Probab. Theory Relat. Fields 119, 433–451 (2001)
Digital Object Identifier (DOI) 10.1007/s004400000091

Dinah Rosenberg · Eilon Solan · Nicolas Vieille

Stopping games with randomized strategies

Received: 10 May 1999 / Revised version: 18 May 2000 /
Published online: 15 February 2001 –

Abstract. We study stopping games in the setup of Neveu. We prove the existence of a uniform value (in a sense defined below), by allowing the players to use randomized strategies. In constrast with previous work, we make no comparison assumption on the payoff processes. Moreover, we prove that the value is the limit of discounted values, and we construct ϵ-optimal strategies.

1. Introduction

Dynkin (1969) introduced the following optimization problem. Two players observe stochastic sequences $(r(n), x(n))_n$. Player 1 (*resp.* player 2) is allowed to stop whenever $x(n) \leq 0$ (*resp.* $x(n) > 0$). The two players choose stopping times μ_1 and μ_2 which obey this rule, and the payoff is given by

$$\gamma(\mu_1, \mu_2) = \mathbf{E}\{1_{\mu_1 < \mu_2} r(\mu_1) + 1_{\mu_1 > \mu_2} r(\mu_2)\}.$$

The goal of player 1 is to maximize $\gamma(\mu_1, \mu_2)$, whereas player 2 tries to minimize $\gamma(\mu_1, \mu_2)$. Dynkin proved that this game has a value if $\sup_n |r(n)| \in L^1$, and constructed ϵ-optimal strategies for the two players.

Kiefer (1971) and Neveu (1975) gave other sufficient conditions for existence of the value in this zero-sum game and in a variant of it. Neveu extended the game by allowing the players to stop simultaneously: a process (a_n, b_n, c_n) is given (with $\sup_n \sup(|a_n|, |b_n|, |c_n|) \in L^1$), the two players choose stopping times μ_1 and μ_2, and the payoff to player 1 is

$$\mathbf{E}\{a_{\mu_1} 1_{\mu_1 < \mu_2} + b_{\mu_2} 1_{\mu_2 < \mu_1} + c_{\mu_1} 1_{\mu_1 = \mu_2 < +\infty}\}.$$

He proved that, under the assumption $a_n = c_n \leq b_n$, the game has a value.

D. Rosenberg: Laboratoire d'Analyse, Géométrie et Applications, Institut Galilée, Université Paris Nord, avenue Jean-Baptiste Clément, 93430 Villetaneuse, France. e-mail: `dinah@math.univ-paris13.fr`

E. Solan: MEDS Department, Kellogg Graduate School of Management, Northwestern University, 2001 Sheridan Rd., Evanston, IL 60208, USA, and School of Mathematical Sciences, Tel Aviv University, Israel. e-mail: `eilons@post.tau.ac.il`

N. Vieille: HEC, Département Economie et Finance, and Laboratoire d'Econométrie de l'Ecole Polytechnique, 1 rue Descartes, 75005 Paris, France. e-mail: `vieille@poly.polytechnique.fr`

We thank an anonymous referee for many useful comments

Mathematics Subject Classification (2000): 60G40, 91A15

There is a broad literature on continuous time Dynkin games giving sufficient conditions for the existence of the value and optimal strategies: Bismut (1979) proved that under the hypothesis $a_n = c_n \leq b_n$, some regularity assumption and Mokobodski's hypothesis (namely that there exist positive bounded supermartingales z and z' satisfying $a \leq z - z' \leq b$) the value exists. The regularity assumption was weakened by Alario-Nazaret, Lepeltier and Marchal (1982), and then Lepeltier and Maingueneau (1984) established the existence of the value and optimal strategies without Mokobodski's hypothesis, assuming only $a_n = c_n \leq b_n$.

In the present paper, we focus on discrete time Dynkin games and we allow the players to use randomized stopping times. We prove the existence of the value, under the single integrability condition.

This result is related to a result due to Maitra and Sudderth (1993), for general stochastic games. In such games, the players receive a payoff in each stage. Maitra and Sudderth define the payoff associated to a play as the lim sup of the payoffs received along the play. They prove that such games have a value, provided the payoffs are bounded and deterministic functions of the state.

It is clear that, under some regularity assumptions on the processes (a_n), (b_n) and (c_n), stopping games may be viewed as general stochastic games with a very specific transition structure (note however that boundedness of the payoff function will not be satisfied). Thus, the result of Maitra and Sudderth has some bite in stopping games. We emphasize that our method bears no relation to their approach (which is based on transfinite induction).

Our contribution is threefold. (i) We prove that the value exists under the single integrability requirement, and, moreover, it is uniform in a sense defined below. (ii) We prove that the value is the limit of the so-called discounted values, studied by Yasuda (1985). In particular, it follows that the discounted values converge. (iii) We construct ϵ-optimal strategies for the players.

Our method is to construct a strategy for player 1 that guarantees him an expected payoff which is, up to an ϵ, the limit of some sequence of discounted values. We provide two different constructions for an ϵ-optimal strategy. In the first construction the player plays at each stage an optimal discounted strategy, where the discount factor may change from time to time. In the second construction, which has the flavor of Dynkin's construction, the player plays almost the limit of the optimal discounted strategies.

The paper is arranged as follows. In section 2 we present the model and the main results, in section 3 we introduce few tools, in section 4 we explain the main ideas of the two constructions, and finally, in sections 5.2 and 5.3 we provide the two constructions of ϵ-optimal strategies. Section 6 concludes the paper by discussing related issues.

2. The model and the main results

Let $(\Omega, \mathscr{A}, \mathbf{P})$ be a probability space, and $(\mathscr{F}_n)$ be a filtration over $(\Omega, \mathscr{A}, \mathbf{P})$ (the information available at stage n). Let (a_n), (b_n), (c_n) be processes, defined over

$(\Omega, \mathscr{A}, \mathbf{P})$. We assume

$$\sup_n |a_n|, \sup_n |b_n|, \sup_n |c_n| \in L^1(\mathbf{P}). \tag{1}$$

We also assume that (a_n), (b_n), and (c_n) are adapted. This assumption can be dispensed with. One needs only replace everywhere (a_n), (b_n), and (c_n) by their conditional expectations given $\mathscr{F}_n$. It is also convenient to assume $\mathscr{A} = \sigma(\mathscr{F}_n, n \geq 0)$.

By properly enlarging the probability space $(\Omega, \mathscr{A}, \mathbf{P})$, one can assume w.l.o.g. that it supports a double sequence $(X_n, Y_n)_{n=0}^{\infty}$ of *iid* variables, uniformly distributed over $[0, 1]$, such that, for each n: (i) (X_n, Y_n) is independent of the process $(a_k, b_k, c_k)_k$; (ii) (X_n, Y_n) is $\mathscr{F}_{n+1}$-measurable, and independent of $\mathscr{F}_n$.

Define the stopping game as follows. A *strategy* for player 1 (*resp.* player 2) is a $[0, 1]$-valued, adapted process $\mathbf{x} = (x_n)$ (*resp.* $\mathbf{y} = (y_n)$): x_n is the probability that player 1 stops at stage n, conditional on stopping occurs after $n-1$. The interpretation of a strategy as a randomized stopping time will be discussed in Section 6.

Given strategies $(\mathbf{x}, \mathbf{y})$, define the stopping stages of players 1 and 2 by $t_1 = \inf\{n \geq 0, X_n \leq x_n\}$, $t_2 = \inf\{n \geq 0, Y_n \leq y_n\}$, and set

$$t = \min(t_1, t_2). \tag{2}$$

Notice that $t + 1$ is a stopping time, but t needs not be.

We set $r(\mathbf{x}, \mathbf{y}) = a_{t_1} 1_{t_1 < t_2} + b_{t_2} 1_{t_2 < t_1} + c_{t_1} 1_{t_1 = t_2 < +\infty}$. The payoff of the game is $\gamma(\mathbf{x}, \mathbf{y}) = \mathbf{E}(r(\mathbf{x}, \mathbf{y}))$. The goal of player 1 is to maximize $\gamma(\mathbf{x}, \mathbf{y})$, and the goal of player 2 is to minimize it.

Definition 2.1. $v \in \mathbf{R}$ *is the* value *of the game if* $v = \sup_{\mathbf{x}} \inf_{\mathbf{y}} \gamma(\mathbf{x}, \mathbf{y}) = \inf_{\mathbf{y}} \sup_{\mathbf{x}} \gamma(\mathbf{x}, \mathbf{y})$. *Let* $\epsilon > 0$. *A strategy* $\mathbf{x}$ *that satisfies* $\inf_{\mathbf{y}} \gamma(\mathbf{x}, \mathbf{y}) \geq v - \epsilon$ *is an* ϵ-optimal strategy for player 1. *A strategy* $\mathbf{y}$ *that satisfies* $\sup_{\mathbf{x}} \gamma(\mathbf{x}, \mathbf{y}) \leq v + \epsilon$ *is an* ϵ-optimal strategy for player 2.

We will establish the following:

Theorem 2.2. *Every zero-sum stopping game that satisfies (1) has a value* v.

