Int J Game Theory (2009) 38:169–181
DOI 10.1007/s00182-008-0143-4

ORIGINAL PAPER

Periodic stopping games

Ayala Mashiah-Yaakovi

Accepted: 28 September 2008 / Published online: 28 October 2008

Abstract Stopping games (without simultaneous stopping) are sequential games in which at every stage one of the players is chosen, who decides whether to continue the interaction or stop it, whereby a terminal payoff vector is obtained. Periodic stopping games are stopping games in which both of the processes that define it, the payoff process as well as the process by which players are chosen, are periodic and do not depend on the past choices. We prove that every periodic stopping game without simultaneous stopping, has either periodic subgame perfect ϵ-equilibrium or a subgame perfect 0-equilibrium in pure strategies.

Keywords Stopping games · Dynkin games · Stochastic games · Subgame-perfect equilibrium

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.

This work is part of the master thesis of the author done under the supervision of Prof. Eilon Solan. I am thankful to Prof. Solan for his inspiring guidance. I also thank two anonymous referees of the International Journal of Game Theory for their comments.

A. Mashiah-Yaakovi
School of Mathematical Sciences, Tel Aviv University, 69978 Tel Aviv, Israel
e-mail: ayalam@post.tau.ac.il

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

Since the game has perfect information, by Mertens (1987) it follows that every multi-player stopping game has an ϵ-equilibrium. Since the ϵ-equilibrium strategies that were constructed by Mertens (1987) employ threats of punishment, which might be non-credible, subsequent work concentrated on the existence of subgame perfect equilibria.

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. 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. Solan leaves open the question whether the same result holds when the process by which players are chosen, as well as the terminal payoff process, are non-stationary.

Stopping games in which players can stop simultaneously have been studied by, e.g., Kiefer (1971), Neveu (1975), Rosenberg et al. (2001) and Shmaya and Solan (2004). These authors provide conditions for the existence of the value in pure strategies, and prove that in general the value in mixed strategies exists, as well as a mixed ϵ-equilibrium in two player non-zero-sum games.

In this paper, we answer the open problem posed in Solan (2005), and show that every periodic stopping game has either a periodic ϵ-equilibrium or a subgame perfect 0-equilibrium in pure strategies. Our approach is to construct a sequence of equilibria in finite stopping games, and using a diagonal extraction argument to construct an equilibrium in the infinite game. The novelty of our approach is that the finite stage games that we consider do not terminate at a given stage t [as done in Rosenberg et al. (2002) and Solan and Vieille (2003)], but rather at a properly chosen stopping time, such that the probability that at least one player stops in an equilibrium of this game is bounded away from 0.

Our study can be useful to the study of both stopping games and stochastic games. First, it seems possible that our approach, together with the technique developed in Shmaya and Solan (2004), may be used to prove the existence of a subgame perfect equilibrium in every stopping game (without simultaneous moves). Second, since periodic stopping games form a class of n-player stochastic games, our approach may prove useful to study equilibria in other classes of n-player stochastic games (cf., e.g., Flesch et al. (2007) and Flesch et al. (2008) for recent results in this area).

The paper is organized as follows: In Sect. 2, we present the model, some basic definitions and the main result, and introduce an example that illustrates some features of the model. In Sect. 3, we present the proof of the result for periodic stopping games.

2 The model and the main results

2.1 The model

A *Markovian stopping game* is given by $\Gamma = (I, (p(t), (a_i(t))_{i \in I})_{t \in \mathbf{N}}, a_\infty)$ where:

- $I = \{1, \ldots, n\}$ is a non-empty finite set of players.
- $p(t) \in \Delta(I)$ is[1] a distribution over the players according to which players are chosen at stage $t \in \mathbf{N}$.
- $a_i(t) \in \Re^n$ is a terminal payoff, for every $i \in I$ and every $t \in \mathbf{N}$.
- $a_\infty \in \Re^n$ is a payoff vector.

The game is played as follows: at every stage $t \in \mathbf{N}$ one player $i_t \in I$ is chosen according to the distribution $p(t)$, and that player decides whether to stop the game or to continue. Thus, players are not allowed to stop simultaneously. The game terminates as soon as one player chooses to stop, in which case the terminal payoff vector is $a_{i_t}(t)$. If the game never terminates, the players receive a payoff according to[2] a_∞. The aim of every player is to maximize his nominal payoff.

Definition 2.1 A Markovian stopping game is called a *periodic stopping game* if there is an integer $C \geq 1$, such that $p(t) = p(t \bmod C)$ and $a_i(t) = a_i(t \bmod C)$ for every $t \in \mathbf{N}$ and every $i \in I$. The minimal C that satisfies these conditions is the *period* of the game.

2.2 Strategies and main results

In a stopping game, the set of actions available to any player at any stage is $A = \{\text{Continue, Stop}\}$. To save notations, we assume that the players choose actions even after the game terminates.

