



ELSEVIER

Available online at www.sciencedirect.com

Journal of Mathematical Economics 1097 (2003) 1–19

**JOURNAL OF
Mathematical
ECONOMICS**

www.elsevier.com/locate/jmateco

Deterministic multi-player Dynkin games

Eilon Solan ^{a,b,*}, Nicolas Vieille ^c

^a *MEDS Department, Kellogg School of Management, Northwestern University,
2001 Sheridan Road, Evanston, IL 60208, USA*

^b *School of Mathematical Sciences, Tel Aviv University, Tel Aviv, Israel*

^c *Département Finance et Economie, HEC, 1 rue de la Libération, 78 351 Jouy-en-Josas, France*

Received 3 September 2002; received in revised form 22 January 2003; accepted 28 January 2003

Abstract

A *multi-player Dynkin game* is a sequential game in which at every stage one of the players is chosen, and that player can decide whether to continue the game or to stop it, in which case all players receive some terminal payoff.

We study a variant of this model, 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.

We prove that a subgame-perfect ϵ -equilibrium in Markovian strategies exists. If the game is not degenerate this ϵ -equilibrium is actually in pure strategies.

© 2003 Elsevier Science B.V. All rights reserved.

JEL classification: C73

Keywords: Dynkin game; Stochastic game; Equilibrium; n -player games

1. Introduction

Dynkin (1969) introduced the following zero-sum game of optimal stopping. The game involves two players, and two stochastic processes: $(i_n)_{n \in \mathbf{N}}$ is a $\{1, 2\}$ -valued process, which indicates which player is active at stage n , and $(r_n)_{n \in \mathbf{N}}$ is a \mathbf{R}^2 -valued process, which indicates the terminal payoff.

At every stage n , the two players are informed of past and current values of the two processes. Player i_n , the *active* player at stage n , decides whether he continues or stops. The game stops at the first stage θ in which the active player chooses to stop. The payoff (paid

* Corresponding author.

E-mail addresses: e-solan@kellogg.northwestern.edu, eilons@post.tau.ac.il (E. Solan), vieille@hec.fr (N. Vieille).

by player 2 to player 1) is r_θ if $\theta < +\infty$ and zero otherwise. A *pure* strategy of player i is a stopping time that is consistent with the rules of the game.

Dynkin proved that this game has a value if $\sup_{n \in \mathbf{N}} |r_n| \in L^1$, and constructed pure ϵ -optimal strategies for the two players. Dynkin's ϵ -optimal strategies are subgame-perfect in the sense that after every finite history, the continuation strategy is ϵ -optimal in the subgame defined by that history.

An extensive literature developed from this seminal work. In a discrete time framework, much attention was paid to the case where the players are allowed to stop simultaneously. In the zero-sum case, several authors, including Kiefer (1971) and Neveu (1975), provided sufficient conditions for the existence of the value, when players are restricted to stopping times. Rosenberg et al. (2001) proved (under a minimal boundedness condition) that the value always exists, provided the players are allowed to use *randomized* stopping times. In the two-player non-zero-sum case, Shmaya and Solan (2002) proved that an ϵ -equilibrium always exists in randomized stopping times (again, under some boundedness condition).

Dynkin's (1969) result implies that in every *multi-player* Dynkin game (without simultaneous moves) an ϵ -equilibrium exists. Indeed, let σ^i be a pure ϵ -optimal strategy of player i in the zero-sum game in which player i maximizes his expected payoff, and all other players try to minimize player i 's payoff. Let σ_i^{-i} be a pure ϵ -optimal strategy of i 's opponents in this game. One can verify that the strategy profile in which each player i follows σ^i until a deviation occurs (since each σ^i is pure, a deviation is detected immediately), and upon deviation of player j all his opponents switch to σ_j^{-j} , is a 2ϵ -equilibrium.

The model of multi-player Dynkin games offers a stylized framework to analyze various issues of timing games. For example, in situations of shrinking markets (see, e.g. Ghemawat and Nalebuff, 1985; Fine and Li, 1989), n firms have to decide when to exit a shrinking market. Once a firm exits, we remain with a market with $n - 1$ firms, which can be solved inductively, hence the overall game reduces to a Dynkin game.

A similar situation occurs in takeover games, where n firms strategically decide to make a takeover attempt on opponent firms.

Another related model is that of multi-player duels, or n -uels (see, e.g. Kilgour, 1975, 1977 or Kilgour and Brams, 1997). In this model, n gunners alternately have the option to shoot one of their opponents or to abstain. Since once a gunner hits one of his opponents we are left with a game with $n - 1$ players, which can be solved inductively, the game is essentially reduced to a Dynkin game where players have *several* stop actions.

As the ϵ -equilibrium we presented above involves threats of punishment, which might be non-credible, it is desirable to know whether a *subgame-perfect* ϵ -equilibrium exists for every $\epsilon > 0$. To this day, it is still not known whether every multi-player Dynkin game has an ϵ -equilibrium.

When $|I| = 2$, the proof of Shmaya and Solan (2002) can be used to show the existence of a subgame-perfect ϵ -equilibrium. Solan (2002) uses the theory of differential inclusions to prove the existence of a subgame-perfect ϵ -equilibrium when (i) the sequence (i_n) is i.i.d., and (ii) r_n depends only on i_n (so that the terminal payoff depends only on the identity of the player who terminates the game).

In the present paper, we analyze the following class of I -player games. A deterministic sequence $(i_n, p_n, r_n) \in I \times [0, 1] \times \mathbf{R}^I$ is given. At each stage n , player i_n chooses whether

to continue or to stop. If he continues, the game continues to the next stage, while if he stops a lottery is performed. With probability p_n the game terminates, yielding the payoff r_n , while with probability $1 - p_n$ the game continues.

The assumption that the order of players is deterministic is restrictive but sometimes relevant. On the other hand, allowing the probability of termination to be strictly less than 1 is quite natural: a takeover attempt is not always successful, and the accuracy of a gunner is not always perfect.

Our main result states that if the sequence (r_n) of payoffs is bounded, a subgame-perfect ε -equilibrium in Markovian strategies exists. Moreover, unless the game is degenerate, this ε -equilibrium is in pure strategies. However, in degenerate cases, a subgame-perfect 0-equilibrium need not exist. Since the subgame-perfect ε -equilibrium we identify is in Markovian strategies, it is robust to the information players receive along the game; all they need to know is the stage of the game. Translated to the n -uel model, this means that there is a subgame-perfect ε -equilibrium which is also a subgame-perfect ε -equilibrium in the silent n -uel, in which players do not observe missed shots.

In degenerate cases, there need not be subgame-perfect ε -equilibria in pure strategies. This is to be contrasted with (i) finite games of perfect information and (ii) two-player zero-sum Dynkin games, where a subgame perfect (ε -) equilibrium in pure strategies always exists.

We hope that the combination of the arguments we use here with the techniques presented by Shmaya and Solan (2002) and Solan (2002) can be used to further study multi-player Dynkin games.

Another motivation to our study is linked to the observation that deterministic Dynkin games form a simple class of stochastic games. By now, some results are available on the existence of equilibrium payoffs in multi-player stochastic games, see Solan (1999) and Vieille (2000). By contrast, apart from few classes of games, there are no results on the existence of subgame-perfect equilibrium payoffs and useful techniques are yet to be found. We hope that this paper will contribute to this emerging literature.

The paper is arranged as follows. In Section 2 we present the model and the main result. Several examples appear in Section 3. The proof of the main result appears in Section 4.

2. The model and the main result

2.1. Deterministic multi-player Dynkin games

A deterministic multi-player Dynkin game $\Gamma = (I, (i_n, p_n, r_n)_{n \in \mathbf{N}})$ is given by

- a finite set I of players;
- for every $n \in \mathbf{N}$, a triplet $(i_n, p_n, r_n) \in I \times [0, 1] \times \mathbf{R}^I$.

The triplet (i_n, p_n, r_n) specifies who is allowed to stop at stage n , the probability that the game terminates if player i_n decides to stop, and the terminal payoff if the game terminates at stage n , respectively.

