Z. Wahrscheinlichkeitstheorie verw. Gebiete
57, 253-264 (1981)

Zeitschrift für
Wahrscheinlichkeitstheorie
und verwandte Gebiete

Optimal Stopping and Supermartingales over Partially Ordered Sets

A. Mandelbaum[1] and R.J. Vanderbei[2]

[1] Dept. of Operations Research, Cornell University, Ithaca, N.Y. 14853, USA
[2] Dept. of Mathematics, Cornell University, Ithaca, N.Y. 14853, USA

1. Introduction

1.1. The subject of this paper is the problem of optimal stopping for discrete multiparameter stochastic processes; in particular, for a family of Markov processes.

In 1953, Snell [12] discovered the relation between optimal stopping of a random sequence and supermartingales. In 1963, Dynkin [3] described the optimal stopping rule for a Markov process in terms of excessive functions. In 1966, Haggstrom [7], motivated by problems of sequential experimental design, extended Snell's results to processes indexed by a tree. All these results are particular cases of the general theory of optimal stopping for a family of random variables indexed by a partially ordered set. This general setting was considered by Krengel and Sucheston [8] who proved a number of general theorems and applied them to the case of functionals of a family of independent identically distributed random variables.

We start by discussing, in the spirit of Snell's theory, the general optimal stopping problem over a partially ordered set. The proofs of the main theorems are similar to Haggstrom's proofs but for completeness we outline them briefly in Sect. 6. Our emphasis is on the nature of stopping points taking values in partially ordered sets. Not all stopping points are appropriate but only a certain subclass which we call predictable. An important result is that the supermartingale sampling theorem holds for predictable stopping points.[1]

The general theory is applied to a family of Markov chains and we get results analogous to Dynkin's results. An example for two independent random walks is considered in Sect. 3.

Finally, we discuss the relation between optimal stopping with time constraints and Walsh's theory of multiharmonic functions (Sect. 4).

1.2. Let Z_t be a sequence of random variables adapted to a filtration $\mathscr{F}_t$. The classical optimal stopping problem is to find a stopping time τ^* which is op-

---

[1] Walsh [15] has extended this to the continuous two parameter case

0044-3719/81/0057/0253/$02.40

timal in the sense that

$$EZ_{\tau^*} = \sup EZ_\tau. \tag{1.1}$$

Here the supremum is taken over all stopping times τ (we use the convention that $Z_\tau = 0$ on $\{\tau = \infty\}$).

This problem was first solved by Snell [12] using the minimal nonnegative supermartingale X_t majorizing Z_t. We call this supermartingale *Snell's envelope*. Thinking of Z_t as the reward obtained by stopping at time t, we can interpret X_t as the highest reward possible if one stops after time t. The sequence X_t satisfies the following stochastic dynamic programming equation

$$X_t = \max\{Z_t, E^{\mathscr{F}_t} X_{t+1}\}. \tag{1.2}$$

The two terms in the max represent two options available at time t: stop and receive Z_t or move to $t+1$ and proceed optimally from there.

Let τ^* be the first time that $X_t = Z_t$. From (1.2) it is clear that X_t is a martingale up to time τ^*. If τ^* is finite with probability one then

$$EZ_{\tau^*} = EX_{\tau^*} = EX_0 = \sup EZ_\tau, \tag{1.3}$$

which shows that it is optimal.

1.3. Let S be a partially ordered set having a minimal element 0. In this case Snell's envelope satisfies

$$X_s = \max\{Z_s; \sup_{u>s} E^{\mathscr{F}_s} X_u\}, \tag{1.4}$$

and, in some sense, it still represents the optimal reward for stopping beyond s. We need to make this precise. If we define a *stopping point* v as a random point in S which satisfies $\{v \leqq s\} \in \mathscr{F}_s$ for all s in S, then X_s is not optimal relative to all stopping points, as the following example shows.

Example. Let S consist of three elements $0, l, r$. The element 0 is smaller than r and l. The elements r and l are not ordered with respect to each other. Let $Z_0 = 0$, Z_l be either 0 or 2 with probability $\frac{1}{2}$ and $Z_r = 2 - Z_l$. The random point v, which equals l if $Z_l = 2$ and r if $Z_r = 2$ is a stopping point and satisfies $EZ_v = 2$. Snell's envelope for this example is $X_0 = 1$ and $X_s = Z_s$ for $s \neq 0$. We observe that $EX_0 < EZ_v$ so we cannot expect the last equality in (1.3) to hold unless we take the supremum over a restricted family of stopping points. To get hold of this family we use (1.4) as a guide.