Let $H_t = (I \times A)^t$ be the space of all histories of length t, $H = \cup_{t \in \mathbf{N}} H_t$ be the space of all finite histories, and $H_\infty = (I \times A)^{\mathbf{N}}$ be the set of all infinite histories. H_∞, equipped with the σ-algebra spanned by the cylinder sets, is a measurable space. We denote by $\mathcal{H}_t$ the sub-σ-algebra induced by the cylinder sets defined by H_t.

Definition 2.2 A *pure strategy* for player $i \in I$ is a function $\sigma^i : H \to \{0, 1\}$. For every $h \in H_t$, $\sigma^i(h) = 1$ if player i stops the game in case history h occurs and player i is chosen at stage $t+1$ (provided the game did not terminate before that stage), while $\sigma^i(h) = 0$ if player i continues if history h occurs and player i is chosen at stage $t+1$.

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 function $\sigma^i : H \to [0, 1]$. For every $h \in H_t$, $\sigma^i(h)$ is the probability that player i stops the game if history h occurs

[1] For every finite set I, $\Delta(I)$ is the set of probability distributions over I.

[2] a_∞ will not be normalized to 0 as could seem natural at this point, because another normalization will be used later.

and player i is chosen at stage $t+1$ (provided the game did not terminate before that stage).

A strategy σ^i of player i is called $\widehat{C}$*-periodic*, for some $\widehat{C} \in \mathbf{N}$, if $\sigma^i(h) = \sigma^i(h', h)$ for every $t \in \mathbf{N}$, every $h \in H_t$, and every $(h', h) \in H_{\widehat{C}+t}$. Here (h', h) is the concatenation of h' and h.

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. Each profile σ induces a probability distribution P_σ over the set of plays.

Let θ be the terminal stage, the first stage in which the chosen player chooses to stop. In case the game never terminates we set $\theta = +\infty$. If $\theta < \infty$, i_θ is the player who terminates the game. A profile σ is called *terminating* if $P_\sigma(\theta < \infty) = 1$.

Let E_σ be the expectation operator that corresponds to P_σ. The expected payoff vector under σ is

$$\gamma(\sigma) = E_\sigma \left[1_{\{\theta<\infty\}} a_{i_\theta}(\theta) + 1_{\{\theta=\infty\}} a_\infty \right].$$

Definition 2.4 Let $\epsilon \geq 0$. A profile σ is an ϵ*-equilibrium* if for every player $i \in I$ and every strategy $\overline{\sigma}^i$ of player i,

$$\gamma^i(\sigma) \geq \gamma^i(\sigma^{-i}, \overline{\sigma}^i) - \epsilon.$$

The vector $\gamma(\sigma)$ is called an ϵ*-equilibrium payoff vector*.

For every finite history $h \in H$, we denote by $\Gamma_{|h}$ the restriction of Γ to the subgame that start after history h occurs. Given history $h \in H$ and strategy σ^i of player $i \in I$, we define a strategy $\sigma^i_{|h}$ of player i in $\Gamma_{|h}$ by

$$\sigma^i_{|h}(h') = \sigma^i(h, h') \quad \forall h' \in H.$$

$\sigma^i_{|h}$ is the continuation strategy of player i given history h occurs. Let $\gamma_{|h}(\sigma)$ be the expected payoff vector which corresponds to the profile $\sigma_{|h}$ in the subgame $\Gamma_{|h}$.

Definition 2.5 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 $\Gamma_{|h}$.

Every stopping game is a game with perfect information, so by Mertens (1987) it has an ϵ-equilibrium. However, the existence of a subgame perfect ϵ-equilibrium is still an open problem. In this paper, we prove the existence of a subgame perfect $\epsilon/0$-equilibrium in the class of periodic stopping games.

The main result of the paper is:

Theorem 2.6 *Every periodic stopping game has either a periodic subgame perfect* ϵ*-equilibrium, or a subgame perfect* 0*-equilibrium in pure strategies.*

The period $\widehat{C}$ of the subgame perfect ϵ-equilibrium that we construct is either 1 or C, the period of the game. Recall that if simultaneous stopping is allowed, the length of the subgame perfect ϵ-equilibrium may be unbounded, as ϵ goes to 0 (Solan 2001).

The result of Theorem 2.6 was proved by Solan (2005) in the case $C = 1$. The proof of Solan uses tools from differential inclusions. It is not clear whether Solan's technique can be used for $C > 1$. Our proof uses a completely different technique, which is closer in spirit to other proofs in the field, as will be clarified in Sect. 2.3.

2.3 Example

We provide here an example that illustrates some features of the model. The example exhibits a periodic stopping game with period 2, that has a subgame perfect 0-equilibrium in pure strategies.