The game is played in stages. At each stage $n \in \mathbf{N}$, provided the game has not terminated yet, player i_n has to choose whether to *Continue* or *Stop*. If he decides to continue, the game continues to stage $n + 1$. If he decides to stop, a lottery takes place (all lotteries in

the game, including random choices by the players, are independent). With probability p_n the game terminates, and the terminal payoff for the players is given by the vector r_n . With probability $1 - p_n$ the game continues to stage $n + 1$. If the game never terminates, the payoff is zero for all players.¹

We denote by θ the termination stage of the game, i.e. the first stage in which a player decides to stop and the game terminates. Thus, the payoff to player $i \in I$ is $r_\theta^i \mathbf{1}_{\theta < \infty}$.

2.2. Strategies and results

A strategy of player $i \in I$ maps the set of information sets of player i to the set of mixed moves of player i . We let $N_i = \{n \in \mathbf{N} | i_n = i\}$ be the set of stages in which player i is active.

We are going to restrict the players to Markovian strategies; namely, strategies that depend only on the stage, and not on the history. We will prove below that the game admits a subgame-perfect ε -equilibrium in Markovian strategies. By a general observation (see, e.g. Fudenberg and Tirole (1991, p. 501)), this subgame-perfect ε -equilibrium remains a subgame-perfect ε -equilibrium without the restriction to Markovian strategies.

In the present context, a (behavior Markovian) strategy of player i is a function $\sigma^i: N_i \rightarrow [0, 1]$, where $\sigma^i(n)$ is the probability assigned by player i to stop at stage n , provided the game does not terminate before that stage. We denote the set of strategies of player i by Σ^i .

A strategy profile (or simply a profile) is a vector $\sigma = (\sigma^i)_{i \in I}$ of strategies, one for each player.

Every strategy profile $\sigma \in \times \Sigma_{i \in I}^i$ induces a probability distribution \mathbf{P}_σ over the space of plays, or infinite histories. The corresponding expectation operator is \mathbf{E}_σ . Thus, the expected payoff to player i given a strategy profile σ is

$$\gamma^i(\sigma) := \mathbf{E}_\sigma[r_\theta^i \mathbf{1}_{\theta < \infty}].$$

Before we state our result, we first recall standard equilibrium notions.

Definition 1. Let $\varepsilon \geq 0$. A strategy profile σ is an ε -equilibrium if for every player $i \in I$ and every strategy $\tau^i \in \Sigma^i$,

$$\gamma^i(\sigma) \geq \gamma^i(\sigma^{-i}, \tau^i) - \varepsilon.$$

We mention that, for any $\varepsilon' > \varepsilon$, an ε -equilibrium is a uniform ε' -equilibrium; that is, it is an ε' -equilibrium (a) in every discounted game, provided the discount factor is sufficiently small, and (b) in every N -stage game, provided N is sufficiently large. Indeed, the proof provided in Solan and Vieille (2001, Proposition 2.13) does adapt to the present framework.

For $n \in \mathbf{N}$, we denote by $\gamma_n(\sigma)$ the expected payoff induced by the strategy profile σ in the subgame starting at stage n .

A strategy profile is a subgame-perfect (ε -) equilibrium of a game if it induces an (ε -) equilibrium in any subgame. In the present context, this amounts to the following definition.

¹ Equivalently, we may assume that, with probability p_n , player i_n is given the opportunity to stop for sure. For each strategy profile, the payoff is the same under both interpretations of the game.

Definition 2. Let $\varepsilon \geq 0$. A strategy profile σ is a *subgame-perfect ε -equilibrium* if for every $n \in \mathbf{N}$, every player $i \in I$, and every $\tau^i \in \Sigma^i$,

$$\gamma_n^i(\sigma) \geq \gamma_n^i(\sigma^{-i}, \tau^i) - \varepsilon.$$

Our main result is the following.

Theorem 1. Let $\Gamma = (I, (i_n, p_n, r_n)_{n \in \mathbf{N}})$ be a deterministic Dynkin game. If the sequence $(r_n)_{n \in \mathbf{N}}$ is bounded, then for every $\varepsilon > 0$ the game Γ admits a subgame-perfect ε -equilibrium in Markovian strategies.

We conclude this section with two comments.

As will be clear from the proof, in most cases, there is a *pure* subgame-perfect ε -equilibrium. However, this is not always true (see Example 3 below). This is in sharp contrast with *finite* extensive games of perfect information and with two-player zero-sum Dynkin games.

Our proof is valid as long as $\gamma(\sigma)$ is uniformly bounded, for every profile σ (which is the case when the sequence $(r_n)_{n \in \mathbf{N}}$ is bounded). If this does not hold, there are strategies σ such that the corresponding payoff for at least one player is infinite, so that the payoff function of the game is not well-defined.

3. Examples

In the present section we provide several examples that illustrate the main features of the model.

Example 1. Take $I = \{1, 2, 3\}$ and

$$(i_n, p_n, r_n) = \begin{cases} (1, 1, (1, 0, 3)), & n = 1 \text{ modulo } 3, \\ (2, 1, (3, 1, 0)), & n = 2 \text{ modulo } 3, \\ (3, 1, (0, 3, 1)), & n = 0 \text{ modulo } 3. \end{cases}$$

In words, at the first stage, player 1 can stop the game, thereby yielding the payoff vector $(1, 0, 3)$. If player 1 chooses to continue, at the second stage player 2 can stop the game, yielding the terminal payoff $(3, 1, 0)$. If player 2 chooses to continue as well, at the third stage player 3 can stop the game, yielding the terminal payoff $(0, 3, 1)$. The process then repeats itself cyclically. This game is a variation upon a game studied by Flesch et al. (1997).

We will characterize all pure subgame-perfect 0-equilibrium profiles of that game, using backward induction.

Let σ be such a 0-equilibrium. Assume that at stage $3n$, for some $n \geq 2$, player 3 stops with probability 1; that is, $\sigma^3(3n) = 1$. In particular, $\gamma_{3n}(\sigma) = (0, 3, 1)$.

Consider the subgame starting at stage $3n - 1$. In that subgame, player 2 receives $\gamma_{3n}^2(\sigma) = 3$ if he chooses to continue at stage $3n - 1$, while he receives only 1 if he chooses to stop. By the subgame-perfect equilibrium condition, player 2 continues at stage $3n - 1$, that is, $\sigma^2(3n - 1) = 0$. Hence, $\gamma_{3n-1}(\sigma) = \gamma_{3n}(\sigma) = (0, 3, 1)$.

We repeat this argument with the subgame starting at stage $3n - 2$. By continuing at stage $3n - 2$ player 1 receives 0, as the game will be terminated at stage $3n$, while by stopping he receives 1. By the subgame-perfect equilibrium condition, $\sigma^1(3n - 2) = 1$ and $\gamma_{3n-2}(\sigma) = (1, 0, 3)$.

Applying this backward induction argument repeatedly, we get that $\sigma^3(3n - 3) = 0$, $\sigma^2(3n - 4) = 1$, $\sigma^1(3n - 5) = 0$ and $\sigma^3(3n - 6) = 1$. The cycle of length 6 then repeats itself.

On the other hand, if $\sigma^3(3n) = 0$ for some $n \geq 2$, then $\sigma^3(3n - 3) = 1$ and the previous analysis holds.

Thus, there are two pure subgame-perfect 0-equilibria: (a) at odd stages the active player stops, and at even stages the active player continues, and (b) at even stages the active player stops, and at odd stages the active player continues.

In each pure equilibrium, the players agree on who shoots first. We believe that the interpretation of these two equilibria is quite appealing: suppose there are three gunners. If gunner 1 thinks that gunner 2 is going to shoot tomorrow gunner 3 (or gunner 1 if gunner 3 is already dead), he has no reason to shoot today: he is better off by letting gunner 2 be done with gunner 3, and shoot gunner 2 the next time he can. On the other hand, if gunner 1 thinks that gunner 2 is not going to shoot tomorrow if gunner 3 is still alive, but shoot gunner 1 if gunner 3 is already dead, and that gunner 3 is going to shoot him the day after, he is indifferent between shooting and not shooting gunner 2, as he is going to die anyway, so he can as well shoot gunner 2 today.

Remark 1. This game admits other subgame-perfect equilibria. In particular, the profile in which each player stops with probability $1/2$ whenever active, is a subgame-perfect equilibrium. In a sense, it corresponds to the cyclic equilibrium constructed by [Flesch et al. \(1997\)](#).