The term Z_s in (1.4) corresponds to the option of stopping at s. If we do not stop we proceed to a point $u > s$ and, of course, we pick the one which attains the supremum (to guarantee that the supremum is attained we assume that each s in S has a finite number of direct successors). In this way we construct a sequence of stopping points in S, $0 = \sigma_0^* \leqq \sigma_1^* \leqq \ldots$, with the following property: σ_{t+1}^* is determined by the information available at the point σ_t^*. Any increasing sequence with this property is called a *strategy*. Let τ^* be the first time that $X_{\sigma_t^*} = Z_{\sigma_t^*}$. The random variable τ^* is a stopping time with respect to $\mathscr{F}_{\sigma_t^*}$. A pair $\pi = (\sigma_t, \tau)$ where σ_t is a strategy starting from 0 and τ is a stopping time with respect to $\mathscr{F}_{\sigma_t}$ is called a *policy*. Corresponding to any policy $\pi=(\sigma_t, \tau)$ is a stopping point $\alpha(\pi)=\sigma_\tau$. We call stopping points of this form *predictable*.

For the policy $\pi^*=(\sigma_t^*, \tau^*)$ constructed above, the sequence $X_{\sigma_t^*}$ is a martingale up to τ^*. If τ^* is finite with probability one we can write the analogue of (1.3) as follows

$$EZ_{\alpha(\pi^*)}=EX_{\alpha(\pi^*)}=EX_0=\sup EZ_v \tag{1.5}$$

where the last supremum is taken over all predictable stopping points. We can now state precisely the problem of optimal stopping over a partially ordered set: find a policy π^* which is optimal in the sense that

$$EZ_{\alpha(\pi^*)}=\sup EZ_{\alpha(\pi)}. \tag{1.6}$$

Here the supremum is taken over all policies π (we use the convention that $Z_{\alpha(\pi)}=0$ on $\{\tau=\infty\}$).

1.4. The solution to the stopping problem (1.6) is given in Theorem 1 and Theorem 2 (Sect. 2). The key to the proofs is the following.

Lemma. *Let X_s, $s\in S$, be a supermartingale over a partially ordered set S. Then $X_{\sigma_t}, t\geqq 0$, is a supermartingale for any strategy σ_t. Hence, if X_s is nonnegative and v is a predictable stopping point,*

$$EX_0\geqq EX_v. \tag{1.7}$$

The first statement is easily proved using the properties of a strategy. The second follows from the first and Doob's optional sampling theorem.

The example in Sect. 1.2 shows that there are stopping points for which (1.7) fails (in this connection it appeared in [9]). It is an open problem as to whether predictable stopping points are the only ones for which (1.7) holds for all nonnegative supermartingales.

1.5. Let $Y_t, t\geqq 0$, be a Markov chain and let $f(y)$ be a function on its state space. The classical optimal stopping problem for $Z_t=f(Y_t)$ has a particularly nice answer. Define the *value function* v by

$$v(y)=\sup_\tau E_y f(Y_\tau).$$

Dynkin [3] proved that v is the minimal excessive majorant of f and if an optimal stopping time exists, it is the first hitting time of $\{y: v(y)=f(y)\}$. Snell's envelope in this case is $X_t=v(Y_t)$.

Now let $Y^1_{s^1}, \dots, Y^k_{s^k}$ be k independent Markov chains. Consider the problem of stopping these processes (possibly at different times) so as to maximize a payoff function $f(y^1, \dots, y^k)$ defined on the product of the state spaces. This may arise, for example, if we wish to control two random walks with absorbing barriers so as to maximize their probability of meeting before either is absorbed (Sect. 3.2).

We extend Dynkin's results to $Z_s=f(Y^1_{s^1}, \dots, Y^k_{s^k})$ (the index set is the nonnegative k-dimensional lattice $s=(s^1, \dots, s^k)$). Indeed let

$$v(y^1, \dots, y^k)=\sup E_{y^1, \dots, y^k} f(Y^1_{\sigma^1_\tau}, \dots, Y^k_{\sigma^k_\tau}).$$

We prove (Theorem 3) that v is the smallest "multiexcessive" majorant of f. This implies that Snell's envelope is $X_s = v(Y^1_{s1}, \ldots, Y^k_{sk})$. In terms of v we can give a simple description of the optimal policy: again we stop at the first hitting time of $\{y: v(y) = f(y)\}$ and off this set we use further properties of v to specify which process to let go while freezing the others.

1.6. Walsh [13] investigated the lower envelope v of all functions which are multiexcessive in a domain R and majorize some function f on the boundary of R. He looked for conditions under which v would be multiharmonic and found that this was true only under very special conditions. In Sect. 4 we show that v always satisfies $\Psi v = 0$ where Ψ is a certain nonlinear operator. To show this we introduce the notion of a *constraint set* and solve a more general stopping problem with time constraints. The definition of a constraint set is similar to the definition of multivalued stopping times in [10], stopping domains in [14] and optional times for random fields in [6].