Example 1 Let $I = \{1, 2, 3, 4\}$, be a set of players. At every odd stage player 1 is chosen with probability 1, and at every even stage one player of $I_2 = \{2, 3, 4\}$ is chosen according to a uniform distribution. The terminal payoff vectors are stationary, and given by

$$\begin{aligned} a_1 &:= a_1(t) = (0, 3, -1, -1), && t = 1, 2, \ldots \\ a_2 &:= a_2(t) = (-4, 0, 6, 6), && t = 1, 2, \ldots \\ a_3 &:= a_3(t) = (1, -1, 0, 2), && t = 1, 2, \ldots \\ a_4 &:= a_4(t) = (1, -1, 2, 0), && t = 1, 2, \ldots \\ a_\infty &= (1, -1, 1, 1) \end{aligned}$$

The stationary profile $(0^i)_{i \in I}$, in which no player ever stops, is not an equilibrium. Indeed, in every subgame player 2 is better off by deviating and stopping whenever he is chosen. It follows that every subgame perfect ϵ-equilibrium is necessarily terminating.

We now construct a profile σ that is a subgame perfect 0-equilibrium, by constructing a sequence of finite games, such that every game terminates by one of the players with a probability which is bounded away from 0. For every positive integer l, every bounded stopping time $\tau > l$, and every real v, let $\Gamma(l, \tau, v)$ be an auxiliary finite-stage stopping game that is the same as the original game with the following changes:

(1) It starts at stage l of the original game: the probability distribution according to which a player is chosen at stage t, and the terminal payoff at that stage, are $p(l + t - 1)$ and $a(l + t - 1)$, respectively.

(2) If not terminated before, the game terminates at stage τ with a terminal payoff v.

In Figs. 1, 2 and 3, we describe three such auxiliary games, together with subgame perfect 0-equilibria in these games. In the figures below c indicates "continue", and s indicates "stop".

The game $\Gamma(l_1, \tau_1, v_1)$ is described in Fig. 1, where $v_1 = (-\frac{22}{27}, -\frac{2}{27}, \frac{70}{27}, \frac{58}{27})$, $l_1 = 1$ and $\tau_1 = 2$. This is a one-shot game with terminal payoff v_1. Since $v_1^1 < 0 = a_1^1$,

Fig. 1 The game $\Gamma(l_1, \tau_1, v_1)$

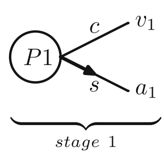

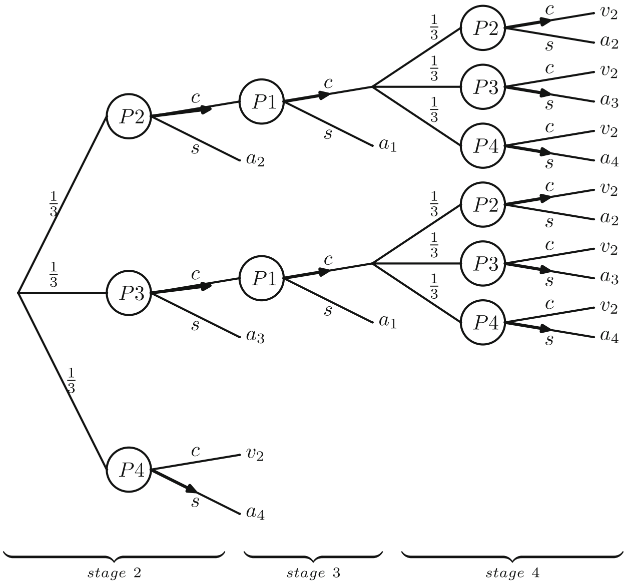

Fig. 2 The game $\Gamma(l_2, \tau_2, v_2)$

Fig. 3 The game $\Gamma(l_3, \tau_3, v_3)$

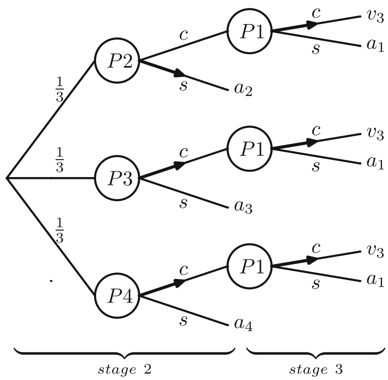

an equilibrium σ_1 in this game is that player 1 stops. The equilibrium payoff under σ_1 is $v_2 = (0, 3, -1, -1)$.

The game $\Gamma(l_2, \tau_2, v_2)$ that is described in Fig. 2 is not a t-stage game. Rather, here $l_2 = 2$, $\tau_2 = 3$ if player 4 is chosen at the first stage, and $\tau_2 = 5$ otherwise. The subgame perfect 0-equilibrium σ_2 is described in Fig. 2 by arrows. The equilibrium payoff under σ_2 is $v_3 = (\frac{7}{9}, -\frac{1}{9}, \frac{8}{9}, \frac{2}{9})$.

The game $\Gamma(l_3, \tau_3, v_3)$ that is described in Fig. 3 is a two stage game: $l_3 = 2$ and $\tau_3 = 4$. The subgame perfect 0-equilibrium σ_3, which can be found using a backward induction procedure, is described in Fig. 3 by arrows. The equilibrium payoff under σ_3 is $v_1 = (-\frac{22}{27}, -\frac{2}{27}, \frac{70}{27}, \frac{58}{27})$.