In the next example, we allow for probabilities of success below 1.

Example 2. Consider the following modification of [Example 1](#), where $I = \{1, 2, 3\}$, and

$$(i_n, p_n, r_n) = \begin{cases} (1, 1, (1, 0, 3)), & n = 1 \text{ modulo } 3, \\ (2, 1/2, (3, 1, 0)), & n = 2 \text{ modulo } 3, \\ (3, 1/2, (0, 3, 1)), & n = 0 \text{ modulo } 3. \end{cases}$$

Thus, when player 1 stops the game terminates with probability 1, while when either player 2 or player 3 stops the game terminates with probability $1/2$.

As we did in [Example 1](#), we characterize the set of subgame-perfect 0-equilibrium in pure strategies. Let σ be such a strategy profile. Let $n > 0$ and i be the active player at stage n . By the subgame-perfect equilibrium condition, $\sigma^i(n) = 1$ if $\gamma_{n+1}^i(\sigma) < 1$ and $\sigma^i(n) = 0$ if $\gamma_{n+1}^i(\sigma) > 1$.

Let $n \geq 3$, and assume that $\sigma^1(3n + 1) = 1$. Then $\gamma_{3n+1}(\sigma) = (1, 0, 3)$, and therefore $\sigma^3(3n) = 0$. This implies that $\gamma_{3n}(\sigma) = \gamma_{3n+1}(\sigma) = (1, 0, 3)$, and therefore $\sigma^2(3n - 1) = 1$.

It follows that

$$\gamma_{3n-1}(\sigma) = \frac{1}{2}(3, 1, 0) + \frac{1}{2}(1, 0, 3) = (2, \frac{1}{2}, \frac{3}{2}),$$

and therefore $\sigma^1(3n - 2) = \sigma^3(3n - 3) = 0$ and $\sigma^2(3n - 4) = 1$.

Then,

$$\gamma_{3n-4}(\sigma) = \frac{1}{2}(3, 1, 0) + \frac{1}{2}(2, \frac{1}{2}, \frac{3}{2}) = (\frac{5}{2}, \frac{3}{4}, \frac{3}{4}),$$

and therefore $\sigma^1(3n - 5) = 0$ and $\sigma^3(3n - 6) = 1$.

One therefore has

$$\gamma_{3n-6}(\sigma) = \frac{1}{2}(0, 3, 1) + \frac{1}{2}(\frac{5}{2}, \frac{3}{4}, \frac{3}{4}) = (\frac{5}{4}, \frac{15}{8}, \frac{7}{8}),$$

and therefore $\sigma^2(3n - 7) = 0$, $\sigma^1(3n - 8) = 0$ and $\sigma^3(3n - 9) = 1$.

Finally,

$$\gamma_{3n-9}(\sigma) = \frac{1}{2}(0, 3, 1) + \frac{1}{2}(\frac{5}{4}, \frac{15}{8}, \frac{7}{8}) = (\frac{5}{8}, \frac{39}{16}, \frac{15}{16}),$$

and therefore $\sigma^2(3n - 10) = 0$, and $\sigma^1(3n - 11) = 1$.

Therefore, any pure subgame-perfect 0-equilibrium must repeat the sequence of actions (starting with player 1) (S, C, S; C, C, S; C, S, C; C, S, C). Along this cycle, player 1 first stops, then player 3 stops twice in a row, then player 2 stops twice in a row. This difference with the subgame-perfect 0-equilibrium of [Example 1](#) arises since the probability of termination is here below 1. By further decreasing the probabilities p_n for $n = 2$ or 3 modulo 3, while keeping $p_n = 1$ for $n = 1$ modulo 3, one can create examples in which all pure subgame-perfect equilibria have cycles of arbitrary length.

This example highlights one effect of low values for (p_n) . Note indeed that the expected payoff, starting from some stage n , is a convex combination of r_n and of the continuation payoff (the expected payoff, starting from stage $n + 1$). The weight of r_n depends on the probability of termination, but cannot exceed p_n . In particular, when the probability of termination is low, the expected payoff is close to the continuation payoff. Therefore, if some player has an incentive to stop only once the continuation payoff reaches a certain threshold, many stages may be required so that this threshold is reached. Thus, if the game has a periodic equilibrium, lowering the probabilities of termination often results in periodic equilibria with longer and longer periods.

We next introduce a two-player game that has no subgame-perfect 0-equilibrium and no pure subgame-perfect ε -equilibrium.

Example 3. Take $I = \{1, 2\}$, and

$$(i_n, p_n, r_n) = \begin{cases} (1, 1, (-1, 2)), & n \text{ is odd,} \\ (2, 1, (-2, 1)), & n \text{ is even.} \end{cases}$$

Fix $\varepsilon \in (0, 1)$, and let σ be the strategy profile defined by $\sigma^1(2n+1) = 1$ and $\sigma^2(2n+2) = \varepsilon$ for every $n \geq 0$. We claim that σ is a subgame-perfect ε -equilibrium. One should verify that player 1 (respectively player 2) cannot profit by deviating in the subgames that start at odd (respectively even) stages. Consider first the subgame that starts at stage $2n + 1$, for some $n \geq 0$. By stopping at stage $2n + 1$ player 1 receives -1 , while, since player 2 eventually stops with probability 1, player 1's payoff is at most -1 , whatever he plays. In the subgame starting at stage $2n + 2$, player 2's expected payoff under σ is $\varepsilon + 2(1 - \varepsilon) = 2 - \varepsilon$, whereas the maximal payoff to player 2 in the game is 2.

We next prove that the game has no subgame-perfect ε -equilibrium in pure strategies. Assume to the contrary that there exists such a profile σ .

We first claim that there is an infinite set of even stages in which player 2 chooses to stop. Otherwise, let N be the maximal integer such that player 2 stops at stage $2N$ (set $N = 0$ if player 2 never stops). Consider now the subgame that starts at stage $2N + 2$. By the definition of N , player 2 never stops in this subgame. Since σ is a subgame-perfect ε -equilibrium, this implies that under σ player 1 never stops in this subgame: by never stopping he receives 0, while by stopping he receives -1 . But this leads to a contradiction, as it implies that player 2 can profit 1 by deviating: by never stopping he receives 0, while by stopping he receives 1.

We next claim that there is at most one even stage in which player 2 chooses to stop. Together with the previous paragraph, this shows that there cannot be a subgame-perfect ε -equilibrium. Assume that player 2 stops at stage $2N$, with $N > 1$. Since σ induces an ε -equilibrium in the subgame that starts at stage $2N - 1$, and since player 2 stops at stage $2N$, under σ player 1 stops at stage $2N - 1$. However, since player 1 stops at stage $2N - 1$, under σ player 2 continues in all stages $2k$ for $k < N$: by continuing in all these stages he receives 2, while his payoff upon stopping is 1.

This example shows that pure subgame-perfect ε -equilibria need not exist. Such a case may arise when there is a player i who by stopping gives everyone else high payoff, but he himself receives low payoff. It is then in the interest of his opponents to threaten him that if he does not stop, one of them will eventually stop and punish player i . The punisher, however, stops with low probability, so that player i has a chance to correct his behavior and stop the game at a later stage.

We finally prove that there is no subgame-perfect 0-equilibrium. We argue by contradiction, and we let σ be a subgame-perfect 0-equilibrium. For $i = 1, 2$, we denote by c^i the strategy that always continues, i.e. $c^i(n) = 0$ for each $n \in N_i$. Note first that, for each $n \in \mathbf{N}$, one has

$$P_\sigma(\theta < +\infty | \theta \geq n) = 1. \tag{1}$$

Indeed, the sequence $(P_\sigma(\theta < +\infty | \theta \geq n))_{n \in \mathbf{N}}$ would otherwise decrease to zero, hence the sequence $(\gamma_n(\sigma))_{n \in \mathbf{N}}$ would converge to zero, and player 2 would have a profitable deviation in the subgame starting at stage n , for n large enough. By (1) the game terminates with probability 1, hence at least one of the players eventually stops with probability 1:

$$P_{\sigma^1, c^2}(\theta < +\infty | \theta \geq n) = 1 \text{ for each } n \in \mathbf{N}, \tag{2}$$

or

$$P_{c^1, \sigma^2}(\theta < +\infty | \theta \geq n) = 1 \text{ for each } n \in \mathbf{N}. \tag{3}$$

If (2) holds, then c^2 is the best reply to σ^1 in all subgames, hence $\sigma^2 = c^2$. Since the unique best reply of player 1 to c^2 is c^1 , one gets $\sigma = (c^1, c^2)$ —a contradiction to (1).