Let $\lambda \in]0, 1[$. Define the λ-discounted payoff by $r_\lambda(\mathbf{x}, \mathbf{y}) = (1 - \lambda)^{t+1} r(\mathbf{x}, \mathbf{y})$ and $\gamma_\lambda(\mathbf{x}, \mathbf{y}) = \mathbf{E}(r_\lambda(\mathbf{x}, \mathbf{y}))$.

Definition 2.3. v_λ *is the* λ-discounted value *of the game if*

$$v_\lambda = \sup_{\mathbf{x}} \inf_{\mathbf{y}} \gamma_\lambda(\mathbf{x}, \mathbf{y}) = \inf_{\mathbf{y}} \sup_{\mathbf{x}} \gamma_\lambda(\mathbf{x}, \mathbf{y}).$$

Yasuda (1985) proves that the λ-discounted value always exists. In the sequel we prove that

Theorem 2.4. $v = \lim_{\lambda \to 0} v_\lambda$.

In particular, $\lim_{\lambda \to 0} v_\lambda$ exists.

Set $\gamma_n(\mathbf{x}, \mathbf{y}) = \mathbf{E}(\frac{n-t}{n} r(\mathbf{x}, \mathbf{y}) 1_{t<n})$. The natural interpretation of $\gamma_n(\mathbf{x}, \mathbf{y})$ is in terms of average payoffs: for $k \in \mathbf{N}$, set $g_k = r(\mathbf{x}, \mathbf{y})$ on $\{t < k\}$ and $g_k = 0$ otherwise. Then $\gamma_n(\mathbf{x}, \mathbf{y}) = \mathbf{E}(\frac{1}{n} \sum_{k=0}^{n-1} g_k)$.

By dominated convergence, $\lim_n \gamma_n(\mathbf{x}, \mathbf{y}) = \gamma(\mathbf{x}, \mathbf{y})$. Therefore, if $\mathbf{x}^*$ is an ϵ-optimal strategy of player 1, then for every $\mathbf{y}$ there exists a stage N such that $\gamma_n(\mathbf{x}^*, \mathbf{y}) \geq v - 2\epsilon$ holds for every $n \geq N$.

We prove that the value v is uniform in the sense below.

Theorem 2.5. *For every $\epsilon > 0$, there exist $\mathbf{x}^*$ and $N \in \mathbf{N}$, such that, for every $\mathbf{y}$ and every $n \geq N$, $\gamma_n(\mathbf{x}^*, \mathbf{y}) \geq v - \epsilon$. A symmetric result holds for player 2.*

Thus, Theorem 2.5 is a strengthening of Theorem 2.2. It can be shown that it also implies Theorem 2.4. We then say that v is the *uniform* value of the game.

Theorem 2.5 was proved by Mertens and Neyman (1981) for general stochastic games with bounded payoffs, in which the function $\lambda \mapsto v_\lambda$ satisfies some bounded variation property. In the case of recursive games with bounded payoffs, Rosenberg and Vieille (2000) proved that Theorem 2.5 holds, if (v_λ) converge uniformly as λ goes to 0 (the uniformity is with respect to the initial state of the game). Our proof does not require any conditions on the discounted values.

3. Local games

3.1. Reminder and definitions

Let $g : A \times B \to \mathbf{R}$, where A and B are finite sets (g is the payoff function of a zero-sum matrix game with action sets A and B). Denote by $\Delta(A)$ and $\Delta(B)$ the sets of probability distributions over A and B, and still by g the bilinear extension of g to $\Delta(A) \times \Delta(B)$.

The min max theorem states that $\sup_{x \in \Delta(A)} \inf_{y \in \Delta(B)} g(x, y) = \inf_{y \in \Delta(B)} \sup_{x \in \Delta(A)} g(x, y)$, which we denote by val g. Any x (*resp.* y) which achieves the sup on the left side (*resp.* inf on the right side) is called an optimal strategy of player 1 (*resp.* player 2). It is well known that the operator val is non-decreasing and non-expansive: val $f \leq$ val g if $f \leq g$, and $|\text{val } f - \text{val } g| \leq \sup_{A \times B} |f - g|$.

For any real-valued $\mathscr{F}_n$-measurable function f, we let $G_n(f)$ be the 0-sum game with ($\mathscr{F}_n$-measurable) payoff matrix

f	b_n
a_n	c_n

in which player 1 chooses a row and player 2 a column.

A strategy of player 1 in this game is a [0, 1]-valued, $\mathscr{F}_n$-measurable variable x_n, to be interpreted as the probability that player 1 chooses the bottom row. A strategy of player 2 is defined analogously.

Define $G_n(x_n, y_n; f)$ to be the ($\mathscr{F}_n$-measurable) payoff to player 1 when the players use strategies x_n and y_n:

$$G_n(x_n, y_n; f) = x_n(1 - y_n)a_n + y_n(1 - x_n)b_n + x_n y_n c_n + (1 - x_n)(1 - y_n)f.$$

By the min max theorem, for every $\omega \in \Omega$ the game with payoff matrix

$f(\omega)$	$b_n(\omega)$
$a_n(\omega)$	$c_n(\omega)$

has a value, denoted by val $G_n(f)(\omega)$, for every $\omega \in \Omega$.

We now argue that each player has an optimal strategy in $G_n(f)$.

Proposition 3.1. *Let f be $\mathcal{F}_n$-measurable and real-valued. There exists a strategy x_n in $G_n(f)$, such that, for every y,*

$$G_n(x_n, y; f) \geq \text{val } G_n(f) \textit{ everywhere}.$$

A symmetric property holds for player 2.

Proof. For every $\omega \in \Omega$, the game with payoff matrix

$f(\omega)$	$b_n(\omega)$
$a_n(\omega)$	$c_n(\omega)$

has optimal strategies for both players. Since f, a_n, b_n and c_n are all $\mathcal{F}_n$-measurable, the map which associates to each ω the set of optimal strategies for player 1 is upper-semi-continuous and $\mathcal{F}_n$-measurable. By Kuratowski and Ryll-Nardzewski (1965) it has an $\mathcal{F}_n$-measurable selection. □

Any x_n that satisfies the conclusion of Proposition 3.1 is said to be optimal in the game $G_n(f)$. If x_n and y_n are optimal strategies in $G_n(f)$, one has $G_n(x_n, y_n; f) =$ val $G_n(f)$ everywhere. In particular, val $G_n(f)$ is $\mathcal{F}_n$-measurable.

3.2. Local games and discounted values

It is useful to extend the notions of discounted values to the game *starting at stage n*.

For $n \in \mathbf{N}$, set $\Sigma_n = \{\mathbf{x}, x_p = 0, \forall p < n\}$, and $T_n = \{\mathbf{y}, y_p = 0, \forall p < n\}$. Those are strategies where the probability that the players stop before stage n is zero. Set

$$\underline{v}_n(\lambda) = \text{esssup}_{\Sigma_n} \text{essinf}_{T_n} \mathbf{E}[(1-\lambda)^{-n} r_\lambda(\mathbf{x}, \mathbf{y}) | \mathcal{F}_n],$$

and

$$\overline{v}_n(\lambda) = \text{essinf}_{T_n} \text{esssup}_{\Sigma_n} \mathbf{E}[(1-\lambda)^{-n} r_\lambda(\mathbf{x}, \mathbf{y}) | \mathcal{F}_n].$$

The proposition below contains obvious properties.

Proposition 3.2. $(\overline{v}_n(\lambda))_n$ *and* $(\underline{v}_n(\lambda))_n$ *are adapted processes. Moreover,* $\sup_n |\overline{v}_n(\lambda)|, \sup_n |\underline{v}_n(\lambda)| \in L^1(\mathbf{P})$.

Yasuda (1985) proves that $(\overline{v}_n(\lambda))_n$ and $(\underline{v}_n(\lambda))_n$ are both solutions of the recursive equation

$$v_n(\lambda) = (1-\lambda)\text{val } G_n(\mathbf{E}[v_{n+1}(\lambda)|\mathcal{F}_n]), \ \mathbf{P} - a.s. \tag{3}$$

He then proves that any solution of this sequence of equations is at most $(\underline{v}_n(\lambda))$ and at least $(\overline{v}_n(\lambda))$. Since $\overline{v}_n(\lambda) \geq \underline{v}_n(\lambda)$ it follows that the two are equal, $\mathbf{P}$-a.s.

We give a shorter argument, adapted from Shapley (1953). Since the value operator is non-expansive,

$$\begin{aligned}|\overline{v}_n(\lambda) - \underline{v}_n(\lambda)| &\leq (1-\lambda)|\mathbf{E}[\overline{v}_{n+1}(\lambda) - \underline{v}_{n+1}(\lambda)|\mathscr{F}_n]| \\ &\leq (1-\lambda)\mathbf{E}[|\overline{v}_{n+1}(\lambda) - \underline{v}_{n+1}(\lambda)||\mathscr{F}_n]\end{aligned}$$

By taking expectations, one obtains

$$\begin{aligned}\|\overline{v}_n(\lambda) - \underline{v}_n(\lambda)\|_1 &\leq (1-\lambda)\|\overline{v}_{n+1}(\lambda) - \underline{v}_{n+1}(\lambda)\|_1 \\ &\leq (1-\lambda)^p\|\overline{v}_{n+p}(\lambda) - \underline{v}_{n+p}(\lambda)\|_1\end{aligned}$$

for each $p \in \mathbf{N}$. Since $\sup_n |\overline{v}_n(\lambda)|, \sup_n |\underline{v}_n(\lambda)| \in L^1(\mathbf{P})$, one obtains by letting $p \to \infty$ that $\overline{v}_n(\lambda) = \underline{v}_n(\lambda)$, $\mathbf{P}$-a.s. We define $v_n(\lambda) = \underline{v}_n(\lambda)$ $(= \overline{v}_n(\lambda))$ to be the λ-discounted value of the game starting at stage n. Notice that $v(\lambda) = \mathbf{E}[v_0(\lambda)]$.

