Discussion Paper No. 1258

Stopping Games with Randomized Strategies

by
Dinah Rosenberg,
Eilon Solan
and Nicolas Vieille

May 4, 1999

http://www.kellogg.nwu.edu/research/math

Stopping Games with Randomized Strategies

Dinah Rosenberg[1], Eilon Solan[2] and Nicolas Vieille[3]

May 4, 1999

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] School of Mathematical Sciences, Sackler Faculty of Exact Sciences, Tel Aviv University, Tel Aviv 69978, Israel.
[2] MEDS Department, Kellogg Graduate School of Management, Northwestern University, 2001 Sheridan Rd., Evanston, IL 60208, USA.
[3] GRAPE, Université Bordeaux-IV Montesquieu, avenue Leon Duguit, 33608 Pessac, and Laboratoire d'Econométrie de l'Ecole Polytechnique.

1 Introduction

Dynkin (1969) presented the following optimization problem. Two players observe stochastic sequences $(r(n), x(n))$. 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 $\sup |r(n)| \in L^1$, 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.

There is a broad literature on continuous time Dynkin games giving sufficient conditions for the existence of the value and optimal strategies: Bismuth (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 Lepeltier, Alario 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 game 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 (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 and 6 we provide the two constructions of ϵ-optimal strategies.

2   The Model and the Main Results

Let $(\Omega, \mathcal{A}, \mathbf{P})$ be a probability space, that is rich enough to support a double sequence $(X_n, Y_n)_{n=0}^{\infty}$ of *iid* variables, uniformly distributed over $[0, 1]$. Let $(\mathcal{F}_n)$ be a filtration over $(\Omega, \mathcal{A}, \mathbf{P})$ (the information available at stage n). We assume that (X_n, Y_n) is $\mathcal{F}_{n+1}$-measurable, and independent of $\mathcal{F}_n$.

Let $(a_n), (b_n), (c_n)$ be processes, defined over $(\Omega, \mathcal{A}, \mathbf{P})$. We assume

$$\sup_n |a|_n, \sup_n |b|_n, \sup_n |c|_n \in L^1(\mathbf{P}). \tag{1}$$

A *strategy* for each of the players is a randomized stopping time. A strategy of player 1 (*resp.* player 2) is a process $\mathbf{x} = (x_n)$ (*resp.* $\mathbf{y} = (y_n)$) adapted to $(\mathcal{F}_n)$, with values in $[0, 1]$: x_n is the probability that player 1 stops at stage n, conditional on stopping occurs after $n - 1$.

3

Given strategies $(\mathbf{x}, \mathbf{y})$, define the stopping stages $t_1 = \inf\{n, X_n \leq x_n\}$, $t_2 = \inf\{n, Y_n \leq y_n\}$, and

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

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

We set $r(\mathbf{x}, \mathbf{y}) = a_{t_1} 1_{t_1 < t_2} + b_{t_2} 1_{t_2 < t_1} + c_{t_1} 1_{t_1 = t_2 < +\infty}$, which we write also $r_t(\mathbf{x}, \mathbf{y})$. 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. ϵ-*optimal strategies for player 2 are defined analogously.*

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

Let $\lambda \in]0, 1[$. Define the λ-discounted payoff by $r_\lambda(\mathbf{x}, \mathbf{y}) = (1-\lambda)^{t+1} r_t(\mathbf{x}, \mathbf{y})$ and $\gamma_\lambda(\mathbf{x}, \mathbf{y}) = \mathbf{E}(r_\lambda(\mathbf{x}, \mathbf{y}))$. Notice the exponent $t + 1$ in the definition of r_λ. This differs from the usual convention which uses t instead. This has no incidence on the results.

DEFINITION 2.2 v_λ *is the* λ-discounted value *of the game if* $v_\lambda = \operatorname{esssup}_{\mathbf{x}} \operatorname{essinf}_{\mathbf{y}} \gamma_\lambda(\mathbf{x}, \mathbf{y}) = \operatorname{essinf}_{\mathbf{y}} \operatorname{esssup}_{\mathbf{x}} \gamma_\lambda(\mathbf{x}, \mathbf{y})$.

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

Theorem 2 $v = \lim v_\lambda$.

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

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 3 *For every* ϵ, *there exists* $\mathbf{x}^*$ *and* 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 3 is a refinement of Theorem 1. It can be shown that it also implies Theorem 2.

Theorem 3 was proved by Mertens and Neyman (1981) for general stochastic games with bounded payoffs, in which the discounted values satisfy some bounded variation property. In the case of recursive games with bounded payoffs, Rosenberg and Vieille (1998) proved that Theorem 3 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.