If (3) holds, there are infinitely many even integers n such that $\sigma^2(n) > 0$. By optimality of σ^1 , and since (3) holds, one has $\sigma^1(n - 1) = 1$ for any such n . Therefore, (2) holds—a contradiction.

4. The Proof of Theorem 1

In the present section we prove Theorem 1.

4.1. Preliminaries

In this subsection, we analyze few degenerate cases, and slightly rephrase the problem. The core of the proof of Theorem 1 is in Section 4.4.

Let $\Gamma = (I, (i_n, p_n, r_n)_{n \in \mathbf{N}})$ be a deterministic Dynkin game. Since the sequence $(r_n)_{n \in \mathbf{N}}$ is bounded, we can assume w.l.o.g. that payoffs are bounded by 1.

Let $\tilde{\Gamma} = (I, (i_n, p_n, \tilde{r}_n)_{n \in \mathbf{N}})$ be another game with the same sequence of active players and the same probabilities of success. Since the payoff functions of the two games differ by at most $\sup_{n \in \mathbf{N}} \|r_n - \tilde{r}_n\|$, any subgame-perfect ε -equilibrium of $\tilde{\Gamma}$ is a subgame-perfect ε' -equilibrium of Γ , where $\varepsilon' = \varepsilon + \sup_{n \in \mathbf{N}} \|r_n - \tilde{r}_n\|$.

Since we are looking for an ε -equilibrium, and since payoffs are bounded, there is no loss of generality in assuming that the range of the sequence $(r_n)_{n \in \mathbf{N}}$ is finite, and that if (i, r) and (j, \tilde{r}) are two distinct elements in that range then $r^k \neq \tilde{r}^k$ for every $k \in I$.

Notice now that Theorem 1 will follow if we prove that there is a subgame-perfect ε -equilibrium in some subgame of Γ . Indeed, the conclusion for Γ will then follow by applying backward induction to the first stages of the game. Moreover, since finite extensive games with perfect information have pure subgame-perfect equilibria, the resulting profile will be pure when the subgame-perfect ε -equilibrium of the subgame is pure.

Let \mathbf{IR} be the finite range of the sequence $(i_n, r_n)_{n \in \mathbf{N}}$. For each $(i, r) \in \mathbf{IR}$ define

$$\pi(i, r) = \sum \{p_n | n \in \mathbf{N}, (i_n, r_n) = (i, r)\}$$

and set $\mathbf{IR}_\infty = \{(i, r) \in \mathbf{IR} | \pi(i, r) = +\infty\}$. If $\pi(i, r) = +\infty$ then if player i stops whenever $(i_n, r_n) = (i, r)$, and all players continue in all other stages, the game will eventually terminate, and the terminal payoff will be r .

We now argue that we may assume w.l.o.g. that

$$\pi(i, r) = +\infty \text{ for each } (i, r) \in \mathbf{IR}. \tag{4}$$

As a first step, we prove that we may assume w.l.o.g. that

$$\pi(i, r) = 0 \text{ for each } (i, r) \notin \mathbf{IR}_\infty. \tag{5}$$

Choose first $N \in \mathbf{N}$ large enough such that

$$\sum_{n \geq N: (i_n, r_n) = (i, r)} p_n < \varepsilon / |\mathbf{IR}| \text{ for each } (i, r) \notin \mathbf{IR}_\infty. \tag{6}$$

Such an N exists since \mathbf{IR} is a finite set. Denote by Γ_N the subgame that starts at stage N . Let $\tilde{\Gamma}_N = (I, (i_n, \tilde{p}_n, r_n)_{n \in \mathbf{N}})$ be the game that coincides with Γ_N except that $\tilde{p}_n = 0$ whenever $(\tilde{i}_n, \tilde{r}_n) \notin \mathbf{IR}_\infty$.

By (6), the payoff functions of the two games Γ_N and $\tilde{\Gamma}_N$ differ by at most 2ε . Therefore, any subgame-perfect ε -equilibrium of $\tilde{\Gamma}_N$ is a subgame-perfect 3ε -equilibrium of Γ_N , and, by backward induction, yields a subgame-perfect 3ε -equilibrium of Γ .

As the game $\tilde{\Gamma}_N$ satisfies (5), one can assume w.l.o.g. that (5) holds.

Thus, we are led to analyze games such that, for each $(i, r) \in \mathbb{IR}$, either $\pi(i, r) = 0$ or $\pi(i, r) = +\infty$. If $\pi(i, r) = 0$ for each $(i, r) \in \mathbb{IR}$ (so that $\mathbb{IR}_\infty = \emptyset$), the payoff function of the game is identically zero and the conclusion of Theorem 1 follows trivially.

Assume now that $\mathbb{IR}_\infty \neq \emptyset$. Consider the game obtained by dropping all stages n such that $\pi(i_n, r_n) = 0$ (and by relabeling stages). Since there are infinitely many stages n such that $\pi(i_n, r_n) = +\infty$, the resulting game is again a deterministic multi-player Dynkin game. Plainly, any subgame-perfect ε -equilibrium of this new game is also a subgame-perfect ε -equilibrium of the initial game (with the proper identification of stages, and with an arbitrary behavior in the stages that have been dropped).

It follows that we can assume w.l.o.g. that (4) holds.

4.2. A partition into blocks

In the present section we fix $\varepsilon < 1/40$. Given ε , we define a partition of the set \mathbf{N} of stages into blocks. This partition will be used in the sequel to prove Theorem 1.

We will use the following technical result.

Lemma 1 (Rosenberg et al., 2002, Lemma 18). *Let $n \in \mathbf{N}$, and let p_1, \dots, p_n be non-negative reals that satisfy $\sum_{i=1}^n p_i < 1/20$. Then,*

$$\sum_{i=1}^n p_i - 20 \left(\sum_{i=1}^n p_i \right)^2 \leq \sum_{i=1}^n p_i \prod_{j=1}^{i-1} (1 - p_j) \leq \sum_{i=1}^n p_i.$$

Observe that $\sum_{i=1}^n p_i \prod_{j=1}^{i-1} (1 - p_j)$ is the probability that the result of at least one out of n coins with parameters p_1, \dots, p_n is Head. In particular it is equal to $1 - \prod_{i=1}^n (1 - p_i)$.

Corollary 1. *Let $\varepsilon < 1/40$, $n \in \mathbf{N}$, and p_1, \dots, p_n be non-negative reals that satisfy $\sum_{i=1}^n p_i \geq \varepsilon$. Then $\sum_{i=1}^n p_i \prod_{j=1}^{i-1} (1 - p_j) \geq \varepsilon/2$.*

Proof. The proof is divided into three cases.

If $\sum_{i=1}^n p_i \leq 1/20$ the claim follows from Lemma 1 and since $\varepsilon < 1/40$.

If there is i such that $p_i \geq \varepsilon$ the claim holds trivially.

Otherwise, there is a subset $I \subset \{1, \dots, n\}$ such that $1/20 - \varepsilon \leq \sum_{i \in I} p_i \leq 1/20$. Then

$$\begin{aligned} \sum_{i=1}^n p_i \prod_{j=1}^{i-1} (1 - p_j) &= 1 - \prod_{i=1}^n (1 - p_i) \geq 1 - \prod_{i \in I} (1 - p_i) \\ &\geq \frac{1}{20} - \varepsilon - 20 \left(\frac{1}{20} - \varepsilon \right)^2 \geq \varepsilon - 20\varepsilon^2 \geq \varepsilon/2, \end{aligned}$$

where the second inequality follows from Lemma 1, and the third one holds since the function $x - 20x^2$ is monotonic decreasing for $x < 1/20$ and since $\varepsilon < 1/40$. \square

We are now ready to define the partition of \mathbf{N} into blocks. Set $n_0 = 1$ and, for $l \in \mathbf{N}$, define the initial stage n_l of block l to be

$$n_l = \min \left\{ n > n_{l-1} \mid \sum_{n_{l-1} \leq k < n, (i_k, r_k) = (i, r)} p_k \geq \epsilon \quad \forall (i, r) \in \mathbf{IR} \right\}.$$

Since $\pi(i, r) = +\infty$ for each $(i, r) \in \mathbf{IR}$, all $n_l, l \in \mathbf{N}$, are finite.

