

A Deterministic Approach to Optimal Stopping

M.H.A. Davis* and I. Karatzas†

*Imperial College, London, U.K.

†Columbia University, U.S.A.

This paper concerns the optimal stopping problem of maximising EY_τ over the set \mathcal{M} of all stopping times of a filtration $(\mathcal{F}_t)_{t \in T}$, where $(Y_t)_{t \in T}$ is a non-negative adapted process. Here $T = Z_+$ or $T = R_+$. It is shown that

$$\sup_{\tau \in \mathcal{M}} EY_\tau = E \left[\sup_{t \in T} (Y_t + M_\infty - M_t) \right],$$

where (M_t) is the martingale component in the Doob–Meyer decomposition $Z_t = M_t - A_t$ of the Snell envelope (Z_t) of (Y_t) . Thus the process (λ_t) defined by $\lambda_\tau := M_\infty - M_\tau$ is the Lagrange multiplier corresponding to the ‘non-anticipativity constraint’ that τ be a stopping time (rather than a general non-adapted random time). The basic ‘prophet inequality’ of Krengel and Sucheston is shown to follow easily from this result.

33.1 INTRODUCTION

In the Preface to *Optimization over Time*, Peter Whittle writes: ‘Most of the classifying dichotomies that I visualized when starting proved either false or grossly asymmetric. For example, the deterministic/stochastic split is not a dichotomy but a contrast between less and more general cases.’ What he means, of course, is that deterministic optimisation is a special case of stochastic optimisation in which, if one represents randomness by the traditional probability triple (Ω, \mathcal{F}, P) , Ω reduces to a single point or, equivalently, P to a Dirac measure. It is, however, possible to argue that the inclusion is the other way round. In a typical finite-horizon discrete-time problem one observes a sequence of random

variables (Y_0, Y_1, \dots, Y_N) and the decision θ_k made at time k must depend only on the past observation record (Y_0, Y_1, \dots, Y_k) , i.e. the process $\Theta = (\theta_k)$ must be adapted to the filtration $\mathcal{Y}_k = \sigma\{Y_0, Y_1, \dots, Y_k\}$. Assume for the sake of argument that each θ_k is an integrable real-valued random variable. Then the set of adapted processes is a *linear subspace* of $X := L_1((\Omega, \mathcal{F}, P); R^{N+1})$, and one can consider optimising over all of X , enforcing by means of a suitable Lagrange multiplier the equality constraint $(I - \Pi)\Theta = 0$, where I is the identity operator and Π the projection onto the subspace of adapted processes. However, for maximising a given reward function J over X we have (ignoring technicalities),

$$\sup_{\Theta \in X} EJ(\Theta, \omega) = E \left[\sup_{\vartheta \in R^{N+1}} J(\vartheta, \omega) \right].$$

In the bracket on the right is a family of *deterministic* optimisation problems indexed by $\omega \in \Omega$; the only role of the probability measure P is to average the result. The original optimisation over adapted decision processes is accomplished if one introduces the appropriate multiplier corresponding to the non-anticipativity constraint.

This approach was initiated by Wets (1975) and Eisner and Olsen (1975) in the context of stochastic programming, and has been pursued by a number of authors, including Dempster (1982, 1988), who studies the multiplier processes in detail, Back and Pliska (1987) who give a general formulation in continuous time, and Davis and Burstein (1992) who extend the approach to nonlinear stochastic control. It is generally not a simple matter to find the right functional-analytic framework for a particular problem; the L_1/L_∞ formulation sketched above is only for purposes of illustration. There is conceptual content in this point of view in clarifying the 'deterministic/stochastic split' and a utilitarian aspect in defining a 'price for information' and in quantifying to what extent uncertainty plays a significant role in a given stochastic decision problem. Indeed, Dempster (1988) and Rockafellar and Wets (1991) have proposed algorithms for solving stochastic programming problems incorporating these ideas.

It should be mentioned that there is no simple extension to problems with incomplete information (except in special cases such as the LQG problem) or non-traditional information patterns such as considered by, for example, Witsenhausen (1968). In such cases the constraint for θ_k depends on the choice of $(\theta_0, \dots, \theta_{k-1})$, and 'non-anticipativity' is not equivalent to a simple subspace constraint as outlined above; this reinforces Peter Whittle's contention that 'the real contrast is not between deterministic and stochastic models but between models with complete and incomplete information'.