For the sake of exposition, we assume that (a_n), (b_n), and (c_n) are adapted. This assumption can be dispensed with, by replacing everywhere (a_n), (b_n), and (c_n) by their conditional expectations given $\mathcal{F}_n$.

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 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 $f \le$ val g if $f \le g$, and $|\text{val } f - \text{val } g| \le \sup_{A \times B} |f - g|$.

For any real-valued $\mathcal{F}_n$-measurable function f, we let $G_n(f)$ be the 0-sum game with ($\mathcal{F}_n$-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 is an $\mathcal{F}_n$-measurable variable x_n in $[0, 1]$, 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 ($\mathcal{F}_n$-measurable) payoff to player 1 when the players use strategies x_n and y_n:

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

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 $G_n(f)(\omega)$.

We now argue that each player has optimal strategies 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,*

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

A symmetric property holds for player 2.

Proof: for each ω, 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 measurable. By Kuratowski and Ryll-Nardzewski (1965) its graph has a 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 $G_n(x_n, y_n; f) = \text{val } G_n(f)$ 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

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

and

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

The proposition below contains obvious properties.

PROPOSITION 3.2 $(\overline{v}_n(\lambda))$ and $(\underline{v}_n(\lambda))$ 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, $\mathbf{P}$-a.s., $(\overline{v}_n(\lambda))$ and $(\underline{v}_n(\lambda))$ are both solutions of the recursive equation

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

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.

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

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

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.

Notice that $v(\lambda) = \mathbf{E}[v_0(\lambda)]$

We now let (λ_p) be any sequence which converges to 0. Set $v_n = \limsup_{p \to \infty} v_n(\lambda_p)$. Note that v_n is $\mathcal{F}_n$-measurable. We shall prove that $v = \mathbf{E}[v_0]$ is the uniform value of the game.

PROPOSITION 3.3 *One has $v_n \leq$ val $G(\mathbf{E}[v_{n+1}|\mathcal{F}_n])$, for every n.*

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

$$v_n(\lambda_q) \leq \alpha_{p_0} \text{val } G_n(\mathbf{E}[\sup_{p \geq p_0} v_{n+1}(\lambda_p)|\mathcal{F}_n]),$$

provided $q \geq p_0$, where $\alpha_{p_0} = 1 - \lambda_{p_0}$ if the val is negative, and 1 otherwise. Since the right-hand side is independent of q, one has

$$\sup_{p \geq p_0} v_n(\lambda_p) \leq \alpha_{p_0} \text{val } G_n(\mathbf{E}[\sup_{p \geq p_0} v_{n+1}(\lambda_p)|\mathcal{F}_n]).$$

Both sides are non-decreasing functions of p_0. The result follows by dominated convergence. ∎

3.3 Locally optimal strategies and martingale properties

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

Thus, for every strategy $\mathbf{y}$ and $n \in \mathbf{N}$, one has

$$G_n(x_n^*, y_n; \mathbf{E}[v_{n+1}|\mathcal{F}_n]) \geq v_n, \quad \mathbf{P}\text{-a.s.} \tag{4}$$

and

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

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.

Equation 3 and Proposition 3.3 provide recursive formulas for (v_n) and $(v_n(\lambda))$. In order to interpret these formulas in terms of submartingale properties, we use auxiliary processes.

We say that an inequality $\alpha \geq \beta$ holds $\mathbf{P}$-a.s. on an event A if $\mathbf{P}(A; \alpha < \beta) = 0$.

Let (α_n) be an adapted integrable process on $(\Omega, \mathcal{A}, (\mathcal{F}_n), \mathbf{P})$, 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 $n \in \mathbf{N}$, the inequality $\mathbf{E}[\alpha_{n+1}|\mathcal{F}_n] \geq \alpha_n$ holds $\mathbf{P}$-a.s. on the event $\{s_1 \leq n < s_2\}$. $(\alpha_n)_n$ 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 $\mathcal{F}_{s_1}$ the σ-algebra of events known at stage s_1. Then one has $\mathbf{E}[\alpha_s|\mathcal{F}_{s_1}] \geq \alpha_{s_1}$, $\mathbf{P}$-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 $\tilde{\alpha}_n = r_t$ 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 $\{t \geq n\}$:

$$\mathbf{E}[\tilde{\alpha}_{n+1}|\mathcal{F}_n] = G(x_n, y_n; \alpha_{n+1}).$$

Let $\mathbf{x}^* = (x_n^*)$ and $\mathbf{x}(\lambda) = (x_n(\lambda))$ be strategies of player 1 that satisfy (4) and (5) respectively.