1.7. In Sect. 1.2 we have constructed a stopping point v^* which is not predictable. We ask: under what conditions are all stopping points predictable? Krengel and Sucheston [8] have given necessary and sufficient conditions, however, when S is the k-dimensional lattice we can give simple conditions. For $k=1$ all stopping points are predictable. For $k=2$ condition (F4) of Cairoli and Walsh [1] is sufficient (Theorem 4 of Sect. 5). This includes the case of two Markov chains and the case of independent identically distributed random variables studied in [8]. For $k \geqq 3$, in the case of Markov chains, there exist stopping points which are not predictable (Sect. 5.2) but this is not so for the i.i.d. case [8].

1.8. The authors are grateful to Professor E.B. Dynkin, both for stimulating this study and for extensive comments that were used in writing this paper.

2. Snell's Problem

2.1. Let S be a countable partially ordered set. Let $[r, s]$ denote the set of all u in S such that $r \leqq u \leqq s$. A *direct successor* of r in S is a point s in S such that $r < s$ and there exist no points between s and r; i.e., $[r, s]$ consists of only two points r and s. Denote by $U(s)$ the set of all direct successors of s.

We assume that our parameter set is locally finite (i.e., $[r, s]$ is finite for all $r \leqq s$) and that there exists a unique element 0 in S such that $0 \leqq s$ for all s in S. We also assume that each s in S has only a finite number of direct successors.

Let $\mathscr{F}_s$ be a family of σ-algebras with index set S such that $\mathscr{F}_r \subset \mathscr{F}_s$ for $r \leqq s$. One can verify that a random variable v is a *stopping point* if and only if $\{v = s\} \in \mathscr{F}_s$ for all s in S.

A *strategy* is an increasing predictable sequence of stopping points; that is, a sequence $\sigma_t, t = 0, 1, \ldots$, with the following properties:

$$\sigma_t \leqq \sigma_{t+1}, \quad \text{for } t = 0, 1, \ldots, \tag{2.1}$$

$$\sigma_{t+1} \in \mathscr{F}_{\sigma_t}. \tag{2.2}$$

(The abbreviation $Y \in \mathscr{F}$ means that a function Y is measurable with respect to the σ-algebra $\mathscr{F}$.)

A strategy is *continuous* if σ_{t+1} is a direct successor of σ_t for all t. To avoid many subscripts put

$$\mathscr{G}_t = \mathscr{F}_{\sigma_t}, \mathrm{t} = 0, 1, \ldots.$$

A *policy* π is a pair (σ_t, τ) where σ_t is a strategy and τ is a stopping time with respect to $\mathscr{G}_t$. A policy $\pi = (\sigma_t, \tau)$ is *continuous* if the strategy σ_t is continuous.

Denote by Π the class of all policies $\pi = (\sigma_t, \tau)$ for which $\sigma_0 = 0$. To every $\pi \in \Pi$ there corresponds the *predictable stopping point* $\alpha(\pi) = \sigma_\tau$ and there is a continuous policy $\pi' \in \Pi$ such that $\alpha(\pi') = \alpha(\pi)$.

2.2. Suppose $\{Z_s, s \in S\}$ is a family of random variables adapted to $\mathscr{F}_s$ such that $E \sup_s Z_s^+ < \infty$ (We denote by a^+ the positive part of a, i.e. $a^+ = a \vee 0$.). Our aim is to find a policy π^* which is *optimal* in the sense that

$$EZ_{\alpha(\pi^*)} = \sup_{\Pi} EZ_{\alpha(\pi)}$$

($Z_{\alpha(\pi)} = 0$ if $\tau = \infty$).

The standard definition of a supermartingale is applicable to a family of random variables indexed by a partially ordered set: a nonnegative family X_s is an $\mathscr{F}$-supermartingale if $X_s \in \mathscr{F}_s$ and $X_s \geqq E^{\mathscr{F}_s} X_u$ for all $s \leqq u$. This last condition is equivalent to $X_s \geqq E^{\mathscr{F}_s} X_u$ for all u in $U(s)$.