In this paper we study optimal stopping from the 'deterministic' point of view. The general problem is to maximise EY_τ over stopping times τ , where (Y_t) is a non-negative valued stochastic process (in either discrete or continuous time). It is well known that the solution to this problem involves in an essential way the so-called Snell envelope—described below—which is the smallest supermartingale (Z_t) dominating (Y_t) . Z_t represents the maximum available

rew
gen
the
ana
tion
time
inte
Lag
at s
S

res
cast
mul
we
inec
is c
the
on
for
into

Let
 \mathcal{F}_k
con
 τ s
 (Y_k)
 Y_k

The
and

an

va
ar
(1
w
ge
{

reward given that the process has not been stopped before time t . Under quite general conditions, (Z_t) has the Doob–Meyer decomposition $Z_t = M_t - A_t$ into the difference of a martingale (M_t) and an increasing process (A_t) . Hitherto, analysts' attention has mostly been focussed on (A_t) , which has a clear interpretation as the loss of available reward occasioned by failure to stop at judicious times. Here we show that the martingale component (M_t) also has a clear interpretation, namely that the process (λ_t) defined by $\lambda_t = M_\infty - M_t$ is the Lagrange multiplier enforcing the constraint that the process must be stopped at stopping times rather than at general, non-adapted, random times.

Section 33.2 below introduces the problem in a discrete-time setting, and the result is established in the following two sections, starting with the finite-horizon case where all computations are explicit and one can see clearly why the multiplier takes the form it does (and why we call it a multiplier!) As a by-product we are able to give, in Section 33.5, an almost trivial proof of the basic 'prophet inequality' of Krengel and Sucheston (1978). In the final Section 33.6 the problem is considered in a general continuous-time setting. In principle the main result there, Theorem 3, includes the earlier results as special cases, but its proof rests on a large substructure of general optimal stopping theory whereas the proofs for the discrete-time case are much more self-contained and provide extra insight into the optimisation process.

33.2 OPTIMAL STOPPING IN DISCRETE TIME

Let (Ω, \mathcal{F}, P) be a complete probability space and $(\mathcal{F}_k)_{k \in \mathbb{Z}_+}$ be a filtration, i.e. \mathcal{F}_k is a sub σ -field of \mathcal{F} for each $k \in \mathbb{Z}_+ = \{0, 1, 2, \dots\}$, $\mathcal{F}_k \subset \mathcal{F}_{k+1}$ and each \mathcal{F}_k contains all the null sets of \mathcal{F} . $\mathcal{M}_n, \mathcal{M}_n^N$ will denote the sets of \mathcal{F}_k -stopping times τ such that, respectively, $P[n \leq \tau < \infty] = 1$ and $P[n \leq \tau \leq N] = 1$. Now let $(Y_k)_{k \in \mathbb{Z}_+}$ be a stochastic process such that Y_k is \mathcal{F}_k -measurable for each $k \in \mathbb{Z}_+$, $Y_k \geq 0$ a.s. and

$$M := E \left[\sup_{k \in \mathbb{Z}_+} Y_k \right] < \infty. \quad (1)$$

The general problem of optimal stopping, as discussed in, for example, Bismut and Skalli (1977), Fakeev (1970), Neveu (1975) and Shiryaev (1978) is to calculate

$$V := \sup_{\tau \in \mathcal{M}_0} E Y_\tau \quad (2)$$

and to find—if one exists—a stopping time $\hat{\tau}$ such that $V = E Y_{\hat{\tau}}$.

The problem is of course trivial if $\mathcal{F}_0 = \mathcal{F}$. Then any \mathbb{Z}_+ -valued random variable is in \mathcal{M}_0 and $V = M$, defined by (1). Optimal times take the form $\hat{\tau}(\omega) \in \arg \max_k Y_k(\omega)$ (assuming the maximum is attained) and, as pointed out by Neveu (1975), the smallest such time can be expressed as $\hat{\tau}(\omega) = \min \{k : Y_k(\omega) = Z_k(\omega)\}$ where $Z_k(\omega) = \sup_{j \geq k} Y_j(\omega)$. The essence of the problem is therefore that in general $\mathcal{F}_k \neq \mathcal{F}$ and the times τ must satisfy the non-activipativity requirement $\{\tau \leq k\} \in \mathcal{F}_k$, $k \in \mathbb{Z}_+$, i.e. the linear equality constraint $P[\tau \leq k | \mathcal{F}_k] = I_{\{\tau \leq k\}}$ a.s.