We now let $(\bar{\lambda}_p)_p$ be any decreasing sequence which converges to 0. Set $v_n = \limsup_{p\to\infty} v_n(\bar{\lambda}_p)$, and $w = \mathbf{E}[v_0]$. We shall prove the next proposition.

Proposition 3.3. *For every $\varepsilon > 0$, there is a strategy $\overline{\mathbf{x}}$ of player 1, and a positive integer N such that*

$$\forall \mathbf{y}, \forall n \geq N, \ \gamma_n(\overline{\mathbf{x}}, \mathbf{y}) \geq w - \varepsilon.$$

We now explain why Proposition 3.3 implies Theorem 2.5 – w is the value of the game. Define $z_0 = \liminf_{p\to\infty} v_0(\bar{\lambda}_p)$, and $z = \mathbf{E}[z_0]$. By symmetry, for each ϵ, there exists a strategy $\overline{\mathbf{y}}$ such that $\gamma_n(\mathbf{x}, \overline{\mathbf{y}}) \leq z + \epsilon$ for each $\mathbf{x}$, provided n is large enough. This readily implies $w - \epsilon \leq z + \epsilon$. Since $z \leq w$, and ϵ is arbitrary, one obtains $w = z$. This shows that w is the uniform value of the game. The claim about the limit of discounted values is now immediate, since the sequence $(\bar{\lambda}_p)$ used to define w is arbitrary.

The following result will be used later.

Proposition 3.4. *One has $v_n \leq \text{val } G_n(\mathbf{E}[v_{n+1}|\mathscr{F}_n])$, for every n.*

Proof. Recall that $v_n(\lambda) = (1-\lambda)\text{val } G_n(\mathbf{E}[v_{n+1}(\lambda)|\mathscr{F}_n])$. By monotonicity of the value operator,

$$v_n(\bar{\lambda}_q) \leq \alpha_q \text{val } G_n(\mathbf{E}[\sup_{p\geq q} v_{n+1}(\bar{\lambda}_p)|\mathscr{F}_n]), \text{ for each } q, \tag{4}$$

where $\alpha_q = 1 - \bar{\lambda}_q$ if the val is negative, and 1 otherwise. By dominated convergence, $\lim_{q\to+\infty} \mathbf{E}[\sup_{p\geq q} v_{n+1}(\bar{\lambda}_p)|\mathscr{F}_n] = \mathbf{E}[v_{n+1}|\mathscr{F}_n]$. Since the val operator is non-expansive, the right-hand side of (4) converges to val $G_n(\mathbf{E}[v_{n+1}|\mathscr{F}_n])$, $\mathbf{P}$-a.s. The result follows. □

3.3. Locally optimal strategies and martingale properties

Denote by $x_n(\lambda)$ and by x_n^* optimal strategies of player 1 in the local games $G_n(\mathbf{E}[v_{n+1}(\lambda)|\mathscr{F}_n])$ and $G_n(\mathbf{E}[v_{n+1}|\mathscr{F}_n])$, which exist by Proposition 3.1.

Thus, for every strategy $\mathbf{y}$ and every $n \geq 0$, one has

$$G_n(x_n^*, y_n; \mathbf{E}[v_{n+1}|\mathscr{F}_n]) \geq v_n, \ \mathbf{P}\text{-a.s.} \tag{5}$$

and

$$(1-\lambda)G_n(x_n(\lambda), y_n; \mathbf{E}[v_{n+1}(\lambda)|\mathscr{F}_n]) \geq v_n(\lambda) \ \mathbf{P}\text{-a.s.} \tag{6}$$

Recall that $v_n(\lambda)$ is to be interpreted as the value of the (discounted) game starting in stage n, *conditional* on the fact that the game has not been stopped. Define the strategies $\mathbf{x}(\lambda) = (x_n(\lambda))_n$ and $\mathbf{x}^* = (x_n^*)_n$.

Equation 3 and Proposition 3.4 provide recursive formulas for $(v_n)_n$ and $(v_n(\lambda))_n$. In order to interpret these formulas in terms of submartingale properties, we use auxiliary processes.

For clarity of exposition, given any two events E and A in $\mathscr{A}$, we say that E holds $\mathbf{P}$-a.s. on A if $\mathbf{P}(A \cap E^c) = 0$. We will frequently omit the qualification $\mathbf{P}$-a.s.

Let $(\alpha_n)_n$ be an adapted integrable process on $(\Omega, \mathscr{A}, (\mathscr{F}_n), \mathbf{P})$, and $s_1 \leq s_2$ two stopping times (with values in $\mathbf{N} \cup \{+\infty\}$). We say that $(\alpha_n)_n$ is a submartingale between s_1 and s_2 if, for every $n \geq 0$, the inequality $\mathbf{E}[\alpha_{n+1}|\mathscr{F}_n] \geq \alpha_n$ holds $\mathbf{P}$-a.s. on the event $\{s_1 \leq n < s_2\}$. The process $(\alpha_n)_n$ is a submartingale up to s_2 if it is a submartingale between 0 and s_2. It is straightforward to adapt the sampling theorem as follows. Let (α_n) be a submartingale between s_1 and s_2. Let s be a stopping time, with $\mathbf{P}$-a.s. finite values, such that $s \leq s_2$. Denote by $\mathscr{F}_{s_1}$ the σ-algebra of events known at stage s_1. Then one has $\mathbf{E}[\alpha_s|\mathscr{F}_{s_1}] \geq \alpha_{s_1}$, $\mathbf{P}$-a.s. on the event $\{s_1 \leq s\}$.

Let $(\mathbf{x}, \mathbf{y})$ be a pair of strategies and t the induced stopping stage defined by (2). We define $(\tilde{\alpha}_n)$ as $\tilde{\alpha}_n = \alpha_n$ on $\{t \geq n\}$ and $\tilde{\alpha}_n = r(\mathbf{x}, \mathbf{y})$ if $t < n$. The process $(\tilde{\alpha}_n)$ depends on $(\mathbf{x}, \mathbf{y})$. To avoid ambiguity, we will sometimes write: under $(\mathbf{x}, \mathbf{y})$, the process $(\tilde{\alpha}_n)$ *etc*, when we wish to emphasize which strategies are being used in the definition of $(\tilde{\alpha}_n)$. With a (convenient) abuse of terminology, we refer to $(\tilde{\alpha}_n)$ as the process (α_n) stopped at t.

We use repeatedly the following relation, which holds $\mathbf{P}$-a.s. on the event $\{t \geq n\}$:

$$\mathbf{E}[\tilde{\alpha}_{n+1}|\mathscr{F}_n] = G_n\left(x_n, y_n; \mathbf{E}\left[\alpha_{n+1}|\mathscr{F}_n\right]\right) \tag{7}$$

if (X_n, Y_n) is independent of α_{n+1}. This latter independence property holds in all cases of interest, for instance if $\alpha_{n+1} = v_{n+1}$ or $\alpha_{n+1} = v_{n+1}(\lambda)$, so that we shall apply (7) without further justification.

Set $\mathscr{F}_n^2 = \sigma(\mathscr{F}_n, Y_n)$, so that $\mathscr{F}_n^2$ includes past and present values of the payoff processes, past "decisions" of the players and the decision of player 2 at stage n.

Lemma 3.5. *Let* $\mathbf{y}$ *be a strategy of player 2, and* $\lambda \in]0, 1[$. *Under* $(\mathbf{x}(\lambda), \mathbf{y})$, $((1-\lambda)^n \tilde{v}_n(\lambda))_n$ *is a submartingale up to* $t+1$. *Under* $(\mathbf{x}^*, \mathbf{y})$, $(\tilde{v}_n)_n$ *is a submartingale, both for* $(\mathscr{F}_n)$ *and* $(\mathscr{F}_n^2)_n$.

Notice that $\sup_n |\tilde{v}_n(\lambda)|$ and $\sup_n |\tilde{v}_n|$ belong to $L^1(\mathbf{P})$, for every choice of $(\mathbf{x}, \mathbf{y})$.

Proof. Let $n \geq 0$. On the event $\{t \geq n\}$,

$$\mathbf{E}[(1-\lambda)\tilde{v}_{n+1}(\lambda)|\mathscr{F}_n] = (1-\lambda)G_n(x_n(\lambda), y_n; \mathbf{E}[v_{n+1}(\lambda)|\mathscr{F}_n]),$$

which is at least $v_n(\lambda)$, by (6). This proves the first claim since $\tilde{v}_n(\lambda) = v_n(\lambda)$ if $t \geq n$.