By [Corollary 1](#), in each block *all* players have a probability at least $\epsilon/2$ to terminate the game with any vector they choose.

4.3. A simple case

Under the assumption that $\pi(i, r) = +\infty$ for each $(i, r) \in \mathbf{IR}$, the proof proceeds by induction over the number of elements in \mathbf{IR} . The conclusion is easy if $|\mathbf{IR}| = 1$, and is left to the reader.

We now analyze a somewhat degenerate case that generalizes [Example 2](#). This is the only place in the proof where we use the induction hypothesis.

Lemma 2. *Assume that there exists $(i, r) \in \mathbf{IR}$ such that*

$$r^j \geq \tilde{r}^j \text{ for every } (j, \tilde{r}) \in \mathbf{IR}.$$

Then for each $\epsilon > 0$ there is a subgame-perfect ϵ -equilibrium.

The lemma states that if there is a terminal payoff r that is preferred by each player i to all terminal payoffs i controls, then a subgame-perfect ϵ -equilibrium exists.

Proof. We assume w.l.o.g. that $\epsilon < 1/40$, and we split the discussion into three cases.

Case 1. $r^i \geq 0$.

Let σ be the pure strategy profile in which player i stops whenever $(i_n, r_n) = (i, r)$, and all players continue in all other stages, i.e.

$$\sigma^{i_n}(n) = 1 \text{ if and only if } (i_n, r_n) = (i, r).$$

Fix $n \in \mathbf{N}$. We prove that σ induces a 0-equilibrium in the subgame that starts at stage n . Since $\pi(i, r) = +\infty$, the game eventually terminates, and therefore the expected payoff is r . Player i cannot gain by deviating, since his payoff is at most r^i if he terminates the game, and $0 \leq r^i$ if he always continues. Every player $j \neq i$ cannot gain by deviating either, since his payoff under σ is r^j , while if he deviates his payoff is in the convex hull of r^j and $\{\tilde{r}^j, (j, \tilde{r}) \in \mathbf{IR}\}$, hence at most r^j .

Case 2. $r^i < 0$, and there is $(j, \tilde{r}) \in \mathbf{IR}$ such that $i \neq j$ and $\tilde{r}^j < r^i$.

In this case, we elaborate upon the construction in [Example 3](#). We will have player i stop at all stages in $\{n \in \mathbf{N}: (i_n, r_n) = (i, r)\}$, and player j stop with some small probability at stages in $\{n \in \mathbf{N}: (i_n, r_n) = (j, \tilde{r})\}$. The choices of the corresponding probabilities should

fulfill two conditions: (i) these values should be small enough so that $\|\gamma_n(\sigma) - r\| < \varepsilon$, for each $n \in \mathbf{N}$ and (ii) they should be high enough so that, if player i were to continue at all stages, the game would still stop a.s. in finite time. These two conditions relate to the two dual aspects of the threat. By condition (i) the threat will be used on the equilibrium path with small probability. By condition (ii) it will provide incentives to player i to act as required.

Recall the partition of \mathbf{N} into blocks that was defined in Section 4.2, and that n_l is the first stage of block l , $l \geq 0$.

Since for every $l \geq 0$ one has $\sum_{n_l \leq n < n_{l+1}: (i_n, r_n) = (j, \tilde{r})} p_n \geq \varepsilon$, there is a function $x: \mathbf{N} \rightarrow [0, 1]$ such that for every $l \geq 0$ one has $\sum_{n_l \leq n < n_{l+1}: (i_n, r_n) = (j, \tilde{r})} x_n p_n = \varepsilon^2$.

We let σ be the strategy profile in which player i stops whenever $(i_n, r_n) = (i, r)$, player j stops with probability x_n whenever $(i_n, r_n) = (j, \tilde{r})$, and all players continue otherwise, i.e.

$$\sigma^{i_n}(n) = \begin{cases} 1, & \text{if } (i_n, r_n) = (i, r), \\ x_n, & \text{if } (i_n, r_n) = (j, \tilde{r}), \\ 0, & \text{otherwise.} \end{cases}$$

We prove that σ is a subgame-perfect 2ε -equilibrium. Let $n \in \mathbf{N}$ and consider the subgame that starts at stage n .

The definition of σ , Lemma 1 and Corollary 1 imply that (a) the probability that player i stops under σ in each block l , conditioned that the game reaches stage n_l , is at least $(1 - \varepsilon^2)\varepsilon/2$, and (b) the probability that player j stops under σ in each block l , conditioned that the game reaches stage n_l , is between $(1 - \varepsilon)\varepsilon^2/2$ and ε^2 .

This implies that $\|\gamma_n(\sigma) - r\| \leq 2\varepsilon$. Furthermore, (a) and (b) imply that under any unilateral deviation the game terminates with probability one.

Since for every player k and every $(k, r') \in \mathbf{IR}$ one has $r'^k \leq r^k \leq \gamma_n^k(\sigma) + 2\varepsilon$, no player $k \neq i$ can profit more than 2ε by deviating from σ in the subgame that starts at stage n . Since $\tilde{r}^i \leq r^i \leq \gamma_n^i(\sigma) + 2\varepsilon$, the same applies to player i .

Case 3. $r^i < 0$, and $\tilde{r}^i \geq r^i$ for every $(j, \tilde{r}) \in \mathbf{IR}$ with $i \neq j$.

In that case, by the assumption of the lemma, the strategy of player i that always continues is a weakly dominant strategy.

Consider the modified game where one sets $p_n = 0$ whenever $i_n = i$, or, alternatively, one drops all stages in which $i_n = i$. Note that player i is a dummy in the modified game.

By the induction hypothesis, the modified game admits a subgame-perfect ε -equilibrium σ' . Extend σ' to a profile σ in the original game, by instructing player i to continue at all stages n . Then σ is a subgame-perfect ε -equilibrium.

4.4. The general case

In view of Lemma 2, Theorem 1 will follow from Proposition 1 below.

Proposition 1. *Let Γ be a deterministic multi-player Dynkin game. Assume that for every $(i, r) \in \mathbf{IR}$, (i) $\pi(i, r) = +\infty$, and (ii) there is $(j, \tilde{r}) \in \mathbf{IR}$ such that $\tilde{r}^j > r^j$. Then, for every $\varepsilon > 0$, the game Γ has a subgame-perfect ε -equilibrium in pure Markovian strategies.*

Note that **Example 3** does not fit into **Proposition 1**. We do not know whether a subgame-perfect 0-equilibrium exists or not. The rest of this section is devoted to the proof of the proposition.

As remarked at the beginning of **Section 4.1**, we can assume w.l.o.g. that for every $(i, r), (j, \tilde{r}) \in \mathbb{IR}$, either $(i, r) = (j, \tilde{r})$, or $r^k \neq \tilde{r}^k$ for every k .

For every $i \in I$ set

$$m^i = \max\{r^i | (i, r) \in \mathbb{IR}\}.$$

This is the maximal terminal payoff player i can receive when stopping alone. Let $\rho_i \in \mathbf{R}^I$ be the unique vector r such that $(i, r) \in \mathbb{IR}$ and $r^i = m^i$ (uniqueness is guaranteed by the preceding paragraph).

Finally, set

$$W = \{w \in \mathbf{R}^N | w^i \leq m^i \text{ for some } i \in I\}.$$

This is the set of all payoff vectors w such that at least one player is better off by stopping at some stage rather than continuing forever and receiving w .