We will show that there is a unique choice of multipliers λ_k such that V defined by (2) is given by

$$V = E \left[\sup_{k \in \mathbb{Z}_+} (Y_k + \lambda_k) \right], \quad (3)$$

while for any stopping time τ , $E[\lambda_\tau] = 0$, so that $E[Y_\tau + \lambda_\tau] = EY_\tau$. Thus optimal or ε -optimal stopping times are actually *pathwise* maximisers of $Y_k(\omega) + \lambda_k(\omega)$. λ_k is given explicitly by $\lambda_k = M_\infty - M_k$, where (M_k) is the martingale component in the Doob decomposition of the Snell envelope (Z_k) of (Y_k) , which we introduce next. All of the following information can be found in Chapter VI of Neveu (1975) (see also Karatzas (1993)).

The Snell envelope of (Y_k) is the process (Z_k) defined by

$$Z_k = \operatorname{ess\,sup}_{\tau \in \mathcal{M}_k} E[Y_\tau | \mathcal{F}_k], \quad k \in \mathbb{Z}_+. \quad (4)$$

(Z_k) is uniformly integrable in view of condition (1). It is a supermartingale, is the smallest supermartingale dominating (Y_k) , and satisfies the recursion $Z_k = Y_k \vee E[Z_{k+1} | \mathcal{F}_k]$ as well as the equality $Z_\infty = Y_\infty$, where both $Y_\infty := \limsup_{k \rightarrow \infty} Y_k$ and $Z_\infty := \limsup_{k \rightarrow \infty} Z_k$ are integrable random variables. The main general result of optimal stopping theory is that an optimal stopping time exists if and only if the stopping time $\sigma_0 := \min\{k: Y_k = Z_k\}$ is a.s. finite (i.e. belongs to \mathcal{M}_0), and then σ_0 is the smallest optimal time. Even if no optimal time exists, $\sigma_\varepsilon := \min\{k: Z_k - Y_k \leq \varepsilon\}$ is in \mathcal{M}_0 for any $\varepsilon > 0$ and satisfies $EY_{\sigma_\varepsilon} \geq V - \varepsilon$, where V is given by (2).

Being a non-negative uniformly integrable supermartingale, (Z_k) has the Doob decomposition $Z_k = M_k - A_k$ where (M_k) is a martingale and (A_k) is a predictable increasing process (i.e. A_k is \mathcal{F}_{k-1} measurable, $k = 1, 2, \dots$). M_k and A_k are given recursively by

$$\begin{aligned} M_k &= M_{k-1} + (Z_k - E[Z_k | \mathcal{F}_{k-1}]), & M_0 &= Z_0 \\ A_k &= A_{k-1} + (Z_{k-1} - E[Z_k | \mathcal{F}_{k-1}]), & A_0 &= 0. \end{aligned} \quad (5)$$

For finite-horizon problems the Snell envelope can be defined in an algorithmic manner as follows: let

$$\begin{aligned} Z_N^N &= Y_N \\ Z_k^N &= Y_k \vee E[Z_{k+1}^N | \mathcal{F}_k], \quad k = N-1, \dots, 0. \end{aligned} \quad (6)$$

Then

$$Z_k^N = \operatorname{ess\,sup}_{\tau \in \mathcal{M}_k^N} E[Y_\tau | \mathcal{F}_k].$$

An optimal stopping time $\hat{\sigma}^N \in \mathcal{M}_0^N$ always exists, the least such time being given by

$$\hat{\sigma}^N = \min\{k \leq N: Y_k = Z_k^N\}.$$

33.3 THE FINITE HORIZON CASE

Here we consider maximising EY_τ over $\tau \in \mathcal{M}_0^N$. It is instructive to consider an enlarged class of *randomised strategies* \mathcal{R}^N as discussed by, for example, Shiryaev (1978). An element of \mathcal{R}^N is an adapted sequence $\theta = \{\theta_0, \dots, \theta_N\}$ with $\theta_k \in [0, 1]$, $k = 0, \dots, N-1$ and $\theta_N \equiv 1$. One flips a coin with heads probability θ_k in order to decide whether to stop the process at time k , given that it was not stopped earlier. An ordinary stopping time τ corresponds to the special case $\theta_k = I_{\{\tau \geq k\}}$. The stopping time distribution is then ρ_k , $k = 0, \dots, N$ given by

$$\rho_k = \theta_k \prod_{j=0}^{k-1} (1 - \theta_j)$$

and the corresponding reward is

$$J(\theta) = E \sum_{k=0}^N Y_k \rho_k,$$

so that $J(\theta) = EY_\tau$ when $\theta_k = I_{\{\tau \geq k\}}$. It is easily shown that the available reward is not increased by enlarging the class of strategies in this way, i.e.