For a similar reason, using (5),

$$\mathbf{E}[\tilde{v}_{n+1}|\mathscr{F}_n] \geq \tilde{v}_n,$$

on the event $\{t \geq n\}$. On $\{t < n\}$, $\tilde{v}_{n+1} = \tilde{v}_n$. The same computation works also for the filtration $(\mathscr{F}^2_n)_n$. This completes the proof. □

Corollary 3.6. *For every* $\mathbf{y}$, $\gamma_\lambda(\mathbf{x}(\lambda), \mathbf{y}) \geq \mathbf{E}(v_0(\lambda))$.

Proof. Fix a strategy $\mathbf{y}$ of player 2. Let $n \geq 0$, and apply the submartingale property with the stopping time $\min(t+1, n)$:

$$\mathbf{E}[(1-\lambda)^{\min(t+1,n)}\tilde{v}_{\min(t+1,n)}] \geq \mathbf{E}(v_0(\lambda)),$$

that is, using the definition of the stopped process $(\tilde{v}_n)_n$:

$$\mathbf{E}[(1-\lambda)^n v_n(\lambda)1_{t\geq n} + (1-\lambda)^{t+1} r(\mathbf{x}(\lambda), \mathbf{y})1_{t<n}] \geq \mathbf{E}(v_0(\lambda)).$$

By dominated convergence, the left-hand side converges to $\gamma_\lambda(\mathbf{x}(\lambda), \mathbf{y})$. □

A similar proof proves the following.

Corollary 3.7. *Let* $n \in \mathbf{N}$. *Let* $\tilde{\mathbf{x}}(\lambda)$ *be the strategy that is identically 0 until stage* n, *and coincides with* $\mathbf{x}(\lambda)$ *afterwards. Let* $\mathbf{y}$ *be any strategy of player 2 that is identically 0 until stage* n. *Then*

$$\mathbf{E}[(1-\lambda)^{t+1-n} r(\mathbf{x}(\lambda), \mathbf{y})|\mathscr{F}_n] \geq v_n(\lambda).$$

Corollary 3.6 implies that in the discounted game it is an optimal strategy for player 1 to play $\mathbf{x}(\lambda)$. No such result holds for the original problem: playing $\mathbf{x}^*$ needs not be an optimal strategy.

Example

	1
1	0

This matrix notation is a shortcut for the stopping game with payoffs $a_n = b_n = 1$, $c_n = 0$, $\mathbf{P}$-a.s. for every n. Clearly v_n and $v_n(\lambda)$ are independent of n and constant, so we simply write v and $v(\lambda)$. The real number $0 \leq v(\lambda) \leq 1$ is a solution to the equation $v(\lambda) = (1-\lambda)\mathrm{val}\, G(v(\lambda))$, from which it is easily derived

$v(\lambda) = 1 - \sqrt{\lambda}$, and $\mathbf{x}(\lambda) = \sqrt{\lambda}/(1 + \sqrt{\lambda})$. Therefore $v = 1$. Denote by $\mathbf{0}$ the strategy (of either player 1 or player 2) that never stops ($0_n = 0$ for all n). Then $\mathbf{x}^* = \mathbf{0}$. However, $\gamma(\mathbf{x}^*, \mathbf{0}) = 0$.

Nevertheless, if t_1 is $\mathbf{P}$-a.s. finite under $\mathbf{x}^*$, then $\mathbf{x}^*$ is optimal for player 1.

Lemma 3.8. *If $\mathbf{P}(t_1 < +\infty) = 1$ under $\mathbf{x}^*$, then $\mathbf{x}^*$ guarantees w for player 1.*

Proof. Let $\mathbf{y}$ be an arbitrary strategy of player 2. By Lemma 3.5, $(\tilde{v}_n)$ is a submartingale under $(\mathbf{x}^*, \mathbf{y})$. Since $\mathbf{P}(t_1 < +\infty) = 1$, $\mathbf{P}(t < +\infty) = 1$ as well, hence it follows that

$$\mathbf{E}[r(\mathbf{x}^*, \mathbf{y})1_{t<+\infty}] = \mathbf{E}[\tilde{v}_\infty] \geq \mathbf{E}[v_0] = w,$$

as desired. □

4. The main ideas of the proofs

We give a detailed sketch of the proofs in the deterministic case. Many technical issues disappear in that case. Therefore the main ideas appear, hopefully more clearly. Assume that $(a_n)_n$, $(b_n)_n$, $(c_n)_n$, and therefore also $(v_n)_n$ and $(v_n(\lambda))_n$, are sequences of real numbers, bounded by 1.

For every $\mathbf{y}$, $(\tilde{v}_n)_n$ is a bounded submartingale under $(\mathbf{x}^*, \mathbf{y})$, thus

$$\mathbf{E}\left[\tilde{v}_\infty\right] \geq \mathbf{E}[v_0] = w \tag{8}$$

with $\tilde{v}_\infty = \lim_n \tilde{v}_n$.

For $\mathbf{y} = \mathbf{0}$, $(\tilde{v}_n)$ coincides with v_n up to t_1. Thus, $t_1 < +\infty$, $\mathbf{P}$-a.s., or $(v_n)_n$ is a convergent sequence. In the first case, $\mathbf{x}^*$ is optimal by (8).

We now assume that (v_n) is a convergent sequence, and given $\varepsilon > 0$, we choose N_0 such that $\sup_{n,m \geq N_0} |v_n - v_m| \leq \varepsilon/2$. We also assume for simplicity $N_0 = 0$ (in the general case, the strategies below would be supplemented by: play $\mathbf{x}^*$ up to N_0). If $w \leq \varepsilon$, $\tilde{v}_\infty \leq 3\varepsilon/2$, so that $\mathbf{x}^*$ is $3\varepsilon/2$-optimal by (8). We are thus led to consider the case $w > \varepsilon$.

First proof. Choose λ_0 such that $v(\lambda_0) \geq w - \varepsilon/3$ and $\varepsilon' \in (0, \varepsilon/6)$. Player 1 starts playing according to $\mathbf{x}(\lambda_0)$. For each $\mathbf{y}$, $((1 - \lambda_0)^n \tilde{v}_n(\lambda_0))_n$ is a submartingale up to t. Set $s_1 = \inf\{n, v_n(\lambda_0) \leq \varepsilon'\}$. Since $(\tilde{v}_n(\lambda_0))_n$ is bounded, $\min(t, s_1)$ is $\mathbf{P}$-a.s. finite. Moreover, since $v_{s_1}(\lambda_0) \leq v(\lambda_0) - (\varepsilon/6 - \varepsilon')$ if $s_1 \leq t$, the probability that $t < s_1$ is bounded away from 0.

At stage s_1, the approximation of (v_n) by $(v_n(\lambda_0))_n$ gets poor, so we switch to a new discount factor: λ_0 is replaced by λ_1, with $v_{s_1}(\lambda_1) \geq v_{s_1} - \varepsilon/3 \geq \varepsilon/6$, and $\mathbf{x}(\lambda_1)$ is played until $s_2 = \inf\{n > s_1, v_n(\lambda_1) \leq \varepsilon'\}$, where we again switch from λ_1 to λ_2, and so on.

Call $\bar{\mathbf{x}}$ the resulting strategy. Under $(\bar{\mathbf{x}}, \mathbf{y})$, t is $\mathbf{P}$-a.s. finite, since for every n, the probability of stopping between s_n and s_{n+1} is bounded away from 0. Introduce the sequence $(w_n)_n$, where $w_n = v_n(\lambda_p)$ if $s_p \leq n < s_{p+1}$. By construction, $w_0 \geq v - \varepsilon/3$ and $(\tilde{w}_n)_n$ is a submartingale. Since $t < +\infty$, it converges to $r(\overline{\mathbf{x}}, \mathbf{y})1_{t<+\infty}$, therefore $\gamma(\overline{\mathbf{x}}, \mathbf{y}) \geq w - \varepsilon/3$.

Second proof. The definition of $\overline{\mathbf{x}}$ here is motivated by the observation

$$\limsup_n a_n \geq w - \varepsilon \tag{9}$$

which is derived as follows. For each λ, under $(\mathbf{x}(\lambda), \mathbf{0})$, $t = t_1$ and $r(\mathbf{x}(\lambda), \mathbf{0}) = a_t$ if $t < +\infty$; thus,

$$\mathbf{E}\left[(1-\lambda)^{t+1} a_t 1_{t<+\infty}\right] = \gamma_\lambda(\mathbf{x}(\lambda), \mathbf{0}) \geq v(\lambda).$$

The left-hand side lies in the closed convex hull of $\{0, a_n, n \in \mathbf{N}\}$. Given any $\delta > 0$, $v(\lambda) \geq w - \delta$, for a suitable λ. Therefore, $\sup_n a_n \geq w - \delta$. Since $v_n \geq w - \varepsilon$ for every n, this proof may be repeated, and (9) holds.