2.3. The solution of Snell's problem is given by Theorem 1 and Theorem 2.

Theorem 1. *Snell's envelope for* Z_s *exists and satisfies*

$$X_s = \max\{Z_s; E^{\mathscr{F}_s} X_u, u \in U(s)\}, \tag{2.3}$$

$$EX_0 = \sup_{\Pi} EZ_{\alpha(\pi)}. \tag{2.4}$$

Equation 2.3 is a form of Bellman's dynamic programming equation. It implies the existence of a policy $\pi = (\sigma_t, \tau)$ with the properties:

$$\pi \text{ is a continuous policy,} \tag{2.5}$$

$$\tau = \inf\{t: X_{\sigma_t} = Z_{\sigma_t}\}, \tag{2.6}$$

$$E^{\mathscr{G}_t} X_{\sigma_{t+1}} = X_{\sigma_t} \quad \text{on} \quad \{t < \tau\}. \tag{2.7}$$

We call policies with this properties *admissible.*

Theorem 2. *If* π^* *is admissible and* $\tau^* < \infty$ a.s., *then* π^* *is optimal. If there exists some optimal policy then there exists an admissible one which is optimal.*

The proofs of Theorem 1 and Theorem 2 are outlined in Sect. 6.

3. Optimal Stopping of Several Markov Chains

3.1. We now apply the results of the previous section to the optimal stopping of several independent Markov chains, $(Y_t^j, \mathscr{F}_t^j, E^j, P^j)$, $j = 1, \ldots, k$. The process

Y^j takes values in the state space E^j, moves according to transition operator P^j and is adapted to the σ-algebras $\mathscr{F}^j_t$. If Y^j starts at $y^j \in E^j$ and stops after s^j steps, then the expected *payoff* is

$$E_y f(Y_s) = E^1_{y^1} \dots E^k_{y^k} f(Y^1_{s^1}, \dots, Y^k_{s^k}) = (P^1)^{s^1} \dots (P^k)^{s^k} f(y).$$

Here f is a bounded measurable function on the product space $E = E^1 \times \dots \times E^k$ and P^j acts on f as a function of y^j with the rest of the coordinates fixed. The index set for this model is the k-dimensional nonnegative lattice.

Put $v_0 = f^+$ and let $v_{t+1} = \max\{v_t, P^1 v_t, \dots, P^k v_t\}$. The sequence v_t is increasing so the limit $v = \lim_t v_t$ exists pointwise. By induction we prove that $v_{t+1} = \max\{f^+, P^1 v_t, \dots, P^k v_t\}$ and therefore

$$v = \max\{f^+, P^1 v, \dots, P^k v\}.$$

Adapting the proof of Neveu [11] (§IV−2) it is easy to show that $v(Y_s)$ is Snell's envelope for the family $Z_s = f(Y_s)$, with respect to every measure E_y.

An admissible policy π can be described in terms of a partition of E into sets Γ and Γ^j, $j = 1, \dots, k$: when Y is in Γ^j, Y^j takes one step, and Y stops the first time it hits Γ. That is $\sigma_{t+1} = \sigma_t + u_j$ if $Y_{\sigma_t} \in \Gamma^j$ and $\tau = \inf\{t \colon Y_{\sigma_t} \in \Gamma\}$. Here u_j is the unit vector with 1 in the j-th coordinate and 0 elsewhere. The sets Γ, Γ^j must satisfy $\Gamma^j \subset \{v = P^j v\} \cup \{v = 0\}$ and $\Gamma \subset \{v = f\}$. By Theorem 2 this policy is optimal if $\tau < \infty$ a.s.

3.2. If we apply (2.4) to the measure E_y, we get

$$v(y) = \sup_{\Pi} E_y f(Y_{\alpha(\pi)}). \tag{3.1}$$

For $k = 1$, Dynkin [3] characterized the *value function* as the minimal excessive function majorizing f. For any k we call a nonnegative function g *multiexcessive* if $g \geqq P^j g$, $j = 1, \dots, k$. Then we have

Theorem 3. *The value function v is the smallest multiexcessive majorant of the payoff function f.*

A similar theorem has been proved by Dubins and Savage [2] (§2.12).

Example. Let $Y^1_{s^1}$ and $Y^2_{s^2}$ be two independent symmetric random walks on $\{0, 1, \dots, a\}$ with absorbing endpoints. Let the payoff function be 1 on $\Lambda = \{(y, y) \colon 0 < y < a\}$ and 0 elsewhere. The optimal strategy is: let Y^1 go on $\{(y, 1) \colon 1 < y < a\} \cup \{(y, a-1) \colon 0 < y < a-1\}$, let Y^2 go on $\{(1, y) \colon 1 < y < a\} \cup \{(a-1, y) \colon 0 < y < a-1\}$; on the remaining states let either go; the first time we reach Λ we stop. In particular an optimal strategy is to freeze the process that is closer to the boundary. The value function is:

$$v(y^1, y^2) = \begin{cases} 1 - \dfrac{y^1 - y^2}{a-1}, & 0 < y^2 \leqq y^1 < a \\[2ex] 1 - \dfrac{y^2 - y^1}{a-1}, & 0 < y^1 \leqq y^2 < a \\[2ex] 0 \qquad , & \text{elsewhere.} \end{cases}$$

3.3. The convex combination

$$Pf(y)=\sum_{j=1}^{k} \lambda_j(y) P^j f(y) \tag{3.2}$$

($\lambda_j(y) \geqq 0$ and $\lambda_1(y)+\ldots+\lambda_k(y)=1$) defines a new transition operator on E. A function f is multiexcessive if and only if it is excessive for all operators of the form (3.2). The admissible policy described in subsection 3.1 corresponds to the Markov operator of the form (3.2) with $\lambda_j(y)=1$ for $y \in \Gamma^j$.