An important property of the set W is that if the continuation payoff at stage n is $w \in W$, and if player i_n prefers to stop rather than continue (that is, $w^{i_n} \leq r_n^{i_n}$), then the expected payoff if player i_n stops at stage n , $(1 - p_n)w + p_n r_n$, is in W . Formally, for every $n \in \mathbf{N}$,

$$w \in W \text{ and } w^{i_n} \leq r_n^{i_n} \text{ imply } (1 - p_n)w + p_n r_n \in W. \tag{7}$$

Indeed, under the assumptions, $(1 - p_n)w^{i_n} + p_n r_n^{i_n} \leq r_n^{i_n} \leq m^{i_n}$, and (7) follows.

We will prove the existence of a subgame-perfect ε -equilibrium. We assume w.l.o.g. that $\varepsilon < 1/40$, and that furthermore $\varepsilon < \frac{1}{2} \min_{(i,r) \neq (j,\tilde{r})} |r^i - \tilde{r}^i|$.

Let $l \in \mathbf{N}$ be given. We will define a pure profile σ_l up to stage n_l . We will simultaneously construct a sequence $(w_l(n))_{n=1}^{n_l}$ of vectors in W . As a first approximation, the vector $w_l(n)$ may be interpreted as the expected payoff under σ from stage n onwards.

As for now, we fix $l \in \mathbf{N}$ and we write σ and w instead of σ_l and w_l , respectively.

We define both σ and w backwards. We let $w(n_l)$ be an arbitrary point in $W \cap [-1, 1]^I$. We deal with each of the blocks inductively (starting with the l th one). Let $k \leq l$. Assuming $w(n_k) \in W$ is already defined, we define now σ and w over the stages $n = n_{k-1}, \dots, n_k - 1$.

Given $w(n + 1)$ and $\sigma^{i_n}(n)$, we set $w(n) = \sigma^{i_n}(n)p_n r_n + (1 - \sigma^{i_n}(n)p_n)w(n + 1)$, so that we need only define $\sigma^{i_n}(n)$. Thus, if $w(n + 1)$ is the expected payoff from stage $n + 1$ onwards, $w(n)$ is the expected payoff from stage n onwards. We will define $\sigma^{i_n}(n)$ such that (i) $\sigma^{i_n}(n)$ is pure, and (ii) $\sigma^{i_n}(n) = 1$ implies $r_n^{i_n} \geq w^{i_n}(n + 1)$. Since $w(n_l) \in W$ and by (7), this implies that $w(n) \in W$ for every $n \leq n_l$.

Case 1. $w^i(n_k) \leq m^i - \varepsilon$ for some $i \in I$.

We define σ by backward induction, with an appropriate tie-breaking rule. Set $\sigma^{i_n}(n) = 1$ if $r_n^{i_n} \geq w^{i_n}(n + 1)$, and $\sigma^{i_n}(n) = 0$ otherwise.

Thus, at stage n , player i_n compares his continuation payoff $w^{i_n}(n + 1)$ to the payoff $r_n^{i_n}$ he would get by stopping, and he continues or stops accordingly.

Case 2. $w^i(n_k) > m^i - \epsilon$ for each $i \in I$.

Fix $i_* \in I$ such that $w^{i_*}(n_k) \leq m^{i_*}$. Since $w(n_k) \in W$, such a player exists. We will define σ so that at the final stages of the block only player i_* will possibly stop. In earlier stages, σ will be defined using backward induction as in Case 1.

Formally, let $n_{k-1} \leq n < n_k$. Assume that σ has been defined for stages $q = n + 1, \dots, n_k - 1$. We define σ at stage n as follows. Denote by $\pi(n + 1, n_k)$ the probability under σ that, starting from stage $n + 1$, the game terminates under σ before stage n_k , i.e.

$$\pi(n_k, n_k) = 0, \text{ and}$$

$$\pi(q, n_k) = \sigma^{i_q}(q) p_q + (1 - \sigma^{i_q}(q)) \pi(q + 1, n_k) \text{ for } n + 1 \leq q < n_k.$$

Then:

- if $\pi(n + 1, n_k) < \epsilon$, we set $\sigma^{i_n}(n) = 1$ if both $i_n = i_*$ and $r_n^{i_*} \geq w^{i_*}(n)$ hold. We set $\sigma^{i_n}(n) = 0$ otherwise;
- if $\pi(n + 1, n_k) \geq \epsilon$, we set $\sigma^{i_n}(n) = 1$ if $r_n^{i_n} \geq w^{i_n}(n)$, and $\sigma^{i_n}(n) = 0$ otherwise.

We now prove that under σ , the probability of termination in any single block is bounded away from zero.

Lemma 3. *For each k such that $0 \leq k < l$, one has*

$$P_\sigma(\theta < n_{k+1} | \theta \geq n_k) \geq \frac{1}{3}\epsilon.$$

Proof. We will prove that $\pi(n_k, n_{k+1}) \geq \epsilon/3$. We consider Cases 1 and 2 in turn.

We first assume that Case 1 holds, and we let $i_* \in I$ be a player such that $w^{i_*}(n_{k+1}) \leq m^{i_*} - \epsilon$.

1. If $\sigma^{i_*}(n) = 1$ whenever $(i_n, r_n) = (i_*, \rho_{i_*})$, one has by [Corollary 1](#) $\pi(n_k, n_{k+1}) \geq \epsilon/2$.
2. If $\sigma^{i_*}(n) = 0$ for some n such that $(i_n, r_n) = (i_*, \rho_{i_*})$, then $w^{i_*}(n + 1) > m^{i_*}$. Observe now that, since payoffs are bounded by one, one has

$$w^{i_*}(n + 1) \leq \pi(n + 1, n_{k+1}) + (1 - \pi(n + 1, n_{k+1}))w^{i_*}(n_{k+1}).$$

By the choice of i_* one has $w^{i_*}(n_{k+1}) \leq m^{i_*} - \epsilon$, so that

$$\pi(n + 1, n_{k+1}) \geq \frac{\epsilon}{1 - m^{i_*} + \epsilon} \geq \frac{\epsilon}{3}.$$

Since $\pi(n_k, n_{k+1}) \geq \pi(n + 1, n_{k+1})$, the conclusion also follows in that case.

We next assume that Case 2 holds and we let $i_* \in I$ be the player distinguished in the definition of σ .

1. Assume first that $\sigma^j(n) = 1$ for some n and some player $j \neq i_*$. By definition of the profile σ , one then has $\pi(n + 1, n_{k+1}) \geq \epsilon/2$, hence $\pi(n_k, n_{k+1}) \geq \pi(n + 1, n_{k+1}) \geq \epsilon/2$.
2. Assume now that $\sigma^{i_n}(n) = 0$ whenever $i_n \neq i_*$. In that case, $w^{i_*}(n) \leq m^{i_*}$ for each n . Indeed, only player i_* stops, and his payoff is the average of $w^{i_*}(n_{k+1}) \leq m^{i_*}$

and $\rho_{i_*}^{i_*} = m^{i_*}$. Therefore, $\sigma^{i_*}(n) = 1$ whenever $(i_n, r_n) = (i_*, \rho_{i_*})$, and one gets $\pi(n_k, n_{k+1}) \geq \varepsilon/2$, as in Case 1, item 1. \square

We will now let l vary and we denote by σ_l and w_l the objects that were defined above. The pure strategy profile σ_l may be identified with a point in $\{0, 1\}^N$ (the n th component being the behavior at stage n of the active player i_n). Since the product space $\{0, 1\}^N$ is compact (and metrizable), the sequence $(\sigma_l)_{l \geq 0}$ has a subsequence that converges to some pure strategy profile σ_* . For notational convenience, we still denote this subsequence by $(\sigma_l)_{l \geq 0}$. Since σ_l is a pure strategy for every $l \in \mathbf{N}$, for every fixed $n \in \mathbf{N}$ the first n components of σ_* coincide with the first n components of σ_l , provided l is sufficiently large. For such l 's, the behavior in the first n stages of the game under the two strategy profiles σ_* and σ_l coincide.

Our goal is to prove that σ_* is a subgame-perfect ε -equilibrium. We first prove that the play terminates \mathbf{P}_{σ_*} -a.s. in each subgame. We will then relate the payoff $\gamma(\sigma_*)$ to the sequence $(w_l)_{l \in \mathbf{N}}$ (Lemma 4) and prove that no player has a profitable one-stage deviation (Lemma 5) under σ_* . The conclusion follows (Proposition 2), after we prove that no single player is responsible for the termination of the game (Lemma 6).