Lemma 4 *Let $\mathbf{y}$ be a strategy of player 2, and $\lambda \in]0, 1[$. Under $(\mathbf{x}(\lambda), \mathbf{y})$, $((1 - \lambda)^n \tilde{v}_n(\lambda))$ is a submartingale up to $t + 1$. Under $(\mathbf{x}^*, \mathbf{y})$, $\tilde{v}_n$ is a submartingale.*

Proof: notice that $\sup_n \tilde{v}_n(\lambda)$ and $\sup_n \tilde{v}_n$ belong to $L^1(\mathbf{P})$, for every choice of $(\mathbf{x}, \mathbf{y})$. We start with the first claim. Let $n \in \mathbf{N}$.

On $\{t \geq n\}$, $\tilde{v}_n(\lambda) = v_n(\lambda)$. Thus

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

which is, $\mathbf{P}$-a.s., at least $v_n(\lambda)$, by definition of $\mathbf{x}(\lambda)$.

For a similar reason,

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

on the event $\{t \geq n\}$. Since $\tilde{v}_{n+1} = \tilde{v}_n$ if $t < n$, the second assertion is also established. ∎

Corollary 5 *For every $\mathbf{y}$, $\gamma_\lambda(\mathbf{x}(\lambda), \mathbf{y}) \geq \mathbf{E}(v_0(\lambda))$.*

Proof: for each n, applying the submartingale property with the stopping time $\min(t+1,n)$ yields

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

that is,

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

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

Remark: denote by $\mathbf{0}$ the strategy which never stops ($x_n = 0$ for all n). Let $\mathbf{x}$ be any strategy which coincides with $\mathbf{x}(\lambda)$ from stage n on. It is clear that, for every $\mathbf{y}$,

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

on the event $t \geq n$. Notice that, for $\mathbf{y} = \mathbf{0}$ the left-hand side belongs $\mathbf{P}$-a.s. to the convex hull of the set $\{0, a_n, a_{n+1}, \ldots\}$. This fact will be used in section 6.

Corollary 5 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 in which playing $\mathbf{x}^*$ needs not be an optimal strategy.

EXAMPLE

	1
1	0

This matrix notation is a shortcut for the game with payoffs $a_n = b_n = 1$, $c_n = 0$, $\mathbf{P}$-a.s. for every n. v_n and $v_n(\lambda)$ are independent of n and constant. We simply write v and $v(\lambda)$. The real number $0 \leq v(\lambda) \leq 1$ is a solution to the equation $v(\lambda) = (1-\lambda)\text{val}G(v(\lambda))$, from which it is easily derived $v(\lambda) = 1 - \sqrt{\lambda}$, and $\mathbf{x}(\lambda) = \sqrt{\lambda}/(1 + \sqrt{\lambda})$. Therefore, $v = 1$, and $\mathbf{x}^* = \mathbf{0}$. However, $\gamma(\mathbf{x}^*, \mathbf{0}) = 0$.

4 The Main Ideas of the Constructions

To explain the ideas that underlie the two constructions of ϵ-optimal strategies, we consider the deterministic case with payoffs bounded by 1. In that

10

case, (v_n) is simply a bounded sequence of real numbers. Since $\tilde{v}_n$ is a sub-martingale w.r.t. $(\mathbf{x}, \mathbf{0})$, and coincides with v_n up to t_1, one deduces that, *either* (v_n) is a convergent sequence, *or* $t_1 < +\infty$, almost surely (or both).

As mentioned in Section 3, $\mathbf{x}(\lambda)$ is a λ-discounted optimal strategy for player 1. However, as the last example shows, $\mathbf{x}$ need not be an optimal strategy.

$\mathbf{x}$ is not optimal since the v_n's are positive (and equal to 1). Thus, player 1 expects that if the game continues, he will receive 1. However, if the game continues forever, he never receives this continuation payoff, and his overall payoff is 0. If, on the other hand, v_n was always *negative*, then, since $\tilde{v}_n$ is a submartingale w.r.t. $(\mathbf{x}, \mathbf{y})$ for every $\mathbf{y}$, following $\mathbf{x}$ is an optimal strategy for player 1.

In the λ-discounted game we do not encounter this problem. If the value is positive and player 1 continues with probability 1, then $v_{n+1}(\lambda)$ must be larger than $v_n(\lambda)$. Since the payoffs are bounded player 1 cannot delay stopping the game too much.