3.4. Since v is the minimal function in the set $\{g: g \geqq f, g \geqq P^j g, j=1, \ldots, \mathrm{k}\}$, it is possible to use linear programming to compute v if E is finite. For $k=1$ this was suggested in [3] as well as the method of successive approximation described in subsection 3.1.

A third method can be used when the minimal P^j-excessive majorant $Q^j f$ for every function f is readily obtained. The multiexcessive majorant v is the limit of the following increasing sequence:

$$\begin{aligned} &v_0=f^+, \quad v_1=Q^1 v_0, \ldots, \quad v_k=Q^k v_{k-1}, \\ &v_{k+1}=Q^1 v_k, \ldots, \quad v_{2k}=Q^k v_{2k-1}, \ldots . \end{aligned}$$

4. The Optimization Problem with Time Constraints

4.1. A "stop constraint" for a random sequence is a time η at which we must stop if we have not already stopped. A "go constraint" is a time η before which we are forbidden to stop. Both constraints can be expressed in terms of certain subsets of S: we must stop on the one-point set $A=\{\eta\}$ and we cannot stop on the set $B=\{0,1, \ldots, \eta-1\}$. If η is a random time, then A and B are random sets. If η is a stopping time, then A and B belong to the family of sets C with the property

$$\{s \in C\} \in \mathscr{F}_s \quad \text{for all } s \text{ in } S. \tag{4.1}$$

Property (4.1) makes sense for any partially ordered index set S, and we call sets with this property *constraint sets.**

A constraint set A is a *stop-set* for a policy $\pi=(\sigma_t, \tau)$ if $\tau \leqq \eta$ where $\eta=\inf\{t: \sigma_t \in A\}$. A constraint set B is a *go-set* for π if $\alpha(\pi) \notin B$.

4.2. Let us consider policies in Π with a given stop-set A and go-set B (A and B are disjoint). The supremum of the expected payoff $EZ_{\alpha(\pi)}$ over all policies is, generally, larger than the supremum over just continuous policies. Thus we have two different constraint problems. The second one, with continuous policies, seems more natural and we concentrate on it. If A and B are empty, both suprema coincide with the supremum considered in subsection 2.2 (see the last paragraph of subsection 2.1).

Let Π^{AB} denote the set of all continuous policies in Π for which A is a stop-set and B is a go-set. We say that a policy $\pi^* \in \Pi^{AB}$ is optimal for the constraint problem if

$$EZ_{\alpha(\pi^*)}=\sup_{\Pi^{AB}} EZ_{\alpha(\pi)}.$$

For convenience we assume that the family Z_s is nonnegative.

Snell's envelope X_s for the constraint problem is the family of random variables which satisfies

$$\begin{aligned}&X_s \text{ is a supermartingale on } S\backslash A;\\ &\text{i.e. } X_s \geqq \max\{E^{\mathscr{F}_s} X_u, u \in U(s)\} \text{ on } \{s \notin A\},\end{aligned} \tag{4.2}$$

$$X_s \text{ majorizes } Z_s \text{ on } S\backslash B, \tag{4.3}$$

$$X_s = Z_s \text{ on } A, \tag{4.4}$$

$$\text{if } Y_s \text{ is a family satisfying (4.2) to (4.4) then } Y_s \geqq X_s \text{ on } S. \tag{4.5}$$

For the constraint problem the equivalent of Theorem 1 is

Theorem 1′. *Snell's envelope exists and satisfies*

$$X_s = \begin{cases} \max\{Z_s; E^{\mathscr{F}_s} X_u, u \in U(s)\}, & \text{on } \{s \notin A \cup B\},\\ \max\{E^{\mathscr{F}_s} X_u, u \in U(s)\}, & \text{on } \{s \in B\},\\ Z_s, & \text{on } \{s \in A\},\end{cases} \tag{4.6}$$

$$EX_0 = \sup_{\Pi^{AB}} EZ_{\alpha(\pi)}. \tag{4.7}$$

Equation (4.6) implies the existence of a policy π in Π^{AB} with the properties (2.5)–(2.7). We call policies in Π^{AB} which have these properties *admissible* for the constraint problem. With this amended definition Theorem 2 holds for the constraint problem as well. The proofs of Theorem 1′ and Theorem 2 for the constraint problem are similar to those of Theorem 1 and Theorem 2 for the unconstrained problem and are omitted.