Corollary 2. For each $k \in \mathbf{N}$, one has

$$\mathbf{P}_{\sigma_*}(\theta < n_{k+1} | \theta \geq n_k) \geq \frac{1}{3}\varepsilon.$$

Proof. Let $l > k$ be large enough so that σ_* coincides with σ_l up to stage n_{k+1} , and apply Lemma 3. \square

Lemma 4. For each $n \in \mathbf{N}$, one has

$$\gamma_n(\sigma_*) = \lim_{l \rightarrow \infty} w_l(n).$$

Proof. We prove the result for $n = 1$. The proof is similar for the subgame that starts at any stage $n \in \mathbf{N}$.

Let $k \in \mathbf{N}$ be given. For each $l \geq k$, one has

$$\gamma(\sigma_*) = \mathbf{E}_{\sigma_*}[r_\theta 1_{\theta < n_k}] + \mathbf{P}_{\sigma_*}(\theta \geq n_k)\gamma_{n_k}(\sigma_*),$$

and

$$w_l(1) = \mathbf{E}_{\sigma_l}[r_\theta 1_{\theta < n_k}] + \mathbf{P}_{\sigma_l}(\theta \geq n_k)w_l(n_k).$$

For l large enough, the two profiles σ_l and σ_* coincide up to stage n_k . In particular, $\mathbf{E}_{\sigma_*}[r_\theta 1_{\theta < n_k}] = \mathbf{E}_{\sigma_l}[r_\theta 1_{\theta < n_k}]$ and $\mathbf{P}_{\sigma_*}(\theta \geq n_k) = \mathbf{P}_{\sigma_l}(\theta \geq n_k)$. By Corollary 2

$$\|\gamma(\sigma_*) - w_l(1)\| \leq 2(1 - \frac{1}{3}\varepsilon)^k$$

provided l is large enough, and the result follows. \square

The next lemma says in substance that no player can increase his payoff by more than 3ε by modifying his strategy in a single stage.

Lemma 5. Let $n \in \mathbf{N}$ be given. The following implications hold.

- If $\sigma_*^{in}(n) = 0$ then $\gamma_{n+1}^{in}(\sigma_*) \geq r_n^{in} - 3\varepsilon$.
- If $\sigma_*^{in}(n) = 1$ then $\gamma_{n+1}^{in}(\sigma_*) \leq r_n^{in}$.

Proof. Let $n \in \mathbf{N}$ be given. Let $l \in \mathbf{N}$ be sufficiently large so that $n_l > n$. We first prove a related statement for the strategy profile σ_l . Let $k < l$ be determined by $n_k \leq n < n_{k+1}$.

By construction, $\sigma_l^{in}(n)$ is defined using backward induction, except in one case where $\sigma_l^{in}(n)$ is required to be zero. In Case 1 one has $\sigma_l^{in}(n) = 1$ if $r_n^{in} \geq w_l^{in}(n+1)$ and $\sigma_l^{in}(n) = 0$ otherwise. In Case 2 one has $w_l^{in}(n_{k+1}) \geq m^{in} - \varepsilon$ and $\pi(n+1, n_{k+1}) < \varepsilon$. Therefore, $|w_l^{in}(n_{k+1}) - w_l^{in}(n+1)| < 2\varepsilon$, which yields $w_l^{in}(n+1) \geq m^{in} - 3\varepsilon \geq r_n^{in} - 3\varepsilon$. Hence, in both cases, one has

$$\begin{aligned} w_l^{in}(n+1) &\leq r_n^{in}, & \text{if } \sigma_l^{in}(n) = 1, \text{ and} \\ w_l^{in}(n+1) &\geq r_n^{in} - 3\varepsilon, & \text{if } \sigma_l^{in}(n) = 0. \end{aligned}$$

The conclusion follows by taking the limit $l \rightarrow +\infty$ and using Lemma 4. □

We now prove that the play terminates a.s., even if a single player chooses to continue whenever active. Recall that c^i is the strategy of player i that always continues.

Lemma 6. For every $i \in I$ and every $n \in \mathbf{N}$, one has

$$P_{c^i, \sigma_*^{-i}}(\theta < +\infty | \theta \geq n) = 1.$$

Proof. We argue by contradiction, and we assume that, for some player $i \in I$, the sequence $P_{c^i, \sigma_*^{-i}}(\theta < +\infty | \theta \geq n)$ converges to zero when n goes to $+\infty$. By Corollary 2 the game eventually terminates, so that $P_{\sigma_*}(\theta < +\infty | \theta \geq n) = 1$ for every n . Therefore, it must be the case that player i terminates the game: $P_{\sigma_*^i, c^{-i}}(\theta < +\infty | \theta \geq n) = 1$ for every n , and $\lim_{n \rightarrow +\infty} \|\gamma_n(\sigma_*) - \gamma_n(\sigma_*^i, c^{-i})\| = 0$.

We first prove that $\lim_{n \rightarrow +\infty} \gamma_n(\sigma_*) = \rho_i$, and then deduce a contradiction with the basic assumption made on Γ .

Step 1. The sequence $(\gamma_n^i(\sigma_*))_{n \in \mathbf{N}}$ has a limit.

Let $n \in \mathbf{N}$ be arbitrary. If $n \in N_i$ then Lemma 5 implies that $\gamma_n^i(\sigma_*) \geq \gamma_{n+1}^i(\sigma_*)$. On the other hand, for $n \notin N_i$, one has $\gamma_n^i(\sigma_*) = \gamma_{n+1}^i(\sigma_*)$ if $\sigma_*^{in}(n) = 0$, and

$$|\gamma_n^i(\sigma_*) - \gamma_{n+1}^i(\sigma_*)| = p_n |r_n^i - \gamma_{n+1}^i(\sigma_*)| \leq 2p_n$$

if $\sigma_*^{in}(n) = 1$. Therefore, for every two positive integers $n \geq m$, one has

$$\gamma_m^i(\sigma_*) \geq \gamma_n^i(\sigma_*) - 2 \sum_{m \leq q < n; q \notin N_i} p_q \mathbf{1}_{\sigma_*^{iq}(q)=1}. \tag{8}$$

Let $\tilde{\varepsilon} \in (0, 1/40)$ be given. Choose $N_{\tilde{\varepsilon}} \in \mathbf{N}$ sufficiently large so that $\mathbf{P}^{c^i, \sigma_*^{-i}}(\theta < +\infty | \theta \geq N_{\tilde{\varepsilon}}) < \tilde{\varepsilon}/2$. For such $N_{\tilde{\varepsilon}}$, one has by Corollary 1 $\sum_{N_{\tilde{\varepsilon}} \leq q < +\infty; q \notin N_i} Pq \mathbf{1}_{\sigma_*^i(q)=1} \leq \tilde{\varepsilon}$. Therefore, by (8),

$$\gamma_m^i(\sigma_*) \geq \gamma_n^i(\sigma_*) - 2\tilde{\varepsilon}, \text{ for every } n \geq m \geq N_{\tilde{\varepsilon}}.$$

This implies the convergence of $(\gamma_n^i(\sigma_*))_{n \in \mathbf{N}}$, since it is a bounded sequence.

Step 2. $\lim_{n \rightarrow +\infty} \gamma_n(\sigma_*) = \rho_i$.

Denote $\lambda := \lim_{n \rightarrow +\infty} \gamma_n(\sigma_*)$. We first prove that $\lambda^i = m^i$.

Fix $\delta > 0$ sufficiently small, and take k sufficiently large so that (i) $|\gamma_{n_k}(\sigma_*) - \lambda| < \delta$, (ii) $|\gamma_{n_{k+1}}(\sigma_*) - \lambda| < \delta$, and (iii) $\mathbf{P}^{c^i, \sigma_*^{-i}}(\theta < +\infty | \theta \geq n_k) < \delta$.

By Lemma 4 and since $\sigma_* = \lim_{l \rightarrow +\infty} \sigma_l$, there is $l > k$ sufficiently large such that (i) $|w_l(n_k) - \lambda| < \delta$, (ii) $|w_l(n_{k+1}) - \lambda| < \delta$, and (iii) $\mathbf{P}^{c^i, \sigma_l}(\theta < n_{k+1} | \theta \geq n_k) < \delta$.