Assume for a moment that the payoff process is periodic. For a given $\epsilon > 0$, one can choose λ sufficiently small such that $\| v(\lambda) - v \| < \epsilon^2$. Construct the following strategy: play according to $\mathbf{x}(\lambda)$ until the value drops below 0. Then switch to $\mathbf{x}$ until the value is above ϵ, continue playing according to $\mathbf{x}(\lambda)$ and so on. One can show that whenever the player follows $\mathbf{x}$, v_n is a submartingale, whereas when he follows $\mathbf{x}(\lambda)$, $v_n(\lambda)$ is a submartingale. Moreover, since the payoffs are bounded by 1, if $v_n > \epsilon$ and $v_m < 0$ (where $m > n$) then the probability that the game is stopped between stages n and m is at least ϵ.

We would like to approximate the lower bound of the expected payoff of player 1. It turns out that this bound can be approximated by v_1, but each switch between $\mathbf{x}$ and $\mathbf{x}(\lambda)$ adds an error term of ϵ^2. Since the game is stopped after $O(1/\epsilon)$ stages, it means that this strategy guarantees player 1 a payoff of at least $v_1 - O(\epsilon)$.

This method was used by Rosenberg and Vieille (1998) (see also Thuijsman and Vrieze (1992) for a precursor) for recursive games.

When the payoff process is not periodic, one needs not be able to choose a λ such that $\| v(\lambda) - v \| < \epsilon^2$. In this case, we choose the appropriate λ *at each switch*. That is, whenever v_n exceeds ϵ we choose λ such that $|v_n(\lambda) - v_n| < \epsilon^2$, and player 1 follows $\mathbf{x}(\lambda)$ until the value falls below 0. The rest of the argument remains the same. This summarizes the method used in the first approach.

The second approach is based on a different idea. Since $\mathbf{x}$ is locally optimal, it follows that at every stage n, if the game was not stopped before that stage then the expected payoff in the local game is at least v_n, whatever player 2 plays. Let us define now a new strategy for player 1. At stage n, player 1 checks whether $a_n \geq v_n - \epsilon$; that is, if he stops and player 2 continues, whether his payoff is at least $v_n - \epsilon$. If $a_n < v_n - \epsilon$, player 1 stops with probability $\mathbf{x}_n$, whereas if $a_n \geq v_n - \epsilon$, player 1 stops with probability $\mathbf{x}_n + \epsilon$.

To see that this strategy guarantees player 1 an expected payoff of v_1, we note that the following points hold:

1. If player 2 stops the game, then the expected payoff of player 1 remains the same (up to an ϵ).

2. In the case that player 2 always continue, since player 1 changes his strategy *only* when a unilateral stopping is favorable for him, $\mathbf{E}v_n \geq v_1 - \epsilon$.

A crucial observation is that $a_n < v_n - \epsilon$ occurs infinitely often, hence whatever player 2 plays, the game will eventually be stopped. Hence one conclude that the expected payoff of player 1 is at least $v_1 - \epsilon$.

5   An ϵ-optimal strategy for player 1 - I

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 $A \in \mathcal{A}$,

$$\mathbf{P}(A) < \eta \Rightarrow \mathbf{E}(m1_A) < \epsilon. \tag{6}$$

Notice that $v_n(\lambda), v_n \leq m$, $\mathbf{P}$-a.s. for every n.

The sequence (v_n) has no convergence properties. 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 such that

$$\mathbf{P}\{ \sup_{n,m \geq N_0} |v_n - v_m| > \epsilon/2, t_1 \geq N_0\} < \eta. \tag{7}$$

We define a sequence (s_p) of stopping times. Set $s_0 = N_0$ on $\{v_{N_0} > \epsilon\} \cap \{t_1 \geq N_0\}$, and $s_0 = +\infty$ otherwise. Choose an $\mathcal{F}_{s_0}$-measurable function λ_0 with $v_{s_0}(\lambda_0) > v_{s_0} - \epsilon^2$ if $s_0 < +\infty$.

Set $s_{p+1} = \inf\{n > s_p, v_n(\lambda_p) \leq 0\}$ and choose an $\mathcal{F}_{s_{p+1}}$-measurable function λ_{p+1}, such that $v_{s_{p+1}}(\lambda_{p+1}) > v_{s_{p+1}} - \epsilon^2$ if $s_{p+1} < +\infty$.

Define $\overline{\mathbf{x}}$ as $\overline{\mathbf{x}}_n = \mathbf{x}^*$ for $n < s_0$ and $\overline{\mathbf{x}}_n = \mathbf{x}_n(\lambda_p)$ for $s_p \leq n < s_{p+1}$.

By Lemma 4, for every $\mathbf{y}$, $(\tilde{v}_n)$ 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 $\min(s_{p+1}, t+1)$.

We prove below the following result.