We define $\overline{\mathbf{x}}$ by $\overline{x}_n = x_n^* + \varepsilon$ if $a_n \geq w - 2\varepsilon$, and $\overline{x}_n = x_n^*$ otherwise. Since (9) holds, $t_1 < +\infty$ $\mathbf{P}$-a.s. under $\overline{\mathbf{x}}$. To see that this strategy guarantees player 1 an expected payoff of w, we note that the following points hold:

1. If player 2 stops the game ($t = t_2$), then the expected payoff of player 1 is at least w (up to an ϵ).
2. In the case that player 2 always continues, since player 1 changes his strategy *only* when a unilateral stopping is favorable for him, $\mathbf{E}[v_n] \geq w - \epsilon$.

5. Two ϵ-optimal strategies

5.1. Preliminaries

For the rest of the section we fix $\epsilon > 0$. Set $m = \sup_n(\sup(|a_n|, |b_n|, |c_n|))$. Since $m \in L^1(\mathbf{P})$, there exists $\eta > 0$ such that, for every $A \in \mathscr{A}$,

$$\mathbf{P}(A) < \eta \Rightarrow \mathbf{E}(m 1_A) < \epsilon. \tag{10}$$

Notice that $|v_n(\lambda)|, |v_n| \leq \mathbf{E}[m \mid \mathscr{F}_n]$, $\mathbf{P}$-a.s. for every n.

The sequence (v_n) needs not converge. On the other hand, the process $(\tilde{v}_n)$, being a submartingale under $(\mathbf{x}^*, \mathbf{y})$ (with $\sup \tilde{v}_n \in L^1(\mathbf{P})$) converges $\mathbf{P}$-a.s. and in $L^1(\mathbf{P})$, for every $\mathbf{y}$.

The stopping time t_1 is a function of player 1's strategy. Under $(\mathbf{x}^*, \mathbf{0})$, $t = t_1$, $\mathbf{P}$-a.s. This implies that (v_n) converges $\mathbf{P}$-a.s. on the set $\{t_1 = +\infty\}$.

Choose $N_0 \in \mathbf{N}$ such that

$$\mathbf{P}\{\sup_{n,m \geq N_0} |v_n - v_m| > \epsilon/2, t_1 \geq N_0\} < \eta. \tag{11}$$

Thus, after stage N_0, with high probability v_n does not change by much.

5.2. An ϵ-optimal strategy for player 1 – I

We first define the switching stages (s_p) and the approximating discount factors (λ_p): $v(\lambda_p)$ approximates v between s_p and s_{p+1}. Set $s_0 = N_0$ if $v_{N_0} > \epsilon$, and $s_0 = +\infty$ otherwise. Choose $\varepsilon' \in (0, \varepsilon/6)$ and an $\mathscr{F}_{s_0}$-measurable function λ_0 with $v_{s_0}(\lambda_0) > v_{s_0} - \epsilon/3$ if $s_0 < +\infty$.

Set $s_{p+1} = \inf\{n > s_p, v_n(\lambda_p) \le \varepsilon'\}$ and choose an $\mathscr{F}_{s_{p+1}}$-measurable function λ_{p+1}, such that $v_{s_{p+1}}(\lambda_{p+1}) > v_{s_{p+1}} - \epsilon/3$ if $s_{p+1} < +\infty$.

Let $\overline{\mathbf{x}}$ be the strategy that coincides with $\mathbf{x}^*$ until s_0, and with $\mathbf{x}(\lambda_p)$ between s_p and s_{p+1}:

$$\overline{x}_n = \begin{cases} x_n^* & n < s_0 \\ x_n(\lambda_p) & s_p \le n < s_{p+1} \end{cases}$$

We shall prove that $\overline{\mathbf{x}}$ is 7ε-optimal.

By Lemma 3.5, for every $\mathbf{y}$, $(\tilde{v}_n)_n$ is a submartingale up to s_0, and $((1-\lambda_p)^n \tilde{v}_n(\lambda_p))_n$ is a submartingale between $\min(s_p, t+1)$ and $\min(s_{p+1}, t+1)$, for each p.

We introduce an auxiliary variable z_n defined as

$$z_n = \begin{cases} v_n - \epsilon/3 & n < s_0 \\ v_n(\lambda_p) & s_p \le n < s_{p+1} \end{cases}$$

Intuitively, z_n is (up to $\epsilon/3$), the parameter we are interested in: the limit v_n before stage s_0, and the λ_p-discounted value for $s_p \le n < s_{p+1}$.

We ultimately wish to get a submartingale. A minor adjustment is needed. Define the stopping time s by $s = +\infty$ if $s_0 = +\infty$ and $s = \inf\{n \ge N_0, v_n \le \varepsilon/2\}$ otherwise. By the definition of N_0, $\mathbf{P}(s < +\infty, t_1 \ge N_0) < \eta$. We use s to define a process (w_n) by

$$w_n = \begin{cases} \mathbf{E}\left[m|\mathscr{F}_n\right] & s \le n \\ z_n & \text{otherwise} \end{cases}$$

Observe that

$$\tilde{w}_{n+1} \ge \tilde{v}_{n+1}(\lambda_p) \text{ on the event } \{s_p \le n < s_{p+1}\}. \tag{12}$$

Indeed, this is clear if $s \le n+1$ or if $t < n+1$. If not :

$$\tilde{w}_{n+1} = \begin{cases} v_{n+1}(\lambda_p) & n+1 < s_{p+1} \\ v_{n+1}(\lambda_{p+1}) & n+1 = s_{p+1} \end{cases}$$

If $n+1 < s_{p+1}$, then $\tilde{w}_{n+1} = \tilde{v}_{n+1}(\lambda_p)$, while if $n+1 = s_{p+1}$,

$$\tilde{w}_{n+1} \ge v_{n+1} - \epsilon/3 \ge \varepsilon' \ge v_{n+1}(\lambda_p).$$

We set $\overline{t}+1 = \min(t+1, s)$. Observe that $\mathbf{P}(t = \overline{t}) \ge 1 - \eta$.

Lemma 5.1. *For every* $\mathbf{y}$, $(\tilde{w}_n)$ *is a submartingale up to* $\overline{t}+1$ *under* $(\overline{\mathbf{x}}, \mathbf{y})$.

Proof. Fix a strategy $\mathbf{y}$ of player 2. Let $n \in \mathbf{N}$. We prove that $\mathbf{E}[\tilde{w}_{n+1}|\mathscr{F}_n] \geq \tilde{w}_n$, $\mathbf{P}$-a.s. on the event $\{\bar{t}+1 > n\}$.

If $n < s_0$, $w_n = v_n - \epsilon/3$, $w_{n+1} \geq v_{n+1} - \epsilon/3$ (with equality if $n+1 < s_0$), and $\bar{x}_n = x_n^*$. Thus $\mathbf{E}[\tilde{w}_{n+1}|\mathscr{F}_n] \geq G_n(x_n^*, y_n; \mathbf{E}[v_{n+1} - \epsilon/3|\mathscr{F}_n]) \geq v_n - \epsilon/3$, where the second inequality follows from the inequality $G_n(x_n^*, y_n; \mathbf{E}[v_{n+1}|\mathscr{F}_n]) \geq v_n$ and since the val operator is non-expansive.

If $s_p \leq n < s_{p+1}$, $w_n = v_n(\lambda_p)$, and $\bar{x}_n = x_n(\lambda_p)$. In that case, by (12),

$$\mathbf{E}(\tilde{w}_{n+1}|\mathscr{F}_n) \geq G_n(\bar{x}_n, y_n; \mathbf{E}[v_{n+1}(\lambda_p)|\mathscr{F}_n]) \geq \frac{1}{1-\lambda_p} v_n(\lambda_p) \geq v_n(\lambda_p) = w_n,$$

where the last inequality holds since $v_n(\lambda_p) > 0$. □

Lemma 5.2. *For every* $\mathbf{y}$*, under* $(\bar{\mathbf{x}}, \mathbf{y})$*,* $\bar{t} < +\infty$*,* $\mathbf{P}$*-a.s. on the event* $s_0 = N_0$*.*

Proof. Fix a strategy $\mathbf{y}$ of player 2. We proceed in two steps. We prove first that $\min(s_{p+1}, t) < +\infty$, $\mathbf{P}$-a.s. on $\{s_p < s\}$. From $\min(s_p, t+1)$ up to $\min(s_{p+1}, t+1)$, $((1-\lambda_p)^n \tilde{w}_n)$ is a submartingale. Thus, for every $N \in \mathbf{N}$ and $n \leq N$, the sampling property applied to the finite stopping time $\min(s_{p+1}, t+1, N)$ yields

$$\begin{aligned} w_n \leq \frac{1}{(1-\lambda_p)^n} \mathbf{E}\Big[& m(1-\lambda_p)^{\min(s_{p+1},t+1)} 1_{\min(s_{p+1},t+1) \leq N} \\ & + \tilde{w}_N (1-\lambda_p)^N 1_{\min(s_{p+1},t+1) > N} |\mathscr{F}_n \Big] \end{aligned}$$

on $\{s_p \leq n < \min(s_{p+1}, t+1)\}$.