$$\sup_{\theta \in \mathcal{R}^N} J(\theta) = \sup_{\tau \in \mathcal{M}_0^N} EY_\tau.$$

It is convenient to define $\zeta_0 = 1$ and

$$\zeta_k = \prod_{j=0}^{k-1} (1 - \theta_j) \quad k = 1, \dots, N.$$

Then

$$\zeta_{m+1} = (1 - \theta_m) \zeta_m \quad m = 0, \dots, N-1, \quad (7)$$

and $\rho_m = \zeta_m \theta_m$, so that the reward is expressed as

$$J(\theta) = E \sum_{k=0}^N Y_k \zeta_k \theta_k.$$

We now wish to consider *pathwise* optimisation, relaxing the adaptedness constraint but introducing a multiplier process $\lambda = (\lambda_0(\omega), \dots, \lambda_N(\omega))$. Thus fix $\omega \in \Omega$, write $Y_k := Y_k(\omega)$ and consider maximising $K(\omega, \theta)$ given by

$$K(\omega, \theta) = \sum_{k=0}^N (Y_k + \lambda_k) \zeta_k \theta_k$$

over arbitrary sequences θ , where $\zeta_k = \zeta_k(\omega)$ is given by (7). This is an optimal control problem, and elementary application of dynamic programming gives the following result.

Lemma 1 Let $V_k(\zeta) = \max_{\theta_k, \dots, \theta_{N-1}} \sum_{j=k}^N (Y_j + \lambda_j) \zeta_j \theta_j$ subject to (7) for $m = k, \dots,$

$N-1$ with $\zeta_k = \zeta$. Then $V_k(\zeta) = \beta_k(\lambda)\zeta$, where

$$\begin{aligned}\beta_N(\lambda) &= Y_N + \lambda_N \\ \beta_k(\lambda) &= (Y_k + \lambda_k) \vee \beta_{k+1}(\lambda), \quad k = N-1, \dots, 0.\end{aligned}\quad (8)$$

This shows that, as before, 'randomisation' has no effect and

$$\beta_0(\lambda) = \max_{0 \leq k \leq N} (Y_k + \lambda_k).$$

Most of the remaining argument is based on the simple identity $a \vee b = b + [a - b]^+$, where $[c]^+ = c \vee 0$. We write $E^k X := E[X | \mathcal{F}_k]$ for integrable random variables X . Take for example $k = N-1$ in (8) and assume $\lambda_N = 0$, to obtain

$$\beta_{N-1}(\lambda) = Y_N + [Y_{N-1} + \lambda_{N-1} - Y_N]^+.$$

If we choose $\lambda_{N-1} = Y_N - E^{N-1} Y_N$ then the optimal decision at time $N-1$ in the optimisation problem of Lemma 1 is $\hat{\theta}_{N-1} = I_{(Y_{N-1} \geq E^{N-1} Y_N)}$ which coincides with the best non-anticipative decision. Further,

$$E \sum_{j=N-1}^N (Y_j + \lambda_j) \zeta_j \theta_j = E \sum_{j=N-1}^N Y_j \zeta_j \theta_j$$

for any adapted θ , and these choices of λ_{N-1}, λ_N are the only ones for which these two properties hold. The general result is as follows.

Proposition 1 Let (Z_k^N) be the Snell envelope for the N -horizon problem, defined by (6). Define

$$\begin{aligned}\alpha_N &= Y_N, \quad \lambda_N = 0 \\ \alpha_k &= \alpha_{k+1} + [Y_k - E^k Z_{k+1}^N]^+, \quad k = N-1, \dots, 0\end{aligned}\quad (9)$$

$$\lambda_k = \alpha_{k+1} - E^k \alpha_{k+1}, \quad k = N-1, \dots, 0.\quad (10)$$

Then, in the notation of Lemma 1, $\alpha_k = \beta_k(\lambda)$,

$$E[\alpha_k | \mathcal{F}_k] = Z_k^N, \quad (11)$$

and the strategy $\hat{\theta}_k(\omega) = I_{\{\hat{\tau}^N(\omega) = k\}}$ maximises $K(\omega, \theta)$ for almost every ω .

Proof This is proved by a simple induction argument. For example, to verify (11), suppose that $E[\alpha_{k+1} | \mathcal{F}_{k+1}] = Z_{k+1}^N$; then from (9)

$$\begin{aligned}E^k \alpha_k &= E^k \alpha_{k+1} + [Y_k - E^k Z_{k+1}^N]^+ \\ &= E^k Z_{k+1}^N + [Y_k - E^k Z_{k+1}^N]^+ \\ &= (E^k Z_{k+1}^N) \vee Y_k = Z_k^N,\end{aligned}$$

showing that (11) holds for all k , since clearly it holds when $k = N$.