PROPOSITION 5.1 *There exists N such that, for every $\mathbf{y}$ and $n \geq N$, one has $\gamma_n(\overline{\mathbf{x}}, \mathbf{y}) \geq v - 7\epsilon$.*

Comments: $\overline{\mathbf{x}}$ is defined in terms of $\mathbf{x}^*$ and $\mathbf{x}(\lambda)$ for various λ's. Intuitively, $\overline{\mathbf{x}}$ coincides with $\mathbf{x}^*$ whenever $v_n \leq 0$, thus preventing (v_n) from decreasing (submartingale property). If $v_{N_0} > 0$ (N_0 is chosen so that $\tilde{v}_n$ no longer oscillates afterwards), a λ is chosen such that $v_n(\lambda)$ is approximately equal to v_n and player 1, by playing $\mathbf{x}(\lambda)$ prevents $(v_n(\lambda))$ from decreasing, until the approximation of v_n by $v_n(\lambda)$ gets poor. Player 1 then switches to a new value for λ. A crucial feature of this construction is that, with high probability, in *any* such switch s_p, the new state variable $v_{s_p}(\lambda_p)$ exceeds $v_{s_p}(\lambda_{p-1})$. Indeed, $v_{s_p}(\lambda_{p-1}) \leq 0$ by definition, whereas $v_{s_p}(\lambda_p) > v_{s_p} - \epsilon^2$, which is higher than $\epsilon/2 - \epsilon^2$ with high probability.

By such a device, player 1 forces the termination of the game when $v_n > 0$, in a sense which will be made precise below.

We translate these comments into submartingale properties of an associated *state variable*.

To avoid "small events" troubles, it is convenient to introduce $s = \inf\{n > s_0, v_n < \epsilon/2\}$. Notice that

$$\mathbf{P}(s < +\infty) < \eta. \tag{8}$$

We introduce the state variable w_n defined as $w_n = v_n - \epsilon^2$ for $n < s_0$ and $w_n = v_n(\lambda_p)$ for $s_p \leq n < s_{p+1}$, except if $n = s = s_p$, for some $p \geq 1$, in which case we set $w_n = v_n(\lambda_{p-1})$.

Lemma 6 *$(\tilde{w}_n)$ is a submartingale up to $\min(s, t+1)$.*

13

Proof: let $n \in \mathbf{N}$. We prove that $\mathbf{E}[\tilde{w}_{n+1}|\mathcal{F}_n] \geq w_n$ when $\min(s, t+1) > n$.

If $n < s_0$, $w_n = v_n - \epsilon^2$, $w_{n+1} \geq v_{n+1} - \epsilon^2$ (with equality if $n+1 < s_0$), and $\overline{x}_n = x_n^*$. Thus, $\mathbf{E}[\tilde{w}_{n+1}|\mathcal{F}_n] \geq G(x_n^*, y_n; \mathbf{E}[v_{n+1} - \epsilon^2|\mathcal{F}_n]) \geq v_n - \epsilon^2$, where the second inequality follows from the inequality $G(x_n^*, y_n; \mathbf{E}[v_{n+1}|\mathcal{F}_n]) \geq v_n$ and since the val operator is non-expanding.

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

$$G(\overline{x}_n, y_n; \mathbf{E}[v_{n+1}(\lambda_p)|\mathcal{F}_n]) \geq \frac{1}{1 - \lambda_p} v_n(\lambda_p) \geq v_n(\lambda_p)$$

since $v_n(\lambda_p) > 0$. The claim follows by noting that $\tilde{w}_{n+1} \geq \tilde{v}_{n+1}(\lambda_p)$. ∎

Lemma 7 *One has* $\mathbf{P}(s_0 = N_0, s = t = +\infty) = 0$.

Proof: we proceed in two steps. We prove first that

$$\mathbf{P}(s_p < s = t = s_{p+1} = +\infty) = 0.$$

¿From $\min(s_p, s, t+1)$ up to $\min(s_{p+1}, s, t+1)$, $((1 - \lambda_p)^n \tilde{w}_n)$ is a submartingale. Thus, for every $n_0 \leq n \leq N$, the sampling property applied to the finite stopping time $\min(s_{p+1}, s, t+1, N)$ yields

$$w_n \leq \frac{1}{(1 - \lambda_p)^n} \mathbf{E}[m(1-\lambda_p)^{\min(s_{p+1}, s, t+1)} 1_{\min(s_{p+1}, s, t+1) \leq N} + w_N(1-\lambda_p)^N 1_{\min(s_{p+1}, s, t+1) > N}|\mathcal{F}_n]$$

on the event $\{s_p \leq n_0 \leq n < \min(s_{p+1}, s, t+1)\}$.