4.3. Consider the case of several Markov chains (Sect. 3). To any subset R of the product space $E = E^1 \times \ldots \times E^k$ there correspondends in a natural way a stop-set $A = \{s\colon Y_s \notin R\}$ and a go-set $B = \{s\colon Y_s \in R\}$.

Suppose that Y^j is a simple random walk on the d^j-dimensional lattice (Y^i moves from a state to any of its neighbours with equal probabilities). Then the smallest function v which majorizes f and is multiexcessive in R is

$$v(y) = \sup_{\Pi^{AB}} E_y f(Y_{\alpha(\pi)}). \tag{4.8}$$

Formula (4.8), in the case of Brownian motion, is due to Walsh [13]. In our case, the function v satisfies

$$\max\{(P^1 - I)v, \ldots, (P^k - I)v\} = 0 \quad \text{on } R. \tag{4.9}$$

Indeed, this follows from the dynamic programming equation

$$v = \begin{cases} \max\{P^1 v, \ldots, P^k v\} & \text{on } R\\ f & \text{on } R^c.\end{cases}$$

The formal analogue of the operator in (4.9) for the Brownian motion is $\Psi v = \max\{\Delta_{x^1} v, \ldots, \Delta_{x^k} v\}$, where Δ_{x^j} is the Laplacian acting on v as a function of x^j.

It would be interesting to investigate a control problem for several Brownian motions related to the operator Ψ. This can also be useful for studying the equation

$$\Psi v = 0 \tag{4.10}$$

in a domain R. Equation (4.10) is just the Laplace equation for $k=1$. Another generalization of the Laplace equation is

$$\Delta_{x^1} \dots \Delta_{x^k} v = 0$$

which has been studied in [5] and [16]. Both are satisfied by multiply harmonic functions studied by many authors (see references in [13]).

5. Predictable Stopping Points

5.1. Recall from subsection 1.3 that a stopping point v is *predictable* if $v=\alpha(\pi)$ for some $\pi \in \Pi$.

Theorem 4. *If S is the 2-dimensional non-negative lattice and if the σ-fields $\mathscr{F}_s$ satisfy*

$$\begin{aligned} &\mathscr{F}_r \textit{ and } \mathscr{F}_s \textit{ are conditionally independent given } \mathscr{F}_{r \wedge s}, \\ &\textit{for all } r, s \textit{ in } S, \end{aligned} \tag{5.1}$$

then all stopping points are predictable.

Remarks. (5.1) is condition (F4) of Cairoli and Walsh [1]. The notation $r \wedge s$ means coordinatewise minimum.

Proof. Let $\sigma_0 = 0$. Suppose we have constructed stopping points $\sigma_0, \dots, \sigma_t$ with the properties 2.1 and 2.2. Let $A=\{v^1=\sigma_t^1, v^2>\sigma_t^2\}$, $B=\{v^1>\sigma_t^1, v^2=\sigma_t^2\}$ and $C=\{v=\sigma_t\}$. Using (5.1) it is easy to see that A and B are conditionally independent given $\mathscr{F}_{\sigma_t}=\mathscr{G}_t$. Since $A \cap B = \emptyset$, we have $E^{\mathscr{G}_t} 1_A E^{\mathscr{G}_t} 1_B = 0$, which implies that there exists a partition $\tilde{A}$, $\tilde{B}$ of C^c which is $\mathscr{G}_t$-measurable and such that $A \subset \tilde{A}$ and $B \subset \tilde{B}$. Let

$$\sigma_{t+1} = \begin{cases} \sigma_t + (1,0) & \text{on } \tilde{A} \setminus C \\ \sigma_t + (0,1) & \text{on } \tilde{B} \setminus C \\ \sigma_t & \text{on } C. \end{cases}$$

Finally put $\tau = \inf\{t \colon \sigma_t = v\}$ and $\pi = (\sigma_t, \tau)$. Then $v = \alpha(\pi)$ and we are done.

5.2. Now we return to the case of k independent Markov chains considered in Sect. 3. In this case $\Omega = \Omega_1 \times \dots \times \Omega_k$, $\mathscr{F}_s = \mathscr{F}^1_{s^1} \times \dots \times \mathscr{F}^k_{s^k}$ and $E_y = E^1_{y^1} \times \dots \times E^k_{y^k}$. The hypotheses of Theorem 4 hold for $k=2$, hence the supremum in (3.1) is over all stopping points: $v(y) = \sup\limits_v E_y f(Y_v)$. The situation is different for $k \geqq 3$: there do exist stopping points which are not predictable. Indeed, suppose Y^i $i=1,2,3$ are simple random walks on $\{-1, 0, 1\}$ with -1 and 1 absorbing states. Also suppose that $Y_0^i = 0$ for $i=1,2,3$. Let A_i be the event that the first step of Y^i is to the left and put

$$v=\begin{cases}(1,1,0) & \text{on } A_1^c\times A_2\times \Omega_3\\ (1,0,1) & \text{on } A_1\times \Omega_2\times A_3^c\\ (0,1,1) & \text{on } \Omega_1\times A_2^c\times A_3\\ (1,1,1) & \text{on the rest.}\end{cases}$$