It is possible to express (9) in more convenient form. Denote temporarily $Y := Y_k$, $Z := Z_k^N$ and $E := E^k Z_{k+1}^N$. Then $Z = Y \vee E = Y + [E - Y]^+$ and hence $[Y - E]^+ = [E - Y]^- = [E - Y]^+ - (E - Y) = (Z - Y) - (E - Y) = Z - E$, so that we have

$$\alpha_N = Z_N^N, \quad \alpha_k = \alpha_{k+1} + Z_k^N - E[Z_{k+1}^N | \mathcal{F}_k]$$

or

$$\alpha_k = \sum_{j=k}^{N-1} (Z_j^N - E^j Z_{j+1}^N) + Z_N^N.$$

33.4 THE GENERAL CASE

We now resume consideration of the infinite-horizon problem. Recall that (Z_k) denotes the Snell envelope, defined by (4).

Lemma 2 $E\{\sum_{k=0}^{\infty} (Z_k - E^k Z_{k+1})\} < \infty$; in particular

$$\sum_{n=0}^{\infty} (Z_k - E^k Z_{k+1}) < \infty \quad \text{a.s.}$$

Proof Each term in the sum is a.s. non-negative and the n 'th partial sum satisfies

$$E \sum_{k=0}^N (Z_k - E^k Z_{k+1}) = E(Z_0 - Z_{N+1}) \leq EZ_0 < \infty.$$

The result follows by monotone convergence.

Since Z_k is a non-negative uniformly integrable supermartingale, the limit $Z_{\infty} = \lim_{k \rightarrow \infty} Z_k$ exists a.s. and in L_1 (it coincides with Z_{∞} as previously introduced), and by analogy with (9), (10) we now define

$$\alpha_k = \sum_{j=k}^{\infty} (Z_j - E^j Z_{j+1}) + Z_{\infty} \quad (12)$$

$$\lambda_k = \alpha_k - Z_k \quad (= \alpha_{k+1} - E^k Z_{k+1}). \quad (13)$$

We see from (5) that λ_k can be expressed as $\lambda_k = M_{\infty} - M_k$ and that $\alpha_k = M_{\infty} - A_k$. This brings us to the main result.

Theorem 1 Suppose that $Y_k \geq 0$ and that condition (1) is satisfied. Define α_k, λ_k by (12), (13), so that $\lambda_k = M_{\infty} - M_k$, where M_k is the uniformly integrable martingale appearing in the Doob decomposition $Z_k = M_k - A_k$ of the Snell envelope (Z_k) of (Y_k) . Then

$$\sup_{\tau \in \mathcal{M}_0} EY_{\tau} = E \left\{ \sup_{k \in \mathbb{Z}_+} (Y_k + \lambda_k) \right\}, \quad (14)$$

while $E\lambda_{\tau} = 0$ for any stopping time $\tau \in \mathcal{M}_0$.

Proof It follows as in the previous section that α_k defined by (12) satisfies

$$\alpha_k = \alpha_{k+1} + [Y_k - E^k Z_{k+1}]^+$$

and hence by a dynamic programming argument that

$$\alpha_k = \sup_{j \geq k} (Y_j + \lambda_j)$$

where λ_j is defined by (13). By monotone convergence, and the a.s. and L_1 convergence of Z_k to Z_∞ ,

$$\begin{aligned} E[\alpha_k | \mathcal{F}_k] &= \lim_{N \rightarrow \infty} E \left[\sum_{j=k}^N (Z_j - E^j Z_{j+1}) + Z_\infty | \mathcal{F}_k \right] \\ &= Z_k + \lim_{N \rightarrow \infty} E[Z_\infty - Z_{N+1} | \mathcal{F}_k] \\ &= Z_k. \end{aligned}$$

In particular (a) $E\alpha_0 = EZ_0$, from which (14) follows, and (b) $E[\lambda_k | \mathcal{F}_k] = 0$, so that for any $\tau \in \mathcal{M}_0$

$$\begin{aligned} E\lambda_\tau &= E \left\{ \sum_{k=0}^{\infty} I_{\{\tau=k\}} \lambda_k \right\} \\ &= E \left\{ \sum_{k=0}^{\infty} I_{\{\tau=k\}} E[\lambda_k | \mathcal{F}_k] \right\} = 0. \end{aligned}$$

This completes the proof.