By taking $N \to +\infty$ and by dominated convergence for conditional expectations, one obtains

$$\varepsilon' < v_n(\lambda_p) = w_n \leq \mathbf{E}\left[m(1-\lambda_p)^{\min(s_{p+1},t+1)-n} 1_{\min(s_{p+1},t+1)<+\infty} |\mathscr{F}_n \right] \quad (13)$$

on the event $\{s_p \leq n < \min(s_{p+1}, t+1)\}$.

By taking the limit $n \to \infty$ in (13), one gets $\limsup w_n \leq 0$, $\mathbf{P}$-a.s. on the event $\{s_p < +\infty, t = s_{p+1} = +\infty\} \cap \{s_p < s\}$. But on this event $w_n \geq \varepsilon'$, $\mathbf{P}$-a.s. for every n. This ends the first step.

One can rephrase the conclusion of the first step as $\min(s_{p+1}, \bar{t}) < +\infty$ if $\min(s_p, \bar{t}) < +\infty$, $\mathbf{P}$-a.s. By induction, $\min(s_p, \bar{t}) < +\infty$ if $s_0 < +\infty$, $\mathbf{P}$-a.s. for every p.

Since $(\tilde{v}_n(\lambda_p))_n$ is a submartingale between $\min(s_p, \bar{t}+1)$ and $\min(s_{p+1}, \bar{t}+1)$, and since $v_{s_{p+1}}(\lambda_p) \leq \varepsilon'$,

$$v_{s_p}(\lambda_p) \leq \mathbf{E}[m 1_{\bar{t}+1 \leq s_{p+1}} + \varepsilon' \cdot 1_{s_{p+1} < \bar{t}+1} |\mathscr{F}_{s_p}]$$

on $\{s_p < \bar{t}+1\}$. Since $v_{s_p}(\lambda_p) \geq \epsilon/6$, it follows by taking expectations that

$$\frac{\epsilon}{6} \mathbf{P}(s_p < \bar{t}+1) \leq \mathbf{E}(m 1_{s_p < \bar{t}+1 < +\infty}) + \varepsilon' \mathbf{P}(s_{p+1} < \bar{t}+1),$$

hence

$$\left(\frac{\epsilon}{6}-\varepsilon'\right)\mathbf{P}(s_p<\bar{t}+1)\le \mathbf{E}(m1_{s_p<\bar{t}+1<+\infty})$$

As p goes to infinity, the left-hand side converges to $(\epsilon/6-\varepsilon')\mathbf{P}(s_0=N_0,\bar{t}=+\infty)$, while the right-hand side converges to 0. The result follows. □

Proposition 5.3. *There exists $N\in\mathbf{N}$ such that, for every $\mathbf{y}$ and $n\ge N$, one has $\gamma_n(\overline{\mathbf{x}},\mathbf{y})\ge w-7\epsilon$.*

Proof. By Lemma 5.2 with $\mathbf{y}=\mathbf{0}$, there exists some positive integer $N_1\ge N_0$ such that under $(\overline{\mathbf{x}},\mathbf{0})$

$$\mathbf{P}(s_0=N_0,t\ge N_1)<\eta. \tag{14}$$

This readily implies that (14) holds under $(\overline{\mathbf{x}},\mathbf{y})$, for *every* $\mathbf{y}$.

Let now N_2 be sufficiently large such that

$$\frac{N_1}{N_2}\mathbf{E}[m]<\epsilon. \tag{15}$$

Using (14), (10) and (15) we have:

$$\left|\mathbf{E}\left[\frac{1}{n+1}\sum_{k=0}^{n}g_k\right]-\mathbf{E}\left[\frac{1}{n-N_1+1}\sum_{k=N_1}^{n}g_k\right]\right|<2\varepsilon,\text{ provided } n\ge N_2. \tag{16}$$

Fix $n\ge N_2$ and any strategy $\mathbf{y}$.

By definition, $\gamma_{n+1}(\overline{\mathbf{x}},\mathbf{y})=\mathbf{E}[\frac{1}{n+1}\sum_{k=0}^{n}g_k]$. We will evaluate $\mathbf{E}\left[\frac{1}{n-N_1+1}\sum_{k=N_1}^{n}g_k\right]$. Let $N_1\le k\le n$.

On $\{t<N_1\}$, $g_k=\tilde{w}_{N_1}$.

On $\{t\ge N_1,s_0=N_0\}$, $|g_k|\le m$, but this event has a probability at most 2η.

Consider now the event $\{t\ge N_1,s_0=+\infty\}$. The event $\{t\ge N_1,s_0=+\infty,\sup_{q\ge N_1}v_q>3\varepsilon/2\}$ has probability at most η. On the event $\{t\ge N_1,s_0=+\infty,\sup_{q\ge N_1}v_q\le 3\varepsilon/2\}$, $g_k=\tilde{w}_k$ if $k>t$, while $g_k=0\ge\tilde{w}_k-3\epsilon/2$ if $k\le t$. Therefore,

$$\begin{aligned}\mathbf{E}\left[g_k1_{t\ge N_1,s_0=+\infty}\right]&\ge\mathbf{E}\left[\tilde{w}_k1_{t\ge N_1,s_0=+\infty}\right]-3\varepsilon/2-\varepsilon\\&\ge\mathbf{E}\left[\tilde{w}_{N_1}1_{t\ge N_1,s_0=+\infty}\right]-5\varepsilon/2,\end{aligned}$$

where the second inequality uses the fact that $\{t\ge N_1,s_0=+\infty\}\in\mathscr{F}_{N_1}$, and the submartingale property of $(\tilde{w}_n)_n$.

Thus,

$$\mathbf{E}\left[g_k\right]\ge\mathbf{E}\left[\tilde{w}_{N_1}\right]-5\varepsilon/2-2\varepsilon\ge w-9\varepsilon/2,$$

where the second inequality uses $w=\mathbf{E}\left[w_0\right]$ and the submartingale property of $(\tilde{w}_n)_n$. The result follows from (16). □

5.3. An ϵ-optimal strategy for player 1 – II

By Lemma 3.8, if $\mathbf{P}(t_1 < +\infty) = 1$ under $\mathbf{x}^*$, then $\mathbf{x}^*$ guarantees w for player 1. Therefore, we assume from now on that under $\mathbf{x}^*$

$$\mathbf{P}(t_1 < +\infty) < 1. \tag{17}$$

Recall that $\varepsilon > 0$ is given, and that $\eta > 0$ is such that

$$\mathbf{P}(A) < \eta \Rightarrow \mathbf{E}[m1_A] < \varepsilon. \tag{18}$$

Assume moreover that $\eta\mathbf{E}[m] \leq \varepsilon$.

Recall also that N_0 is such that, under $(\mathbf{x}^*, \mathbf{0})$,

$$\mathbf{P}\left(\sup_{n,m\geq N_0} |v_n - v_m| > \varepsilon/2, t_1 \geq N_0\right) < \eta. \tag{19}$$

By (17) we can assume w.l.o.g. that N_0 is sufficiently large so that under $\mathbf{x}^*$,

$$\mathbf{P}(t_1 < +\infty \mid t_1 \geq N_0) < \eta.$$

Define the strategy $\hat{\mathbf{x}}$ by

$$\hat{x}_n = \begin{cases} \min\{x_n^* + \eta, 1\} & \text{if } n \geq N_0 \text{ and } \epsilon < v_{N_0} < a_n + \epsilon \\ x_n^* & \text{otherwise.} \end{cases}$$

We will prove that $\gamma(\hat{\mathbf{x}}, \mathbf{y}) \geq w - 9\varepsilon$, for every $\mathbf{y}$. The stronger statement: $\gamma_n(\hat{\mathbf{x}}, \mathbf{y}) \geq v - 6\varepsilon$, for every $n \geq N_1$ and every $\mathbf{y}$ also holds, provided N_1 is large enough. We will not provide a proof.

Lemma 5.4. *One has*

$$\limsup_n a_n \geq \limsup_n v_n \textit{ on the event } \{\limsup_n v_n > 0\}.$$

Proof. Let $\lambda > 0$ and $q \in \mathbf{N}$ be given. Denote by $\tilde{\mathbf{x}}(\lambda)$ the strategy that coincides with $\mathbf{0}$ for $n < q$, and with $\mathbf{x}(\lambda)$ for $n \geq q$. From Corollary 3.7, under $(\tilde{\mathbf{x}}(\lambda), \mathbf{0})$,

$$(1-\lambda)^{-q}\mathbf{E}\left[r(\tilde{\mathbf{x}}(\lambda), \mathbf{0})(1-\lambda)^{t+1}1_{t<+\infty}|\mathscr{F}_q\right] \geq v_q(\lambda). \tag{20}$$

Since player 2 never stops, $t = t_1$ and $r = a_{t_1}$ on $t < +\infty$. Since $(1-\lambda)^{t+1-q} \leq 1$, $\mathbf{P}$-a.s., the left-hand side of (20) is at most

$$\mathbf{E}\left[a_t^+ 1_{t<+\infty}|\mathscr{F}_q\right] \leq \mathbf{E}\left[\sup_{n\geq q} a_n^+|\mathscr{F}_q\right],$$

with $a_n^+ = \max(a_n, 0)$. Using (20),

$$\mathbf{E}\left[\sup_{n\geq q} a_n^+|\mathscr{F}_q\right] \geq v_q(\lambda).$$

By letting λ go to zero, one obtains $\mathbf{E}\left[\sup_{n\geq q} a_n^+ | \mathscr{F}_q\right] \geq v_q$. The sequence $(\mathbf{E}\left[\sup_{n\geq q} a_n^+ | \mathscr{F}_q\right])_q$ converges $\mathbf{P}$-a.s. to $\limsup_q a_q^+$. Therefore $\limsup_q a_q^+ \geq \limsup_q v_q$. Thus, on the event $\{\limsup_n v_n > 0\}$, $\limsup_n a_n \geq \limsup_n v_n$, as desired. □

Set $\Omega_1 = \{t \geq N_0, v_{N_0} > \varepsilon\} \in \mathscr{F}_{N_0}$.