By taking $N \to +\infty$ and by dominated convergence for conditional expectations, one obtains, for every $n \geq n_0$,

$$w_n \leq \mathbf{E}[m(1 - \lambda_p)^{\min(s_{p+1}, s, t+1) - n} 1_{\min(s_{p+1}, s, t+1) < +\infty}|\mathcal{F}_n] \tag{9}$$

on the event $\{s_p \leq n_0 \leq n < \min(s_{p+1}, s, t+1)\}$.

By taking the limit $n \to \infty$ in (9), one gets $\limsup w_n \leq 0$ **P**-a.s. on the event $\{s_p \leq n_0 < s = t = s_{p+1} = +\infty\}$. But w_n is at least $\frac{\epsilon}{4}$, **P**-a.s. for every n. Therefore $\mathbf{P}(s_p \leq n_0 < s = t = s_{p+1} = +\infty) = 0$. This ends the first step.

By the first step, applied inductively, $\mathbf{P}(s_0 = N_0, s_p = t = s = +\infty) = 0$ for every p.

By the submartingale property of $\tilde{v}_n(\lambda_p)$ between $\min(s_p, s, t+1)$ and $\min(s_{p+1}, s, t+1)$, one has, since $v_{s_{p+1}}(\lambda_p) < 0$,

$$v_{s_p}(\lambda_p) \leq \mathbf{E}[m 1_{\min(s,t+1)\leq s_{p+1}} + 0 \cdot 1_{s_{p+1} < \min(s,t+1)} | \mathcal{F}_{s_p}]$$

on the event $\{s_0 = N_0, s_p < \min(s, t+1)\}$. Since $v_{s_p}(\lambda_p) \geq \frac{\epsilon}{4}$, it follows by taking expectations that

$$\frac{\epsilon}{4} \mathbf{P}(s_0 = N_0, s_p < \min(s, t+1)) \leq \mathbf{E}(m 1_{s_p < \min(s,t+1) < +\infty}).$$

As p goes to infinity, the left-hand side converges to $\frac{\epsilon}{4} \mathbf{P}(s_0 = N_0, s = t = +\infty)$, while the right-hand side converges to 0. The result follows. ∎

Proof of Proposition 5.1:

Apply first the previous lemma with $\mathbf{y} = 0$: for some $N_1 \in \mathbf{N}$,

$$\mathbf{P}(s_0 = N_0, t \geq N_1, s = +\infty) < \eta. \tag{10}$$

This readily implies that the same holds for *every* $\mathbf{y}$.

On the other hand, one derives from the definition of N_0 and s_0 that

$$\mathbf{P}(s_0 = +\infty, t \geq k, v_k > 3\epsilon/2) < \eta, \tag{11}$$

for every $k > N_0$.

By (6), there is $N_2 \geq N_1$ such that for every $\mathbf{y}$ and every $n \geq N_2$,

$$\left| \gamma_n(\overline{\mathbf{x}}, \mathbf{y}) - \frac{1}{n - N_1 + 1} \mathbf{E}\left(\sum_{k=N_1}^{n} g_k \right) \right| < \epsilon \tag{12}$$

(for n suficiently large, the average payoff up to stage n is close to the average payoff between stages N_1 and n).

For $k \geq N_2$, we estimate $\mathbf{E}(g_k)$ as follows. Set $A_k = \{s < +\infty\}$, $B_k = \{t < k, s = +\infty\}$, $C_k = \{t \geq k, s_0 = N_0, s = +\infty\}$, $D_k = \{t \geq k, v_k \leq 3\epsilon/2, s_0 = s = +\infty\}$, and $E_k = \{t \geq k, v_k > 3\epsilon/2, s_0 = s = +\infty\}$. Notice that $(A_k, B_k, C_k, D_k, E_k)$ is a partition of Ω.

On A_k and on E_k, $g_k \geq -m \geq \tilde{w}_{\min(k,s,t+1)} - 2m$. On B_k, $g_k = \tilde{w}_{\min(k,s,t+1)}$. On C_k, $g_k = 0 \geq \tilde{w}_{\min(k,s,t+1)} - m$. On D_k, $g_k = 0 \geq \tilde{w}_{\min(k,s,t+1)} - 3\epsilon/2$.

Therefore,

$$\mathbf{E}[g_k] \geq \mathbf{E}[\tilde{w}_{\min(k,s,t+1)}] - 3\epsilon/2 - \mathbf{E}[2m 1_{A_k}] - \mathbf{E}[2m 1_{E_k}] - \mathbf{E}[m 1_{C_k}]. \tag{13}$$

15

Since $(\tilde{w}_{\min(n,s,t+1)})$ is a submartingale, $\mathbf{E}[\tilde{w}_{\min(k,s,t+1)}] \geq \mathbf{E}[\tilde{w}_0] = \mathbf{E}[v_0] = v$. On the other hand, $\mathbf{P}(A_k), \mathbf{P}(C_k), \mathbf{P}(E_k) < \eta$ by inequalities (8), (10) and (11).

