stage τ_t and τ_{t+1}. Furthermore, a player might use the fact that during the construction of σ_* we used an approximate continuation payoff (cf. Step (iii)), so by postponing the termination of the game to $\tau_{t'+Q}$ at least, he can gain at most $2Q\frac{1}{2W}$.

By applying Lemmas 4.5 and 4.6 to each of the games $T_t(\omega)$, it follows that if a player deviates in Γ given $F_{4,t'}^{J,\overline{YL}}$ occurs and the game is terminated between stage K and $\tau_{t'+Q}$, then he can gain at most

$$2\sum_{t=t'}^{t'+Q} \Lambda_t + \xi + 2Q\frac{1}{2W} < \frac{\epsilon}{2}.$$

We next claim that, if all the players, except perhaps for one, follow σ_*, the game Γ given $F_{4,t'}^{J,\overline{YL}}$ is terminated between stage K and $\tau_{t'+Q}$ with probability at least $1 - \frac{\epsilon}{4}$. In fact, by Theorem 3.3, it follows that the game T is terminated before it reaches $\tau_{t'+Q}$ with probability at least $1 - \left(1 - \frac{\xi^6}{72B}\right)^{\frac{Q}{3B}-1}$. In addition, by Lemma 4.5, for every $t \in \mathbf{N}$, the difference between the termination probability between stages τ_t and $\tau_{t+1} - 1$ under any strategy profile, in $(T_t(\omega))_{\omega \in F_{4,t'}^{J,\overline{YL}}}$, and in Γ given $F_{4,t'}^{J,\overline{YL}}$ occurs, is at most Λ_{τ_t}. Hence, under σ_* the game Γ given $F_{4,t}^{J,\overline{YL}}$ is terminated between stage K and $\tau_{t'+Q}$ with probability at least $1-\left(1-\frac{\xi^6}{144B}\right)^{\frac{Q}{3B}-1} > 1-\frac{\epsilon}{4}$

To summarize, every player can gain by deviation in $\Gamma_{|K}$ given $F_{4,t}^{J,\overline{YL}}$ at most $\frac{\epsilon}{2}$, if the game is terminated between stage K and $\tau_{t'+Q}$, while he can gain at most 2, if the game terminates after stage $\tau_{t'+Q}$, therefore, he can gain by deviating at most

$$\frac{\epsilon}{2}\cdot\left(1-\frac{\epsilon}{4}\right)+2\cdot\frac{\epsilon}{8}<\epsilon.$$

■

References

[1] Dynkin, E.B. (1969) Game Variant of a Problem on Optimal Stopping, *Soviet Mathematics Doklady*, **10**, 270-274.

[2] Fine, C.H., L. Li. (1989) Equilibrium Exit in Stochastically Declining Industries, *Games and Economic Behavior*, **1**, 4059.

[3] Flesch, J., F. Thuijsman, O.J. Vrieze. (1997) Cyclic Markov Equilibria in a Cubic Game, *International Journal of Game Theory*, **26**, 303-314.

[4] Flesch, J., J. Kuipers, G. Schoenmakers, K. Vrieze. (2008) Subgame-Perfection in Stochastic Games with Perfect Information and Recursive Payoffs. Preprint.

[5] Ghemawat, P., B. Nalebuff. (1985) Exit. RAND, *Journal of Economics*, **16**, 184194.

[6] Karlin, S. (1959) Mathematical Methods and Theory in Games, *Programming and Economics Vol. II: The Theory of Infinite Glames*, **Addison-Wesley, Reading, MA**.

[7] Kifer, Y. (1971) Optimal Stopped Games, *Theory of Probability and its Application*, **16**, 185-189.

[8] Kifer, Y (2000) Game Options, *Finance and Stochastics*, **4**, 443-463.

[9] Maitra, A.P., W.D. Sudderth (2007) Subgame-Perfect Equilibria for Stochastic Games, *Mathematics of Operations Research*, **32**, 711-722.

[10] Mashiah-Yaakovi, A.(2008) Periodic Stopping Games, *International Journal of Game Theory*, to apear.

[11] Mertens, J. F. (1987) Repeated Games, *Proceedings of the International Congress of Mathematicians, Berkeley, 1986, Gleason A.M. eds, American Mathematical Society*, 1528-1577.

[12] Neveu, J.(1975) Discrete-Parameter Martingales, *North-Holland, Amsterdam.*

[13] Rosenberg, D., E. Solan, N. Vieille. (2001) Stopping Games with Randomized Strategies, *Probability Theory and Related Fields*, **119**, 433-451.

[14] Shmaya, E., E. Solan. (2004) Two Player Non Zero-sum Stopping Games in Discrete Time, *The annals of probability*, **32**, 2733-2764.

[15] Solan, E. (2005) Subgame-Perfection in Quitting Games with Perfect Information and Differential Equations, *Mathematics of Operations Research*, **30**, 51-72.

[16] Solan, E., N. Vieille. (2003), Deterministic Multi-Player Dynkin Games, *Journal of Mathematical Economics*, **1097**, 1-19.