33.5 PROPHET INEQUALITIES

'Prophet inequalities' are statements of the form $M \leq cV$, where M, V are defined by (1), (2), c is a constant and the inequality holds for all probability measures P in some set \mathcal{C} . The interpretation is that a 'prophet' possessed of complete clairvoyance can obtain a reward which is only c times that of a player restricted to non-anticipative strategies. An excellent survey of result and techniques in this area can be found in Hill and Kertz (1992). The following classical result of Krengel, Sucheston and Garling (see Krengel and Sucheston, 1978) follows easily from Theorem 1.

Theorem 2 Suppose that $Y_k \geq 0$ for all k , that condition (1) holds and that for each $k \in Z_+$ and $j > k$ the random variable Y_j is independent of \mathcal{F}_k . Then $M \leq 2V$.

Remark The conditions imply that the random variables Y_k are mutually independent and constitute a *semiamart* in the terminology of (Krengel and Sucheston, 1978). It is clear that under these conditions

$$V = \sup_{\tau \in \mathcal{M}'_0} EY_\tau$$

where \mathcal{M}'_0 denotes the set of finite-valued stopping times of the *natural filtration* $\mathcal{F}'_k = \sigma\{Y_0, Y_1, \dots, Y_k\}$. Thus the number V in the statement of Theorem 2 is the same for any filtration satisfying the stated conditions; M , of course, does not involve the filtration.

Proof Under the independence condition stated, it follows from the property $Z_k = Y_k \vee E(Z_{k+1} | \mathcal{F}_k)$ that there is a sequence of constants $v_0 \geq v_1 \geq v_2 \dots \geq 0$ such that

$$Z_k = Y_k \vee v_{k+1}$$

and

$$v_k = E[Y_k \vee v_{k+1}].$$

Thus, from (12) and (13)

$$\begin{aligned} Y_k + \lambda_k &= \sum_{j=k}^{\infty} [Y_j - v_{j+1}]^+ + Y_k - Y_k \vee v_{k+1} + Z_{\infty} \\ &= \sum_{j=k+1}^{\infty} [Y_j - v_{j+1}]^+ + Y_k - v_{k+1} + Z_{\infty}, \end{aligned} \quad (15)$$

where the second equality follows from the representation $Y_k \vee v_{k+1} = v_{k+1} + [Y_k - v_{k+1}]^+$. Thus from (14) and using (15)

$$\begin{aligned} V &= E \left\{ \sup_{k \in \mathbb{Z}_+} (Y_k + \lambda_k) \right\} \\ &\geq E \left\{ \sup_{k \in \mathbb{Z}_+} (Y_k - v_{k+1}) \right\} \\ &\geq E \left\{ \sup_{k \in \mathbb{Z}_+} (Y_k - v_0) \right\} \\ &= M - V. \end{aligned}$$

Thus $M \leq 2V$, as claimed.

33.6 CONTINUOUS TIME

A result analogous to Theorem 1 also holds in a continuous-time setting described as follows. Let (Ω, \mathcal{F}, P) be a complete probability space and $(\mathcal{F}_t)_{t \in \mathbb{R}_+}$ a filtration satisfying the 'usual conditions': each \mathcal{F}_t is completed with all null sets of \mathcal{F} and (\mathcal{F}_t) is right-continuous, i.e. $\bigcap_{\epsilon > 0} \mathcal{F}_{t+\epsilon} = \mathcal{F}_t$, $t \in \mathbb{R}_+$. We denote by \mathcal{M} the set of all \mathcal{F}_t -stopping times (possibly taking the value $+\infty$). Now let $(Y_t)_{t \in \mathbb{R}_+}$ be a real-valued stochastic process satisfying the following conditions:

$$Y_t \geq 0 \quad \text{for all } t, \text{ a.s.} \quad (16)$$

$$(Y_t) \text{ has RCLL paths and is adapted to } (\mathcal{F}_t) \quad (17)$$

$$E \left[\sup_{t \in R_+} Y_t \right] < \infty. \tag{18}$$

Condition (18) implies that $\sup_{t \in R_+} Y_t < \infty$ a.s., and we define $Y_\infty := \limsup_{t \rightarrow \infty} Y_t$.