One thus obtains from (13) and (6)

$$\mathbf{E}(g_k) \geq v - 7\epsilon.$$

The result then follows from (12). ∎

Recall that v is defined as $v = \mathbf{E}[v_0]$, where $v_0 = \limsup v_0(\lambda_p)$. Define $z_0 = \liminf v_0(\lambda_p)$, and $z = \mathbf{E}[z_0]$. By symmetry, for each ϵ, there exists a strategy $\overline{\mathbf{y}}$ such that $\gamma_n(\mathbf{x}, \overline{\mathbf{y}}) \leq z + 5\epsilon$, provided n is large enough. This readily implies $v - 5\epsilon \leq z + 5\epsilon$. Since $z \leq v$, and ϵ is arbitrary, one obtains $v = z$. This shows that v is the uniform value of the game. The claim about the limit of discounted values is now immediate, since the sequence (λ_p) used to define v is arbitrary.

6   An ϵ-optimal strategy for player 1 - II

Recall the definitions of η and N_0 from section 5 (equations (6) and (7)). Denote

$$\Omega' = \Omega \setminus \left\{ \sup_{n,m \geq N_0} |v_n - v_m| > \epsilon/2, t_1 \geq N_0 \right\}.$$

By (7), $\mathbf{P}(\Omega') > 1 - \eta$.

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

$$\hat{x}_n = \begin{cases} \min\{x_n^* + \eta, 1\} & \text{if } n \geq N_0 \text{ and } 5\epsilon < v_{N_0} < a_n + 5\epsilon \\ x_n^* & \text{otherwise.} \end{cases}$$

We will prove that $\hat{\mathbf{x}}$ is an ϵ-optimal strategy for player 1. We set $\Omega_1 = \Omega' \cap \{v_{N_0} > 5\epsilon\}$.

Lemma 8 *On the set Ω_1, one has, $\mathbf{P}$-a.s, $a_n > v_{N_0} - 5\epsilon$ for infinitely many n.*

Proof: Otherwise, there exists a subset $\Omega'' \subset \Omega_1$ such that $\mathbf{P}(\Omega'') > 0$ and $a_n(\omega) > v_{N_0}(\omega) - 5\epsilon$ finitely many times for each $\omega \in \Omega''$. Let n' be sufficiently large such that

$$\mathbf{P}(\Omega'' \cap \{\sup_{n > n'} a_n \leq v_{N_0} - 5\epsilon\}) > (1 - \epsilon)\mathbf{P}(\Omega'') > 0.$$

16

Since $\Omega'' \subseteq \Omega'$, $|v_{N_0} - v_{n'}|$ on Ω'', so that

$$\mathbf{P}(\Omega'' \cap \{\sup_{n>n'} a_n \leq v_{n'} - 4\epsilon\}) > (1 - \epsilon)\mathbf{P}(\Omega'') > 0.$$

Choose $\lambda > 0$ such that $v_{n'}(\lambda) \geq v_{n'} - \epsilon$ with positive probability on $\Omega'' \cap \{\sup_{n>n'} a_n \leq v_{n'} - 5\epsilon\}$.

On $\Omega'' \cap \{\sup_{n>n'} a_n \leq v_{n'} - 4\epsilon\}$, one then has with positive probability,

$$v_{n'} - \epsilon \leq v_{n'}(\lambda) \leq v_{n'} - 4\epsilon$$

a contradiction. Note that the second inequality holds by (6) and by the fact that $v_{n'}(\lambda)$ lies in the convex hull of $\{0, a_{n'}, a_{n'+1}, \ldots\}$. $\blacksquare$

Corollary 9 t is **P**-a.s. finite on Ω_1.

Lemma 10 For every $\mathbf{y}$, $\gamma(\hat{\mathbf{x}}, \mathbf{y}) \geq v - 14\epsilon$ on $t > N_0$.

Proof: Fix a strategy $\mathbf{y}$ of player 2. Set $\Omega_2 = \{t > N_0, v_{N_0} > 5\epsilon\}$, and $\Omega_3 = \Omega_2 \cap \{\sup_{n,m \geq N_0} |v_n - v_m| \leq \epsilon\}$.