The random variable v is a stopping point which is not predictable. To see this just note that, since $\mathscr{F}_{000}=\{\emptyset,\Omega\}$, all $\mathscr{F}_{000}$-measurable stopping points must be constants.

6. Proof of Theorem 1 and Theorem 2

6.1. We supply only the parts of the proofs which differ from those in the classical theory. The reader is referred to Neveu [9] for the missing details.

Denote by Π_s the class of all policies $\pi=(\sigma_t,\tau)$ for which $\sigma_0=s$.

6.2. We now prove Theorem 1. We show that Snell's envelope X_s is given by

$$X_s=\operatorname*{ess\,sup}_{\Pi_s} E^{\mathscr{F}_s}Z_{\alpha(\pi)}.$$

Clearly $X_s\geqq Z_s$ and $X_s\geqq 0$. Also $X_s\leqq E^{\mathscr{F}_s}\sup_r Z_r^+$, so X_s is integrable. Since the family $\{E^{\mathscr{F}_s}Z_{\alpha(\pi)},\ \pi\in\Pi_s\}$ is directed upwards, there exists a sequence $\pi_n\in\Pi_s$ such that $X_s=\lim_n\uparrow E^{\mathscr{F}_s}Z_{\alpha(\pi_n)}$ and $E^{\mathscr{F}_s}Z_{\alpha(\pi_n)}\geqq Z_s^+$. By the monotone convergence theorem,

$$E^{\mathscr{F}_r}X_s=\lim_n E^{\mathscr{F}_r}Z_{\alpha(\pi_n)}\leqq X_r \qquad \text{for } r\leqq s. \tag{6.1}$$

Here we use the fact that for any $\pi=(\sigma_t,\tau)$ in Π_s there exists a $\pi'=(\sigma_t',\tau')$ in $\Pi_r (r\leqq s)$ such that $\alpha(\pi)=\alpha(\pi')$ ($\tau'=\tau+1$, $\sigma_0'=r$, $\sigma_t'=\sigma_{t-1}$ for $t\geqq 1$). Since $X_s\geqq Z_s$ it follows from (6.1) that

$$X_s\geqq \max\{Z_s;\ E^{\mathscr{F}_s}X_u,\ u\in U(s)\}. \tag{6.2}$$

To get the reversed inequality write, for $\pi\in\Pi_s$,

$$Z_{\alpha(\pi)}=Z_s 1_{\{\alpha(\pi)=s\}}+\sum_{u\in U(s)} Z_{\alpha(\pi_u)}1_{A_u}.$$

Here $A_u, u\in U(s)$ is any $\mathscr{F}_{\sigma_0}=\mathscr{F}_s$ measurable partition of $\{\alpha(\pi)\neq s\}$ such that $\sigma_1\geqq u$ on A_u, and $\pi_u=(\sigma_t^u,\tau^u)$ is a policy satisfying on A_u

$$\sigma_t^u=\begin{cases}\sigma_t & t\geqq 1\\ u & t=0\end{cases}$$

and $\tau^u = \tau$. It is easy to see that $\pi_u \in \Pi_u$. Hence

$$\begin{aligned} E^{\mathscr{F}_s} Z_{\alpha(\pi)} &= Z_s 1_{\{\alpha(\pi)=s\}} + \sum_{u \in U(s)} 1_{A_u} E^{\mathscr{F}_s} E^{\mathscr{F}_u} Z_{\alpha(\pi_u)} \\ &\leqq \max \{Z_s; E^{\mathscr{F}_s} X_u, u \in U(s)\} \end{aligned}$$

for all $\pi \in \Pi_s$, which proves 2.3.

Since $X_s \geqq 0$, we see from (6.2) that X_s is a nonnegative supermartingale majorizing Z_s. Let Y_s be another such supermartingale. By the lemma in Sect. 1.4 and Doob's optional sampling theorem,

$$Y_s \geqq E^{\mathscr{F}_s} Y_{\alpha(\pi)} \geqq E^{\mathscr{F}_s} Z_{\alpha(\pi)} \quad \text{for all } \pi \in \Pi_s,$$

hence $Y_s \geqq X_s$ a.s., which proves the minimality of X_s.