We use σ_1, σ_3 and σ_2 to construct a subgame perfect 0-equilibrium σ in the periodic stopping game: follow the profile σ_1, then σ_3 and then σ_2, and so forth. Since v_k, the terminal payoff in the auxiliary game $\Gamma(l_k, \tau_k, v_k)$, is equal to the equilibrium payoff in the auxiliary game $\Gamma(l_{k-1}, \tau_{k-1}, v_{k-1})$, and since in these strategies at least two players stop with positive probability (actually all players stop with positive probability), it follows that this is indeed a subgame perfect 0-equilibrium.

In the general case, we construct a denumerable set of auxiliary finite games, such that the equilibrium payoff in one game is the terminal payoff in the other. We then construct the equilibrium in the periodic stopping game by concatenating the equilibria in the auxiliary finite-stage games.

There is some similarity between our proof and that of Mertens and Parthasarathy (1987) or Solan (1998). In these papers, there is some parameter, the continuation payoff, which changes at every stage, and the players act corresponding to its value. Here the parameter does not change at every stage, but only occasionally.

3 The proof of Theorem 2.6

Let Γ be a periodic stopping game with period C, so that $p(t) = p(t \mod C)$ and $a_i(t) = a_i(t \mod C)$ for every $t \in \mathbf{N}$ and every $i \in I$.

3.1 Preliminaries

We assume w.l.o.g. that for every player $i \in I$, $\sum_{t=1}^{C} p_i(t) > 0$; Otherwise, player i is never chosen, and can be eliminated. Let

$$\delta := \min\{p_i(t) : i \in I,\ t \in \{1, \ldots, C\};\ p_i(t) > 0\}. \tag{1}$$

In addition, we assume w.l.o.g. that the absolute values of the payoffs are bounded by 1, and the maximal payoff which every player can gain by terminating the game by himself is zero, that is

$$\max\{a_i^i(t) : t \in \{1, \ldots, C\},\quad p_i(t) > 0\} = 0, \quad \forall\, i \in I.$$

Definition 3.1 A player $i \in I$ is called a *dummy player* if $a_\infty^i \geq 0$ and $a_j^i(t) \geq 0$ for every player $j \in I$, and every $t \in \mathbf{N}$ such that $p_j(t) > 0$.

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 $\overline{\Gamma}$ in which all dummy players were recursively eliminated. Every subgame perfect ϵ-equilibrium in $\overline{\Gamma}$ can be extended to a subgame perfect ϵ-equilibrium in Γ, by instructing all dummy players to continue whenever chosen. Therefore, we assume w.l.o.g. that there is no dummy player in Γ.

We now study some cases where a periodic subgame perfect ϵ-equilibrium exists. We then turn to the case that a subgame perfect 0-equilibrium in pure strategies exists.

3.2 Periodic equilibria

In this section, we identify two cases where a periodic equilibrium exists. We first discuss the case in which the vector a_∞ is non-negative.

Lemma 3.2 *If $a^i_\infty \geq 0$ for every $i \in I$, then $(0^i)_{i \in I}$, where all players always continue, is a stationary* 0*-equilibrium.*

From now on we assume the following:

A.1. There is $j \in I$ such that $a^j_\infty < 0$.

For the second type of equilibrium we need the next definition.

Definition 3.3 A *social welfare player* is a non-dummy player $i \in I$ who has a pure C-periodic strategy μ^i, such that

$$\gamma^i_{|h}(0^{-i}, \mu^i) = 0 \quad \forall h \in H \text{ and,} \tag{2}$$

$$\gamma^j_{|h}(0^{-i}, \mu^i) \geq \gamma^j_{|h}(0^{-i,-j}, \mu^i, \mu^j) \quad \forall j \neq i,\ \mu^j \neq 0^j \text{ and } h \in H. \tag{3}$$

By Assumption A.1, since μ^i is periodic, it is necessarily terminating. Furthermore, by Eq. (2) player i can gain 0 by following μ^i, and therefore under the strategy μ^i, player i stops only at stages t where $a^i_i(t) = 0$.

By Eq. (3), if i is a social welfare player then every player $j \neq i$ prefers that player i terminates the game according to the strategy μ^i rather he himself does so. This leads us to next type of equilibrium:

Lemma 3.4 *If there is a social welfare player, then there is a periodic subgame perfect ϵ-equilibrium, with period C.*

Proof Let $i \in I$ be a social welfare player. Since i is not dummy, either:

1. $a^i_\infty < 0$, in which the case the profile "player i follows μ^i and all the other players continue whenever chosen" is a periodic subgame perfect 0-equilibrium, or
2. $a^i_j(t') < 0$ for some $j \in I$, $t' \in \{1, 2, \ldots, C\}$ and $p_j(t') > 0$, in which the profile "player i follows μ^i, player j stops with some sufficient small probability ϵ' at every stage t such that $(t \mod C) = t'$ if chosen, and all the other players continue whenever chosen" is a periodic subgame perfect ϵ-equilibrium. □

The result of Lemma 3.4 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.2. There are no social welfare players.