The Snell envelope (Z_t) of (Y_t) is the smallest RCLL supermartingale dominating (Y_t) . Existence of (Z_t) is established, and its properties studied, in for example Bismut and Skalli (1977), El Karoui (1981), Fakeev (1970) and Karatzas (1993). It satisfies, in particular,

$$Z_t = \operatorname{ess\,sup}_{\sigma \in \mathcal{M}_t} E[Y_\sigma | \mathcal{F}_t], \quad t \in R_+,$$

where $\mathcal{M}_t = \{\sigma \in \mathcal{M} : \sigma \geq t \text{ a.s.}\}$. It is evident that under conditions (16) and (18) the supermartingale (Z_t) is of class (D), and it therefore has the Doob–Meyer decomposition $Z_t = M_t - A_t$, where (M_t) is a uniformly integrable martingale and (A_t) is an integrable, predictable increasing process. Furthermore, if (Y_t) has quasi-left-continuous paths, i.e. if we have

$$\limsup_{n \rightarrow \infty} Y_{\tau_n} \leq Y_\tau \quad \text{a.s. for all } \tau_n, \tau \in \mathcal{M}, \tau_n \uparrow \tau \text{ a.s.}, \tag{19}$$

then (Z_t) is a regular supermartingale and (A_t) has continuous paths. See Rogers and Williams (1987), Section VI.29.

For $t, \varepsilon \geq 0$ define

$$U_t^\varepsilon := \inf \{s \geq t : Y_s \geq Z_s - \varepsilon\}.$$

Then $U_t^\varepsilon \in \mathcal{M}$ and it is a key property of the Snell envelope that

$$Z_t = E[Z_{U_t^\varepsilon} | \mathcal{F}_t], \quad \text{a.s.} \tag{20}$$

Using the optional sampling theorem and the decomposition $Z_t = M_t - A_t$ we see that (20) implies

$$E[A_{U_t^\varepsilon} - A_t | \mathcal{F}_t] = 0, \quad \text{a.s.}$$

and hence, since $A_{U_t^\varepsilon} - A_t \geq 0$ a.s., that

$$A_{U_t^\varepsilon} = A_t, \quad \text{a.s.} \tag{21}$$

If, in addition, (Y_t) has the quasi-left-continuity property (19), then U_0^0 achieves the supremum $V := \sup_{t \in \mathcal{M}} EY_t$, and we have

$$Z_t = E[Y_{U_t^0} | \mathcal{F}_t] = E[Z_{U_t^0} | \mathcal{F}_t], \quad \text{a.s.},$$

and

$$\int_0^\infty 1_{\{Z_t > Y_t\}} dA_t = 0 \quad \text{a.s.}, \tag{22}$$

showing that the continuous increasing process (A_t) is almost surely ‘flat off $\{t \geq 0 : Z_t(\omega) = Y_t(\omega)\}$ ’.

We can now state the main result.

Theorem 3 Suppose that (Y_t) satisfies (16)–(18) and define $\lambda_t := M_\infty - M_t$. Then

(i)
$$E\lambda_\tau = 0 \quad \text{for all } \tau \in \mathcal{M}$$

(ii)
$$\sup_{\tau \in \mathcal{M}} EY_\tau = E \left[\sup_{t \in R_+} (Y_t + \lambda_t) \right]$$

(iii) For almost every $\omega \in \Omega$ the time $t = U_0^0(\omega)$ is ε -optimal for maximisation of $Y_t(\omega) + \lambda_t(\omega)$ over $t \in R_+$; the time $t = U_0^0(\omega)$ is optimal if (19) is satisfied.

Proof Part (i) is evident from the definition of λ_t . Define $Q_t := Y_t + \lambda_t$ and $\alpha_t := \sup_{s \geq t} Q_s$; we will show that

$$\alpha_t - Q_{U_t^0} \leq \varepsilon \quad (23)$$

and that

$$E[\alpha_t | \mathcal{F}_t] = Z_t, \quad \text{a.s.} \quad (24)$$

from which (ii) and (iii) follow. Indeed,

$$\begin{aligned} Q_s &= Y_s + M_\infty - M_s \\ &= Y_s + (Z_\infty + A_\infty) - (Z_s + A_s) \\ &= Z_\infty - (Z_s - Y_s) + (A_\infty - A_s). \end{aligned} \quad (25)$$

(Here and below, the equalities and inequalities hold almost surely.) Now $(Z_s - Y_s) \geq 0$ and $(A_\infty - A_s)$ is decreasing, so clearly

$$Q_s \leq Z_\infty + (A_\infty - A_t) = M_\infty - A_t, \quad s \geq t. \quad (26)$$

On the other hand, A_s is constant on $[t, U_t^0]$ so that from (25) and the definition of U_t^0 :

$$Q_{U_t^0} \geq Z_\infty - \varepsilon + (A_\infty - A_t) = M_\infty - A_t - \varepsilon. \quad (27)$$

Thus $\alpha_t = \sup_{s \geq t} Q_s = M_\infty - A_t$ and

$$E[\alpha_t | \mathcal{F}_t] = E[M_\infty - A_t | \mathcal{F}_t] = M_t - A_t = Z_t.$$

Together with (27), this establishes (23) and (24) and completes the proof for the general case. If the quasi-left-continuity condition (19) is satisfied then (22) holds, and this implies that $A_t = A_{U_t^0}$ a.s., from which the last part of (iii) follows.