Proposition 5.5. *Let* $\mathbf{y}$ *be given. One has*

$$\mathbf{E}\left[r(\hat{\mathbf{x}}, \mathbf{y}) 1_{\Omega_1} 1_{t_2 = t < +\infty}\right] \geq \mathbf{E}\left[v_{N_0} 1_{\Omega_1} 1_{t_2 = t < +\infty}\right] - 4\varepsilon$$

under $(\hat{\mathbf{x}}, \mathbf{y})$.

Proof. We explicit the idea that, if player 2 stops at stage n, the corresponding expected payoff (where the expectation is taken with respect to player 1's decision) is at least v_n, up to ηm, since player 1 plays x_n^* up to η.

Recall that $\mathscr{F}_n^2 = \sigma(\mathscr{F}_n, Y_n)$, so that $\mathscr{F}_n^2$ includes past and present values of the payoff processes, past "decisions" of the players and the decision of player 2 at stage n. Observe that $\{t_2 = t = n\} \in \mathscr{F}_n^2$, and that by assumption X_n is independent of $\mathscr{F}_n^2$. Therefore, on the event $\{t_2 = t = n\}$,

$$\mathbf{E}\left[r(\hat{\mathbf{x}}, \mathbf{y}) | \mathscr{F}_n^2\right] = G_n\left(\hat{x}_n, 1; \mathbf{E}\left[v_{n+1} | \mathscr{F}_n\right]\right)$$

(Note that the variable $\mathbf{E}\left[v_{n+1} | \mathscr{F}_n\right]$ is here irrelevant). Since x_n is an optimal strategy in the local game $G_n(\mathbf{E}\left[v_{n+1} | \mathscr{F}_n\right])$, by Lemma 3.4,

$$G_n\left(x_n, 1; \mathbf{E}\left[v_{n+1} | \mathscr{F}_n\right]\right) \geq \operatorname{val} G_n\left(\mathbf{E}\left[v_{n+1} | \mathscr{F}_n\right]\right) \geq v_n.$$

Since $|x_n - \hat{x}_n| \leq \eta$,

$$\left|G_n\left(x_n, 1; \mathbf{E}\left[v_{n+1} | \mathscr{F}_n\right]\right) - G_n\left(\hat{x}_n, 1; \mathbf{E}\left[v_{n+1} | \mathscr{F}_n\right]\right)\right| \leq \eta m,$$

so that $\mathbf{E}\left[r(\hat{\mathbf{x}}, \mathbf{y}) | \mathscr{F}_n^2\right] \geq v_n - \eta m$ on the event $\{t_2 = n = t\}$. In other words,

$$\mathbf{E}\left[r(\hat{\mathbf{x}}, \mathbf{y}) 1_{t_2 = n = t} | \mathscr{F}_n^2\right] \geq (v_n - \eta m) 1_{t_2 = n = t}, \mathbf{P}\text{-a.s.}$$

By first taking conditional expectations given $\mathscr{F}_{N_0}$, and then summing over $n \geq N_0$, one obtains

$$\begin{aligned}\mathbf{E}\left[r(\hat{\mathbf{x}}, \mathbf{y}) 1_{N_0 \leq t_2 = t < +\infty} | \mathscr{F}_{N_0}\right] \geq\ & \mathbf{E}\left[\inf_{n \geq N_0} v_n 1_{t_2 = t = n} | \mathscr{F}_{N_0}\right] \\ & - \eta \mathbf{E}\left[m 1_{N_0 \leq t_2 = t < +\infty} | \mathscr{F}_{N_0}\right],\end{aligned}$$

which yields

$$\mathbf{E}\left[r(\hat{\mathbf{x}}, \mathbf{y}) 1_{\Omega_1} 1_{t_2 = t < +\infty}\right] \geq \mathbf{E}\left[1_{\Omega_1} \inf_{n \geq N_0} v_n 1_{t_2 = t = n}\right] - \eta \mathbf{E}\left[m 1_{\Omega_1} 1_{t_2 = t < +\infty}\right].$$

Define $\Omega_2 = \Omega_1 \cap \{\sup_{n,m \geq N_0} |v_n - v_m| \leq \varepsilon/2\}$. Thus, $\mathbf{P}(\Omega_1 \setminus \Omega_2) < \eta$, therefore

$$\left|\mathbf{E}\left[1_{\Omega_1} \inf_{n \geq N_0} v_n 1_{t_2=t=n}\right] - \mathbf{E}\left[1_{\Omega_2} \inf_{n \geq N_0} v_n 1_{t_2=t=n}\right]\right| \leq \mathbf{E}\left[1_{\Omega_1 \setminus \Omega_2} m\right] \leq \varepsilon.$$

On Ω_2, $\inf_{n \geq N_0} v_n \geq v_{N_0} - \varepsilon/2$. One finally gets

$$\begin{aligned}\mathbf{E}\left[r(\hat{\mathbf{x}}, \mathbf{y}) 1_{\Omega_1} 1_{t_2=t<+\infty}\right] \geq \mathbf{E}&\left[v_{N_0} 1_{\Omega_1} 1_{t_2=t<+\infty}\right] \\ &-\frac{\varepsilon}{2}\mathbf{P}(\Omega_1 \cap \{t_2 = t < +\infty\}) - 3\varepsilon. \qquad (21)\end{aligned}$$

Proposition 5.6. *Let* $\mathbf{y}$ *be given. One has*

$$\mathbf{E}\left[r(\hat{\mathbf{x}}, \mathbf{y}) 1_{\Omega_1} 1_{t_1<t_2}\right] \geq \mathbf{E}\left[v_{N_0} 1_{\Omega_1} 1_{t_1<t_2}\right] - 2\varepsilon.$$

under $(\hat{\mathbf{x}}, \mathbf{y})$.

Proof. Fix a strategy $\mathbf{y}$. Note that $\Omega_1 \cap \{t_1 < t_2\} = \{N_0 \leq t_1 < t_2, v_{N_0} > \epsilon\}$, and on this set, $r(\hat{\mathbf{x}}, \mathbf{y}) = a_{t_1}$.

By the definition of N_0, $\mathbf{P}(t_1 = t \geq N_0, a_t \leq v_{N_0} - \epsilon) < \eta$. In particular, $\mathbf{P}(N_0 \leq t_1 < t_2, a_{t_1} > v_{N_0} - \epsilon > 0) > \mathbf{P}(\Omega_1 \cap \{t_1 < t_2\}) - \eta$. The result follows from (18).

Lemma 5.7. *For every* $\mathbf{y}$*,* $\gamma(\hat{\mathbf{x}}, \mathbf{y}) \geq w - 9\epsilon$.

Proof. Define the stopping time θ by $\theta = N_0$ on $\Omega_1 = \{t \geq N_0, v_{N_0} > \epsilon\}$, and $\theta = +\infty$ otherwise. The strategy $\hat{\mathbf{x}}$ coincides with $\mathbf{x}^*$ up to θ. Therefore, $(\tilde{v}_n)$ is a submartingale up to θ.

Notice that $\theta = +\infty$ if $\theta > N_0$; therefore $(\tilde{v}_n)$ converges, $\mathbf{P}$-a.s. on the event $\{\theta > N_0\}$, say to $\tilde{v}_\infty$.

Given the integrability properties of $(\tilde{v}_n)$, one has

$$\mathbf{E}(\tilde{v}_\theta) \geq \mathbf{E}(\tilde{v}_0) = w. \qquad (22)$$

By definition of $(\tilde{v}_n)$, one has $\tilde{v}_\infty = r(\hat{\mathbf{x}}, \mathbf{y})$ if $t < +\infty$, $\tilde{v}_\infty \leq 3\epsilon/2$ if $t = +\infty$ and $\sup_{n,m \geq N_0} |v_n - v_m| \leq \epsilon/2$, and $\tilde{v}_\infty \leq m$ otherwise. Thus, by (18) and (19),

$$\mathbf{E}[\tilde{v}_\infty 1_{\theta > N_0}] \leq \mathbf{E}[r(\hat{\mathbf{x}}, \mathbf{y}) 1_{t<+\infty} 1_{\theta > N_0}] + 3\epsilon/2 + \epsilon.$$

The inequality (22) may be rewritten as

$$\mathbf{E}[v_{N_0} 1_{\Omega_1} + \tilde{v}_\infty 1_{\theta > N_0}] \geq w.$$

and therefore, using Propositions 5.5 and 5.6, $\mathbf{E}(r(\hat{\mathbf{x}}, \mathbf{y}) 1_{t<+\infty}) \geq w - \frac{17}{2}\epsilon$, and the result follows. $\square$

6. Concluding remarks

This section contains a discussion of related issues. We first discuss an alternative way of introducing randomization into stopping games[1]. We then discuss a fairly easy extension of our main result.