Assumption A.1 ensures that the profile $(0^i)_{i \in I}$ (according to which all players always continue) is not a subgame perfect ϵ-equilibrium. Furthermore, by Assumption A.2, every player $i \in I$ is not a social welfare player. Therefore, for every C-periodic strategy μ^i of player i, there is some player $j \in I$ and a history $h \in H$ such that $\gamma^j_{|h}(0^{-i}, \mu^i) < 0$.

Under these assumptions we prove the next theorem:

Theorem 3.5 *If* A.1 *and* A.2 *hold, then the game has a subgame perfect* 0*-equilibrium in pure strategies.*

3.3 Proof of Theorem 3.5

To prove Theorem 3.5, we define a sequence of finite games.

As we have seen in Sect. 2.3, for every $l \in \mathbf{N}$, every bounded stopping time $\tau > l$, and every vector $v \in \Re^n$, we define an auxiliary finite-stage game $\Gamma(l, \tau, v)$ as the stopping game that starts at stage l and, if not terminated before by one of the players, terminates at stage τ with terminal payoff v. Observe that the terminal payoff is independent of past history. Furthermore, since Γ is a periodic game, we can restrict the initial stages to stages $1, \ldots, C$. Since τ is bounded, by a backward induction argument, $\Gamma(l, \tau, v)$ has a subgame perfect 0-equilibrium, σ. Denote by $\gamma^i(\sigma; \Gamma(l, \tau, v))$ the payoff of player i under σ in $\Gamma(l, \tau, v)$.

We will construct a sequence $(l_k)_{k\in\mathbf{N}}$ of initial stages, a sequence $(\tau_k)_{k\in\mathbf{N}}$ of bounded stopping times, a sequence $(v_k)_{k\in\mathbf{N}}$ of real vectors, and a sequence $(\sigma_k)_{k\in\mathbf{N}}$ of subgame perfect 0-equilibria in $(\Gamma(l_k, \tau_k, v_k))_{k\in\mathbf{N}}$ in pure strategies, such that the following conditions will hold:

- C.1. The stopping time τ_{k+1} is restricted to stages t such that $(t \mod C) = l_k$, for every $k \in \mathbf{N}$—the auxiliary finite-stage game $\Gamma(l_{k+1}, \tau_{k+1}, v_{k+1})$ ends in stages where $\Gamma(l_k, \tau_k, v_k)$ can begin.
- C.2. $v_{k+1} = \gamma(\sigma_k; \Gamma(l_k, \tau_k, v_k))$, for every $k \in \mathbf{N} - v_{k+1}$, the terminal payoff in $\Gamma(l_{k+1}, \tau_{k+1}, v_{k+1})$, is the equilibrium payoff in $\Gamma(l_k, \tau_k, v_k)$.
- C.3. The probability that the game $\Gamma(l_k, \tau_k, v_k)$ terminates before stage τ_k under σ_k is at least δ [cf. Eq. (1)], for every $k \in \mathbf{N}$.
- C.4. For every $i \in I$, for infinitely many k's, the probability that $\Gamma(l_k, \tau_k, v_k)$ terminates before stage τ_k under $(\sigma_k^{-i}, 0^i)$ is at least δ.

The proof of Theorem 3.5 is proven in two steps.

Proposition 3.6 *If* A.1 *and* A.2 *hold, then there are sequences* $(l_k)_{k\in\mathbf{N}}$, $(\tau_k)_{k\in\mathbf{N}}$, $(v_k)_{k\in\mathbf{N}}$ *and* $(\sigma_k)_{k\in\mathbf{N}}$ *in pure strategies such that* C.1–C.4 *hold.*

Proposition 3.7 *If there are sequences* $(l_k)_{k\in\mathbf{N}}$, $(\tau_k)_{k\in\mathbf{N}}$, $(v_k)_{k\in\mathbf{N}}$ *and* $(\sigma_k)_{k\in\mathbf{N}}$ *in pure strategies such that* C.1–C.4 *hold, then the game has a subgame perfect* 0*-equilibrium in pure strategies.*

We start by explaining the main ideas of the proofs of these two steps. The formal proof follows.