Conditional on $\{t_2 = n = t\}$, the expectation of r_t is $G(\hat{x}_n, 1; \mathbf{E}[v_{n+1}|\mathcal{F}_n])$. Since $\|\hat{x}_n - x_n^*\| < \eta$, this expectation is at least $v_n - \epsilon$. In particular,

$$\mathbf{E}[r_t 1_{\Omega_3} 1_{t=t_2<+\infty}] \geq \mathbf{E}[v_{N_0} 1_{\Omega_3} 1_{t=t_2<+\infty}] - 2\epsilon; \qquad (14)$$

On the event $\{t = t_1 = n < t_2\}$, $r_t = a_n$. By definition of $\hat{\mathbf{x}}$, this can happen on Ω_2 only if $a_n \geq v_{N_0} - 5\epsilon$. Thus

$$\mathbf{E}[r_t 1_{t=t_1<t_2} 1_{\Omega_3}] \geq \mathbf{E}[v_{N_0} 1_{t=t_1<t_2} 1_{\Omega_3}] - 6\epsilon. \qquad (15)$$

Since $t_1 < +\infty$, **P**-a.s. on Ω_3, and $\mathbf{P}(\Omega_2 \setminus \Omega_3) < \eta$, one deduces from (14) and (15) that
$$\mathbf{E}[r_t 1_{t<+\infty} 1_{\Omega_2}] \geq \mathbf{E}[v_{N_0} 1_{\Omega_2}] - 7\epsilon.$$

Finally, define the stopping time θ by $\theta = N_0$ on $\{t > N_0\} \cap \{v_{N_0} > 5\epsilon\}$, 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, **P**-a.s. on the event $\{\theta > N_0\}$, say to $\tilde{v}_\infty$.

17

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

$$\mathbf{E}(\tilde{v}_\theta) \geq \tilde{v}_0 = v_0. \tag{16}$$

By definition of $(\tilde{v}_n)$, one has $\tilde{v}_\infty = r_t$ if $t < +\infty$, $\tilde{v}_\infty \leq 6\epsilon$ if $t = +\infty$ and $\sup_{n,m \geq N_0} |v_n - v_m| \leq \epsilon$, and $\tilde{v}_\infty \leq m$ otherwise. Thus, by (6),

$$\mathbf{E}[\tilde{v}_\infty 1_{\theta > N_0}] \leq \mathbf{E}[r_t 1_{t < +\infty} 1_{\theta > N_0}] + 6\epsilon + \epsilon.$$

The inequality (16) may be rewritten as

$$\mathbf{E}[r_t 1_{t \leq N_0} + v_{N_0} 1_{\Omega_2} + \tilde{v}_\infty 1_{\theta > N_0}] \geq v_0.$$

and therefore $\mathbf{E}(r_t 1_{t < +\infty}) \geq v_0 - 14\epsilon.$ ■

This does not prove that v is the *uniform* value. The uniformity can be obtained along similar lines as in section 5.

References

[1] Bismuth, J.M. (1977) Sur un problème de Dynkin, *Z. Warsch. V. Geb.* 39, 31-53

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

[3] Kiefer Y.I. (1971) Optimal Stopped Games, *T. Prob. Appl.*, **16**, 185-189

[4] Kuratowski K. and Ryll-Nardzewski C. (1965) A General Theorem on Selectors, *Bull. Acad. Polon. Sci. Ser. Sci. Math. Astronom. Phys.*, **13**, 397-403

[5] Lepeltier, J.P., M. Alario and B. Marchal (1982) Jeux de Dynkin, 2e Bad Honnef Workshop on stochastic processes, *Lecture notes in Control and Information Sciences*, Springer Verlag.

[6] Lepeltier, J.P. and M.A. Maingueneau (1984) Le jeu de Dynkin en théorie générale sans l'hypothèse de Mokobodski, *Stochastics*, 13, 25-44.

[7] Maitra A. and Sudderth W. (1993) Borel stochastic games with lim sup payoff, *Ann. Probab.*, **21**, 861-885

[8] Mertens J.F. and Neyman A. (1981) Stochastic Games, *Int. J. Game Th.*, **10**, 53-66

[9] Neveu (1975) Discrete-Parameter Martingales, North-Holland, Amsterdam

[10] Rosenberg D. and Vieille N. (1998) The maxmin of recursive games with incomplete information on one side, to appear in *Mathematics of Operations Research*

[11] Shapley L.S. (1953) Stochastic Games, *Proc. Nat. Acad. Sci. U.S.A.*, **39**, 1095-1100

[12] Thuijsman F. and Vrieze K. (1992) Note on Recursive Games, in *Game Theory and Economics Applications*, Lecture notes in Economics and Mathematical Systems 389, Springer-Verlag, 133-145

[13] Yasuda (1985) On a Randomized Strategy in Neveu's Stopping Problem, *Stochastic Processes Appl.*, **21**, 159-166