ACKNOWLEDGMENTS

Some preliminary discussions with Victor de la Peña at Columbia University stimulated us to pursue this project. The work of IK was supported in part by the U.S. Army Research Office under Grant DAAL 03-91-G-0170. The paper was written at the Institute for Mathematics and its Applications, University of Minnesota, where both authors were supported by funds provided by the NSF.

REFERENCES

- Back, K. and Pliska, S.R. (1987) The shadow price of information in continuous-time decision problems, *Stochastics* **22**, 151–186.
- Bismut, J.-M. and Skalli, B. (1977) Temps d'arrêt optimal, théorie générale des processus et processus de Markov, *Z. Wahrscheinlichkeitstheorie ver. Geb.* **39**, 301–313.
- Davis, M.H.A. and Burstein, G. (1992) A deterministic approach to stochastic optimal control, with application to anticipative control, *Stochastics and Stochastics Reports* **40**, 203–256.
- Dempster, M.A.H. (1982) The expected value of perfect information in the optimal evolution of stochastic systems, in *Stochastic Differential Systems*, eds. M. Arato, D. Vermes and A.V. Balakrishnan, *Lecture Notes in Control and Information Sciences* **36** Springer-Verlag, Berlin, pp 25–40.
- Dempster, M.A.H. (1988) On stochastic programming II: dynamic problems under risk, *Stochastics* **25**, 15–42.
- Eisner, M.J. and Olsen, P. (1975) Duality for stochastic programming interpreted as LP in L_p space, *SIAM J. Applied Mathematics* **28**, 779–792.
- El Karoui, N. (1981) Les aspects probabilistes du contrôle stochastique, *Lecture Notes in Mathematics* **876**, Springer-Verlag, Berlin.
- Fakeev, A.G. (1970) Optimal stopping rules for stochastic processes with continuous parameters, *Theory of Probability and its Applications* **15**, 324–331.
- Hill, T.P. and Kertz, R.P. (1992) A survey of prophet inequalities in optimal stopping theory, in *Strategies for Sequential Search*, eds. F.T. Bruss, T.S. Ferguson and S.M. Samuels, *Contemporary Mathematics* **25**, 191–207, *American Math. Soc.*, Providence RI.
- Karatzas, I. (1993) Probabilistic aspects of optimal stopping and stochastic control, Class notes, Department of Statistics, Columbia University, New York, Spring 1993.
- Krengel, V. and Sucheston, L. (1978) On semiamarts, amarts and processes of finite value, in *Probability on Banach Spaces*, ed. J. Kuelbs, Marcel Dekker, New York, pp 197–266.
- Neveu, J. (1975) *Discrete-Parameter Martingales*, North Holland, Amsterdam.
- Rockafellar, R.T. and Wets, R.J.-B. (1991) Scenarios and policy aggregation in optimization under uncertainty, *Math. of Operations Research* **16**, 119–147.
- Rogers, L.C.G. and Williams, D. (1987) *Diffusions, Markov Processes and Martingales II: Ito Calculus*, Wiley, Chichester.
- Shiryaev, A.N. (1978) *Optimal Stopping Rules*, Springer-Verlag, Berlin.
- Wets, R.J.-B. (1975) On the relation between stochastic and deterministic optimization, in *Control Theory, Numerical Methods and Computer Systems Modelling*, ed. J.L. Lions, *Lecture Notes in Economics and Mathematical Systems* **107**, Springer-Verlag, Berlin, pp 350–361.
- Witsenhausen, H.S. (1968) A counterexample in stochastic optimum control, *SIAM J. Control* **6**, 131–140.

Dept of Electrical and Electronic Engineering, Imperial College of Science, Technology and Medicine, Exhibition Road, LONDON SW7 2BT.

Dept of Statistics, Columbia University, New York, NY 10027, U.S.A.

CHA

Web
space
infini
so th
times
cation
condi
the t
optim
appr

One
Whi
(Whi
and
recu
ming

Proba
© 1995