In order to prove Proposition 3.6, we will simultaneously construct by induction the four sequences. The sequence $(v_k)_{k\in\mathbf{N}}$ will also satisfy that $v_k \not\geq \overrightarrow{0}$ for every $k \in \mathbf{N}$. The proof uses the idea presented in the example of Sect. 2.3. By Condition C.1, the $k+1$ auxiliary finite-stage game only terminates at stages t, such that $(t \mod C) = l_k$, with terminal payoff v_{k+1} which, by Condition C.2, equal to the equilibrium payoff in $\Gamma(l_k, \tau_k, v_k)$. Suppose $v_{k+1} \not\geq \overrightarrow{0}$, hence there is a set of players $I' \neq \emptyset$, such that

$v^i_{k+1} < 0$, for every player $i \in I'$. Therefore, there is some player $i_* \in I'$, such that at some stage t_*, player i_* will strongly prefer stopping to continuing and getting a payoff $v^{i_*}_{k+1} < 0$. To simplify the presentation, assume $t_* < l_k$. Observe the finite-stage stopping game $\Gamma(t_*, l_k, v_{k+1})$ which starts at stage t_* and terminates at stage l_k with terminal payoff v_{k+1}. Using a backward induction, we can find an equilibrium in $\Gamma(t_*, l_k, v_{k+1})$ where i_* stops at stage t_*. Denote by v_* the equilibrium payoff in $\Gamma(t_*, l_k, v_{k+1})$.

If $v_* \not\geq \overrightarrow{0}$, then take $\Gamma(t_*, l_k, v_{k+1})$ as the $k+1$ auxiliary finite-stage stopping game.

If $v_* \geq \overrightarrow{0}$, we add some stages to the game $\Gamma(t_*, l_k, v_{k+1})$ before stage t_*, where we force player i_* to stop (e.g., in Fig. 2 in Example 1, Player 4 is forced to stop); If i_* is chosen at some stage t such that $(t \bmod C) = t_*$ and does not stop, the game terminates with terminal payoff v_{k+1}, so he receives a lower payoff than the payoff he receives by stopping, otherwise, the game continues to the next stage. Since i_* is not a social welfare player, adding enough stages, but not too much, where only player i has the incentive to stop, ensures that $v_{k+2} \not\geq \overrightarrow{0}$.

The proof of Proposition 3.7 uses a diagonal extraction argument—a limit of concatenations of $(\sigma_t)_{t\in\mathbf{N}}$ is the desired subgame perfect 0-equilibrium.

To state the formal construction we need additional notions.

By Assumption A.2 every player is not a social welfare player. Therefore the following quantity m, which represents the "minimal possible punishment", is well defined:

$$\underline{m} = \max\{\gamma^j_{|h}(0^{-i}, \mu^i) < 0 \mid i, j \in I,\ \mu^i \neq 0^i \ \text{ a pure C-periodic strategy},\ h \in H\}. \tag{4}$$

Define

$$B = \min\left\{t \in \mathbf{N}\ \middle|(1-\delta)^t \cdot 1 + (1-(1-\delta)^t) \cdot m < 0\right\}. \tag{5}$$

Suppose $\gamma^j(0^{-i}, \mu^i) < 0$, where $\mu^i \neq 0^i$ is a pure C-periodic strategy. In particular $\gamma^j(0^{-i}, \mu^i) \leq \underline{m}$. If during $B \cdot C$ stages, all players other than i continue, and if i follows μ^i, player j's expected payoff if he continues whenever chosen is at most

$$(1-\delta)^B \cdot 1 + (1-(1-\delta)^B) \cdot \underline{m},$$

and it is, therefore, negative (The quantity 1 is an upper bound on the continuation payoff after stage $B \cdot C$).

3.3.1 Proof of Proposition 3.6

Assume A.1 and A.2 hold. We simultaneously construct by induction the four sequences as needed in the proposition. As mentioned before, the sequence $(v_k)_{k\in\mathbf{N}}$ will also satisfy that for every $k \in \mathbf{N}$ there is $i \in I$ such that $v^i_k < 0$.

Let $l_0 = \tau_0 = 1$ and $v_0 = a_\infty$; that is, $\Gamma(l_0, \tau_0, v_0)$ is a degenerate game, such that the game terminates at the same stage it starts and the players do not have the opportunity to act. In that game $\sigma_0 = \emptyset$ is a subgame perfect 0-equilibrium, so that

$\gamma(\sigma_0; \Gamma(l_0, \tau_0, v_0)) = a_\infty$. By Assumption A.1 there is at least one player $i \in I$ whose expected payoff is negative.

Assume we already defined the first k elements of each sequence as required in the proposition. We also assume that $\gamma(\sigma_k; \Gamma(l_k, \tau_k, v_k)) \not\geq \overrightarrow{0}$.

Define $v_{k+1} := \gamma(\sigma_k; \Gamma(l_k, \tau_k, v_k))$, as required by Condition C.2. We now define l_{k+1}, τ_{k+1} and σ_{k+1}.

Since $v_{k+1} \not\geq \overrightarrow{0}$, there is a player i such that $v^i_{k+1} < 0$. Furthermore, the maximal payoff every player can gain by terminating the game by himself is 0. In particular, there is $1 \leq t \leq C$ where $a^i_i(t) = 0 > v^i_{k+1}$.