Consider now the block that is played between stages n_k and n_{k+1} under σ_l . By (iii), the probability that the game terminates by a player $j \neq i$ is smaller than δ . Therefore, player i never stops at a stage n such that $i_n = i$ and $r_n^i < \lambda^i - 2\delta$. However, the probability that player i stops at a stage n such that $i_n = i$ and $r_n^i = m^i$ is at least $\varepsilon/2$. Therefore,

$$2\delta > w_l^i(n_k) - w_l^i(n_{k+1}) > \frac{1}{2}\varepsilon(m^i - \lambda^i) - 3\delta \text{ and } -\delta < m^i - \lambda^i,$$

so that $-\delta < m^i - \lambda^i < 10\delta/\varepsilon$. As δ is arbitrary, the first claim follows.

Hence, $\lim_{n \rightarrow +\infty} \gamma_n^i(\sigma_*) = m^i$. This yields $\lim_{n \rightarrow +\infty} \gamma_n^i(\sigma_*^i, c^{-i}) = m^i$. Since $\rho_i \in \mathbf{R}^I$ is the unique vector such that $(i, r) \in \mathbf{IR}$ and $r^i = m^i$, and since $\gamma_n(\sigma_*^i, c^{-i})$ is in the convex hull of $\{\tilde{r}: (i, \tilde{r}) \in \mathbf{IR}\}$, one has $\lim_{n \rightarrow +\infty} \gamma_n(\sigma_*^i, c^{-i}) = \rho_i$. Finally, this implies $\lim_{n \rightarrow +\infty} \gamma_n(\sigma_*) = \rho_i$.

Step 3. The contradiction.

By assumption, there exists $(j, \tilde{r}) \in \mathbf{IR}$ such that $\tilde{r}^j > \rho_j^j$. Since $\mathbf{P}^{c^i, \sigma_*^{-i}}(\theta < +\infty | \theta \geq m) < 1$ for some $m \in \mathbf{N}$, and since $\pi(j, \tilde{r}) = +\infty$, there are infinitely many stages n such that $(i_n, r_n) = (j, \tilde{r})$ and $\sigma_*^j(n) = 0$. For each such n , by Lemma 5, one has $\tilde{r}^j = r_n^j < \gamma_{n+1}^j(\sigma_*)$. Therefore, $\limsup_{n \rightarrow +\infty} \gamma_n^j(\sigma_*) \geq \tilde{r}^j$. Since $\tilde{r}^j > \rho_j^j = \lim_{n \rightarrow +\infty} \gamma_n^j(\sigma_*)$, we get a contradiction.

Proposition 2. σ_* is a subgame-perfect 3ε -equilibrium.

Proof. Let $i \in I$ be given. We prove that player i cannot gain more than 3ε by deviating from σ_* . The same proof will hold in any subgame, thereby showing the subgame-perfectness property.

Define the sequence $(X_n)_{n \in \mathbf{N}}$ of random variables by $X_n = r_\theta^i$ if $\theta < n$ and $X_n = \gamma_n^i(\sigma_*)$ if $\theta \geq n$. Let τ^i be an arbitrary strategy of player i . By Lemma 6, the sequence $(X_n)_{n \in \mathbf{N}}$ converges $\mathbf{P}_{\tau^i, \sigma_*^{-i}}$ -a.s. to $X_\infty := r_\theta \mathbf{1}_{\theta < +\infty}$, hence

$$\lim_{n \rightarrow +\infty} \mathbf{E}_{\tau^i, \sigma_*^{-i}}[X_n] = \mathbf{E}_{\tau^i, \sigma_*^{-i}}[r_\theta \mathbf{1}_{\theta < +\infty}] = \gamma^i(\tau^i, \sigma_*^{-i}). \tag{9}$$

On the other hand, let $n \in \mathbf{N}$, and denote by \mathcal{H}_n the past play up to stage n . We shall prove that

$$\mathbf{E}_{\tau^i, \sigma_*^{-i}}[X_{n+1} | \mathcal{H}_n] \leq X_n + 3\varepsilon \mathbf{1}_{\theta=n}, \text{ a.s.} \quad (10)$$

On the event $\theta < n$, both X_n and X_{n+1} are equal to r_θ^i . Consider now the event $\theta \geq n$. If $i_n \neq i$, one has

$$X_n = \gamma_n^i(\sigma_*) = \mathbf{E}_{\sigma_*}[X_{n+1} | \mathcal{H}_n] = \mathbf{E}_{\tau^i, \sigma_*^{-i}}[X_{n+1} | \mathcal{H}_n],$$

where the last equality follows since the two profiles (τ^i, σ_*^{-i}) and σ_* coincide at stage n . In both cases, (10) follows trivially. Finally, if $i_n = i$, one has $X_n = X_{n+1} = \gamma_{n+1}^i(\sigma_*)$ if $\sigma_*^i(n) = 0$ and $\mathbf{E}_{\tau^i, \sigma_*^{-i}}[X_{n+1} | \mathcal{H}_n] = p_n r_n^i + (1 - p_n) \gamma_{n+1}^i(\sigma_*)$ otherwise. Inequality (10) then follows by Lemma 5.

By taking expectations in (10), and by summing over n , one obtains $\lim_{n \rightarrow +\infty} \mathbf{E}_{\tau^i, \sigma_*^{-i}}[X_n] \leq X_1 + 3\varepsilon$ which yields, by (9),

$$\gamma^i(\tau^i, \sigma_*^{-i}) \leq \gamma^i(\sigma_*) + 3\varepsilon. \quad \square$$

Acknowledgements

We acknowledge the financial support of the Arc-en-Ciel/Keshet program for 2001/2002. The research of the first author was supported by the Israel Science Foundation (Grant No. 03620191).

References

- Dynkin, E.B., 1969. Game variant of a problem on optimal stopping. *Soviet Mathematics Doklady* 10, 270–274.
- Fine, C.H., Li, L., 1989. Equilibrium exit in stochastically declining industries. *Games and Economic Behavior* 1, 40–59.
- Flesch, J., Thuijsman, F., Vrieze, K., 1997. Cyclic Markov equilibria in stochastic games. *International Journal of Game Theory* 26, 303–314.
- Fudenberg, D., Tirole, J., 1991. *Game Theory*. MIT Press, Cambridge, MA.
- Ghemawat, P., Nalebuff, B., 1985. Exit. *RAND Journal of Economics* 16, 184–194.
- Kiefer, Y.I., 1971. Optimal stopped games. *Theory of Probability and its Application* 16, 185–189.
- Kilgour, D.M., 1975. The sequential truel. *International Journal of Game Theory* 4, 151–174.
- Kilgour, D.M., 1977. Equilibrium points in infinite sequential truels. *International Journal of Game Theory* 6, 167–180.
- Kilgour, D.M., Brams, S.J., 1997. The truel. *Mathematics Magazine* 70, 315–325.
- Neveu, J., 1975. *Discrete-Parameter Martingales*. North-Holland, Amsterdam.
- Rosenberg, D., Solan, E., Vieille, N., 2001. Stopping games with randomized strategies. *Probability Theory and Related Fields* 119, 433–451.
- Rosenberg, D., Solan, E., Vieille, N., 2002. On the maxmin value of stochastic games with imperfect monitoring. Discussion Paper 1337. The Center for Mathematical Studies in Economics and Management Science, Kellogg School of Management, Northwestern University.
- Shmaya, E., Solan, E., 2002. Two-player non-zero-sum stopping games in discrete time. Discussion Paper 1347. The Center for Mathematical Studies in Economics and Management Science, Kellogg School of Management, Northwestern University.

Solan, E., 1999. Three-person absorbing games. *Mathematics of Operations Research* 24, 669–698.

Solan, E., 2002. Subgame perfection in quitting games with perfect information and differential equations. Discussion Paper 1356. The Center for Mathematical Studies in Economics and Management Science, Kellogg School of Management, Northwestern University.

Solan, E., Vieille, N., 2001. Quitting games. *Mathematics of Operations Research* 26, 265–285.

Vieille, N., 2000. Equilibrium in 2-person stochastic games. II. The case of recursive games. *Israel Journal of Mathematics* 119, 93–126.