6.3. We now outline the proof of Theorem 2.

The sequence $M_t = X_{\sigma^*_{\tau^*} \wedge t}$ is a martingale by (2.7), hence $EX_0 = EM_t$. Using Fatou's lemma we show that $EX_0 = \overline{\lim}\, EM_t \leqq EX_{\alpha(\pi^*)}$. If $\tau^* < \infty$, then $EX_{\alpha(\pi^*)} = EZ_{\alpha(\pi^*)}$ by the definition of τ^*. Clearly $EZ_{\alpha(\pi^*)} \leqq \sup\limits_{\Pi_0} EZ_{\alpha(\pi)} = EX_0$ and so π^* is optimal.

Now let π' be any optimal policy. In view of the concluding remark of subsection 2.1 we may assume that π' is continuous. By (1.7) $EX_0 \geqq EX_{\alpha(\pi')} \geqq EZ_{\alpha(\pi')}$ and by the optimality of π', $EZ_{\alpha(\pi')} = \sup\limits_{\Pi_0} EZ_{\alpha(\pi)} = EX_0$. Hence

$$X_{\alpha(\pi')} = Z_{\alpha(\pi')} \quad \text{a.s.,} \tag{6.1}$$

and so π' satisfies (2.7), which means that $X_{\sigma'_{\tau'} \wedge t}$ is a martingale.

Finally let $\tau^* = \inf\{t\colon X_{\sigma'_t} = Z_{\sigma'_t}\}$. By (6.1), $\tau^* \leqq \tau'$, thus $EZ_{\alpha(\pi')} = EZ_{\alpha(\pi^*)}$. This implies that $\pi^* = (\sigma'_t, \tau^*)$ is an admissible policy which is optimal.

Remark. For $\varepsilon > 0$, let $\tau_\varepsilon = \inf\{t\colon X_{\sigma_t} \leqq Z_{\sigma_t} + \varepsilon\}$. The same proof carries over to give: $\tau_\varepsilon < \infty$ implies $EZ_{\alpha(\pi_\varepsilon)} \geqq \sup\limits_{\Pi_0} EZ_{\alpha(\pi)} - \varepsilon$ where $\pi_\varepsilon = (\sigma^*_t, \tau_\varepsilon)$, but in contrast to the traditional optimal stopping problem, we cannot guarantee that τ_ε is finite.

References

1. Cairoli, R., Walsh, J.B.: Stochastic integrals in the plane. Acta. math. **134**, 111–183 (1975)
2. Dubins, L.E., Savage, L.J.: How to Gamble if You Must; Inequalities for Stochastic Processes. New York: McGraw-Hill 1965
3. Dynkin, E.B.: The optimum choise of the instant of stopping a Markov process. Dokl. Acad. Nauk SSSR **150**, 238–240 (1963). In Russian. (English translation in: Soviet Math. Dokl **4**, 627–629)
4. Dynkin, E.B.: Markov processes and Random fields. Bull. Amer. Math. Society. **3** (New Series), 975–999 (1980)
5. Dynkin, E.B.: Harmonic functions associated with several Markov processes. Adv. Appl. Math. To appear.
6. Evstigneev, I.V.: Markov times for random fields. Teor. Verojatnost. i Primenen. **22** no. 3, 563–569 (1977). In Russian. (English translation in: Theor. Probab. Appl. **22**, 563–569 (1978))
7. Haggstrom, G.W.: Optimal stopping and experimental design. Ann. Math. Statist. **37**, 7–27 (1966)
8. Krengel, U., Sucheston, L.: Stopping rules and tactics for processes indexed by a directed set J. Multivariate Anal. To appear

9. Kurtz, T.G.: The optional sampling theorem for martingales indexed by direct sets. Ann. Probab. **8**, 675–681 (1980)
10. Millet, A., Sucheston, L.: A characterization of Vitali conditions in terms of maximal inequalities. Ann. Probab. **8**, 339–349 (1980)
11. Neveu, J.: Discrete Parameter Martingales. Amsterdam: North-Holland 1975
12. Snell, J.L.: Applications of Martingale systems theorems. Trans. Am. Math Soc. **73**, 293–312 (1952)
13. Walsh, J.B.: Probability and a Dirichlet problem for multiply superharmonic functions. Ann. Inst. Fourier (Grenoble) **18**, no. 2 221–279 (1968)
14. Walsh, J.B.: Martingales with a multi-dimensional parameter and stochastic integrals in the plane. Cours de 3 éme Cycle, Laboratoire de Calcul des Probabilités, Université Paris VI, Année 76–77
15. Walsh, J.B.: Optional increasing paths. To appear
16. Vanderbei, R.J.: Toward a stochastic calculus for several Markov processes. In preparation.

Received September 5, 1980; in revised form February 19, 1981