Let $t_* \in \{1, \ldots, l_k - 1\}$ be the maximal integer which is strictly smaller than l_k such that there is some player i who satisfies

$$a^i_i(t_*) > v^i_{k+1}. \tag{6}$$

If there is no t_* as required above, let t_* be the maximal integer in $\{l_k, \ldots, C\}$ such that there is some player i who satisfies (6). In the example of Sect. 2.3, if $v^i_{k+1} = a_\infty$ and $l_k = 1$ then $t_* = 2$. Denote by I_* the set of all players who satisfy (6). Consider the game that starts at stage t_* and terminates at the earliest stage t such that $(t \mod C) = l_k$ with payoff v_{k+1}. The game has a subgame 0-equilibrium in which all the player in I_* stop at stage t_*, and all the other players continue whenever chosen. The equilibrium payoff is v_* as follows:

$$v_* := \sum_{i \in I_*} p_i\,(t_*)\, a_i\,(t_*) + \left(1 - \sum_{i \in I_*} p_i\,(t_*)\right) v_{k+1}.$$

Case 1: $v_* \not\geq \overrightarrow{0}$.
Define

$$\begin{aligned}
&l_{k+1} := t_*.\\
&\tau_{k+1} := \min\{t \geq t_* : (t \mod C) = l_k\}.\\
&\sigma_{k+1} := \text{at stage } t_*, \text{ all players in } I_* \text{ stop, all the other players continue.}
\end{aligned}$$

Case 2: $v_* \geq \overrightarrow{0}$.
Choose a player, i_* in I_*. Then $v^{i_*}_{k+1} < 0$. For every $r \geq 0$, define an auxiliary finite-stage game $\Gamma(r)$ which starts at stage t_* and terminates at stage $t_* + r \cdot C$ with terminal payoff v_*, except for, whenever player i_* is chosen at stage t such that $(t \mod C) = t_*$ and chooses to continue, the game terminates at once with payoff v_{k+1}. For every $r \geq 0$, the game $\Gamma(r)$ has an equilibrium in which at least i_* stops whenever he is chosen at stage t where $(t \mod C) = t_*$. Furthermore, for a sufficient small r (at least for $r = 0$ and $r = 1$), all the players except i_* continue whenever chosen, so the equilibrium payoff is as follows:

$$a_*(r) = \left(1 - \left(1 - p_{i_*}(t_*)\right)^r\right) a_{i_*}\,(t_*) + \left(1 - p_{i_*}\,(t_*)\right)^r v_*.$$

For $r = 0$, $a_*(r) = v_* \geq \overrightarrow{0}$. For every $r \geq B$, $a_*(r) \ngeq \overrightarrow{0}$, by Eq. (5), however, $a_*(r)$ is not necessarily an equilibrium payoff in $\Gamma(r)$. Let r_* be the minimal integer in $\{1, \ldots, B\}$ such that $a_*(r_*) \ngeq \overrightarrow{0}$. Notice, for every $r < r_*$, $a_*(r) \geq \overrightarrow{0}$, hence, all the players, except i_*, prefer to continue whenever chosen in $\Gamma(r_*)$, and $a_*(r_*)$ is an equilibrium payoff.

Define
$l_{k+1} := t_*$.
τ_{k+1} := During $r_* \cdot C$ stages, whenever player i_* is chosen at stage t such that $(t \mod C) = t_*$, and chooses to continue, the game terminates at stage $\widetilde{t} = \min\{t' > t : (t' \mod C) = l_k\}$. If the game does not terminate before the end of those $r_* \cdot C$ stages, it terminates at stage $\widehat{t} = \min\{t' > t_* + r_* \cdot C : (t' \mod C) = l_k\}$. σ_{k+1} :=

(a) player i_* stops at every stage t such that $t_* = (t \mod C)$.
(b) at stage $r_* \cdot C + 1$ all players in I_* stop.
(c) otherwise, the chosen player continues.

The construction implies that Conditions C.1–C.3 are satisfied. It is left to prove that Condition C.4 is satisfied as well.

Assume to the contrary that Condition C.4 does not hold—there is some player i and an integer k_0 such that for every $k > k_0$, the auxiliary finite-stage game $\Gamma(l_k, \tau_k, v_k)$, under the profile σ_k, terminates by player i with probability at least δ, while all the other players always continue.

If for every $k > k_0$ player i is the only one who stops in $\Gamma(l_k, \tau_k, v_k)$ then $v_k^i < 0$ for every $k > k_0$. It might happen that player i will stop even if he receives a negative payoff. However, such a case happens only finitely many times: if only player i stops, and if $v_k^i < 0$ for every $k > k_0$ then player i always stops whenever $a_i^i(t \mod C) = 0$. Therefore, for every $k > k_0$, v_k^i is a weighted average of $v_{k_0}^i$, 0 and possibly negative payoffs that player i may find more worthwhile than the continuation payoffs. Since the weight of 0 in that average increases with k, v_k^i increases with k as well. Thus, there is a bounded $k_1 > k_0$ such that at every finite game $\Gamma(l_k, \tau_k, v_k)$ where $k > k_1$, player i stops only at stages where his payoff is 0.