We introduced randomization by allowing the players, at any stage, to stop with a probability between zero and one. These strategies are usually called *behavior strategies* in the game theory literature. We might as well consider the possibility for a player to select randomly a (deterministic) stopping time at the beginning of the game, thereby extending differently the set of available strategies. These strategies are called *mixed strategies.* For many classes of games, the two extensions are equivalent in a strong sense. The first equivalence result is due to Kuhn (1953).

For stopping games (as for many other games), the definition of mixed strategies as suggested here is problematic, since it requires to define a convenient measurable structure on the set of stopping times. There are two ways to avoid this problem.

Following Aumann (1964), one may enlarge the probability space from $(\Omega, \mathscr{A}, \mathbf{P})$ to $(\Omega \times [0, 1], \mathscr{A} \otimes \mathscr{B}, \mathbf{P} \otimes \lambda_1)$, where λ_1 is the Lebesgue measure. A mixed strategy (for Player 1) is then defined as an $\mathscr{A} \otimes \mathscr{B}$-measurable function ϕ from $\Omega \times [0, 1]$ to $\mathbf{N} \cup \{+\infty\}$ such that

$$\text{for } \lambda_1\text{-a.e. } r \in [0, 1],\ \ \phi(r, \cdot) \text{ is a stopping time.}$$

Intuitively, $([0, 1], \lambda_1)$ is a randomizing device for player 1. We introduce an independent copy $([0, 1], \lambda_2)$ for player 2.

We claim that these mixed strategies are equivalent to behavioral strategies. Denote $\sigma_r = \phi(r, \cdot)$. Then σ_r is λ_1-a.e. a stopping time. For each mixed strategy ϕ and every $n \in \mathbf{N}$, define $H(\phi)_n = \int 1_{\{\sigma_r \leq n\}} \lambda_1(dr)$ the probability under ϕ that player 1 stops prior to stage $n + 1$. Clearly, $(H(\phi)_n)$ is $(\mathscr{F}_n)$-adapted. It can be viewed as the (random) distribution function corresponding to some behavior strategy $\mathbf{x}$, that we denote by $h(\phi)$. The map h from mixed to behavior strategies is onto. Indeed, given a behavior strategy $\mathbf{x}$, denote by $F^{\mathbf{x}}$ the distribution function of t_1. Set $\phi^{\mathbf{x}}(r, \omega) := \inf\{n \geq 0,\ F^{\mathbf{x}}(n, \omega) \geq r\}$. Then $\phi^{\mathbf{x}}$ is a mixed strategy, such that $h(\phi^x) = \mathbf{x}$. It is easy to verify that, for each pair (ϕ, ψ) of mixed strategies, the expected payoff under (ϕ, ψ) coincides with the expected payoff under the pair $(h(\phi), h(\psi))$ of behavior strategies. For more details, see Touzi and Vieille (1999).

Another approach to define mixed strategies is due to Bismut (1977): it consists of interpreting such a strategy as an element of the dual space of a Banach space containing the stopping times, and of using functional analysis methods.

We argue now that the first proof of the main result can be extended to handle a larger class of stochastic games.[2] The class of games we consider now is the following. Each player has finitely many actions. The sets of actions are respectively A and B for the two players. The two players choose repeatedly elements from A and B. For each pair $(a, b) \in A \times B$, two processes $(g_n^{a,b})_n$ and $(p_n^{a,b})_n$ are given: $p_n^{a,b}$ is the probability that the game stops in stage n, if (a, b) is played in that stage

[1] We thank a referee for pointing out the issue.
[2] We thank Sylvain Sorin for suggesting this generalization.

and the game has not stopped earlier; $g_n^{a,b}$ is the payoff that is received by player 1 in that case. The payoff is zero if the game never stops.

In words, those are games where the actions of the players may influence the probability of termination and the terminal payoff, but, if the game continues, they do not influence the information of the players at the next stage.

Clearly, stopping games belong to this class, with $A = B = \{\text{stop, continue}\}$, and $p_n^{a,b} = 0$ if $a = b =$ continue, and $p_n^{a,b} = 1$ otherwise. To specify properly the game, we need to tell what is known at stage n about past choices of the players. This turns here to be irrelevant (in contrast with other classes of stochastic games).

We briefly sketch how the proof in Section 5.2 has to be adapted. All notations are the same. The only difficulty lies in defining N_0, since, loosely speaking, there exists no least terminating strategy. Partition Ω into Ω_c and Ω_d, where $\Omega_c \in \mathscr{A} = F_\infty$ is the convergence set of the sequence $(v_n)_n$, and $\Omega_d = \Omega \backslash \Omega_c$. We choose an integer N_0 large enough and an event $F \in \mathscr{F}_{N_0}$ such that $\mathbf{P}(F \triangle \Omega_c) < \eta$.

We define a strategy $\hat{\mathbf{x}}$ that has the following features: it coincides with $\mathbf{x}^*$ unless F occurs and $v_{N_0} > \epsilon$; in that case, it switches at stage N_0 to the strategy we defined in section 5.2, *i.e.*, it plays a sequence of locally optimal strategies in properly chosen discounted games if $v_{N_0} > \varepsilon$.

It can be shown that $\hat{\mathbf{x}}$ guarantees w up to 7ε.

It is not clear whether the second proof can be generalized to this class of games.

We conclude with a brief discussion on our assumptions related to the filtration $(\mathscr{F}_n)_n$. We assumed that the payoff processes (a_n), (b_n) and (c_n) are adapted, and that X_n and Y_n are independent of $\mathscr{F}_n$ but $\mathscr{F}_{n+1}$-measurable. As we argued previously, the first assumption can be totally dispensed with. Informally, the second assumption means that (i) in any stage, each player has no information about the action the other player is about to choose, and (ii) past choices are observed. The first part of the assumption is crucial, but the second is irrelevant. Observe indeed that our ε-optimal strategies make no use of the past actions of the opponent. Finally, it is crucial that both players have the same filtration. The existence of the value does not extend to the situation where the payoff processes are constant, and one of the players has more information than the other about their value.

References

1. Alario-Nazaret, M., Lepeltier, J.P., Marchal, B.: Dynkin Games, Stochastic Differential Systems (Bad Honnef), 23–32, Lecture notes in Control and Information Sciences, **43**, Springer Verlag (1982)
2. Aumann, R.J.: Mixed and behavior strategies in infinite extensive games, in *Advances in Game Theory*, Dresher, M, Shapley, L.S. and Tucker, A.W. (eds), Annals of Mathematics Study 52, Princeton University Press (1964)
3. Bismut, J.M.: Sur un problème de Dynkin, Z. Warsch. V. Geb., **39**, 31–53 (1977)
4. Bismut, J.M.: Temps d'arrêt optimal, quasi-temps d'arrêt et retournement du temps, Annals of Probability, **6**, 933–964 (1979)
5. Dynkin, E.B.: Game Variant of a Problem on Optimal Stopping, Soviet Math. Dokl., **10**, 270–274 (1969)
6. Kiefer, Y.I.: Optimal Stopped Games, T. Prob. Appl., **16**, 185–189 (1971)

7. Kuhn, H.W.: Extensive games and the problem of information, in Contributions to the Theory of Games, Kuhn, H.W. and Tucker, A.W. (eds), Annals of Mathematics Study 28, Princeton University Press (1953)
8. Kuratowski, K., Ryll-Nardzewski, C.: A General Theorem on Selectors, Bull. Acad. Polon. Sci. Ser. Sci. Math. Astronom. Phys., **13**, 397–403 (1965)
9. Lepeltier, J.P., Maingueneau, M.A.: Le jeu de Dynkin en théorie générale sans l'hypothèse de Mokobodski, Stochastics, **13**, 25–44 (1984)
10. Maitra, A., Sudderth, W.: Borel stochastic games with lim sup payoff, Ann. Probab., **21**, 861–885 (1993)
11. Mertens, J.F., Neyman, A.: Stochastic Games, Int. J. Game Th., **10**, 53–66 (1981)
12. Neveu: Discrete-Parameter Martingales, North-Holland, Amsterdam (1975)
13. Rosenberg, D., Vieille, N.: The maxmin of recursive games with incomplete information on one side, Mathematics of Operations Research, **25**, 23–35 (2000)
14. Shapley, L.S.: Stochastic Games, Proc. Nat. Acad. Sci. U.S.A., **39**, 1095–1100 (1953)
15. Touzi, N., Vieille, N.: Continuous-time Dynkin games with mixed strategies, mimeo (1999)
16. Yasuda: On a Randomized Strategy in Neveu's Stopping Problem, Stochastic Processes Appl., **21**, 159–166 (1985)