Since i is not a social welfare player, there are necessarily some k_2 such that $k_1 < k_2 < k_1 + B$, and a player $j \neq i$ such that $v_k^j < 0$, for every $k > k_2$. Therefore j is better off by deviating and terminating one of the games $\Gamma(l_k, \tau_k, v_k)$, for some $k > k_2$, a contradiction. □

3.3.2 Proof of Proposition 3.7

Assume there are sequences $(l_k)_{k\in\mathbf{N}}$, $(\tau_k)_{k\in\mathbf{N}}$, $(v_k)_{k\in\mathbf{N}}$ and $(\sigma_k)_{k\in\mathbf{N}}$ in pure strategies such that Conditions C.1–C.4 hold.

The play consists of the sequence of chosen players $(i_n)_{n=1}^{\infty}$ and the terminal stage $\theta \in \mathbf{N} \cup \{\infty\}$. For every two stopping times $\widehat{\tau}_1$ and $\widehat{\tau}_2$ define

$$(\widehat{\tau}_1 \oplus \widehat{\tau}_2)(i_1, i_2, \ldots, \theta) = \widehat{\tau}_1(i_1, i_2, \ldots, \theta) + \widehat{\tau}_2(i_{\widehat{\tau}_1+1}, i_{\widehat{\tau}_1+2}, \ldots, \widehat{\theta})$$

where

$$\widehat{\theta} = \max\{1, \theta - \widehat{\tau}_1(i_1, i_2, \ldots, \theta)\}.$$

In words, at stage $\widehat{\tau_1}$ forget past play, and stop $\widehat{\tau_2}$ stages later (where the history to use to calculate $\widehat{\tau_2}$ is the partial history after stage $\widehat{\tau_1}$). Let

$$\tau_k^* = \tau_k \oplus \tau_{k-1} \oplus \cdots \oplus \tau_1.$$

By Condition C.1 the auxiliary finite-stage game $\Gamma(l_k, \tau_k^*, v_0)$ is the concatenation of the games $\Gamma(l_k, \tau_k, v_k)$, $\Gamma(l_{k-1}, \tau_{k-1}, v_{k-1}), \ldots, \Gamma(l_0, \tau_0, v_0)$.

One can verify that, by C.1 and C.2, the following profile σ_k^* is a subgame perfect 0-equilibrium in $\Gamma(l_k, \tau_k^*, v_0)$: first follow σ_k, then σ_{k-1} and so on until σ_1.

By Condition C.3, the probability that under σ_k^* the game terminates before stage τ_k^* is at least $1 - (1 - \delta)^k$.

By Condition C.4, for every player $i \in I$ and for a large enough k, the game $\Gamma(l_k, \tau_k^*, v_0)$ terminates under $(\sigma_k^{*,-i}, 0^i)$ before stage τ_k^* with probability that is as close to 1 as we wish, as k goes to infinity.

The profile σ_k^* is an element in the space $[0, 1]^{n_k}$ for some $n_k \in \mathbf{N}$.σ_k^* can be identified with a point in $[0, 1]^{\mathbf{N}}$, by identifying the n_k initial coordinates with σ_k^* and the other coordinates with 0.

The space $[0, 1]^{\mathbf{N}}$ is compact, therefore $(\sigma_k^*)_{kt\in\mathbf{N}}$ has a subsequence which converges to the limit σ^*.

Using a limiting argument, it is standard to prove that σ^* is a subgame perfect 0-equilibrium (cf. Solan and Vieille 2003). □

References

Dynkin EB (1969) Game variant of a problem on optimal stopping. Sov Math Dokl 10:270–274

Flesch J, Thuijsman F, Vrieze OJ (2007) Stochastic games with additive transitions. Eur J Oper Res 179:483–497

Flesch J, Scheonmakers G, Vrieze K (2008) Stochastic games on a product state space. Math Oper Res 33(2):403–420

Kiefer YI (1971) Optimal stopped games. Theory Probab Appl 16:185–189

Mertens JF (1987) Repeated games. In: Gleason AM (ed) Proceedings of the international congress of mathematicians, Berkeley, 1986. American Mathematical Society, pp 1528–1577

Mertens JF, Parthasarathy T (1987) Equilibria for discounted stochastic games. CORE Discussion Paper 8750

Neveu J (1975) Discrete-parameter martingales. North-Holland, Amsterdam

Rosenberg D, Solan E, Vieille N (2001) Stopping games with randomized strategies. Probab Theory Relat Fields 119:433–451

Rosenberg D, Solan E, Vieille N (2002) Blackwell optimality in markov decision processes with partial observation. Ann Stat 30:1178–1193

Shmaya E, Solan E (2004) Two player non zero-sum stopping games in discrete time. Ann Probab 32:2733–2764

Solan E (1998) Discounted stochastic games. Math Oper Res 23:1010–1021

Solan E (2001) The dynamics of the nash correspondence and n-player stochastic games. Int Game Theory Rev 3:291–300

Solan E (2005) Subgame-Perfection in quitting games with perfect information and differential equations. Math Oper Res 30:51–72

Solan E, Vieille N (2003) Deterministic multi-player Dynkin games. J Math Econ 1097:1–19