A Probabilistic Formula for the Concave Hull of a Function

Robert J. Vanderbei *

Program in Statistics & Operations Research
Princeton University
Princeton, NJ 08544

November 15, 1994

Technical Report SOR-94-06

Abstract

Let D be a compact, convex domain in d-dimensional Euclidean space and let f be a nonnegative real-valued function defined on D. The classical optimal stopping problem is to find a stopping time τ^* that attains the following supremum:

$$v(x) = \sup_{\tau} E_x f(B(\tau)).$$

Here, B is a d-dimensional Brownian motion with absorption on the boundary of D and the supremum is over all stopping times. It is well-known that v is characterized as the smallest superharmonic majorant of f.

In this paper, we modify this problem by allowing B to be essentially any drift-free diffusion (with absorption, as before, on the boundary of D). For example, it could be a Brownian motion diffusing on some lower dimensional affine set. In addition, one is allowed

*Research supported by AFOSR through grant AFOSR-91-0359.

to switch among these diffusions at any time. The problem is to find a stopping time and a switching strategy that together attain the supremum over all stopping times and all switching strategies. For this problem, we show that v is characterized as the smallest concave majorant of f. The domain D can be decomposed into a disjoint union of relatively open convex sets on each of which the function v is affine. Furthermore, the union of the 0-dimensional convex sets is contained in the set on which $v = f$. An optimal switching strategy is any strategy that at all times diffuses in the affine hull of the current convex set. When the diffusion reaches the boundary of the current convex set it will lie on a lower dimensional convex set and must then diffuse on the affine hull of this new set. This process continues until the set on which $v = f$ is reached, which is the optimal stopping time.

Key words: Hamilton-Jacobi-Bellman equation, optimal control, optimal stopping, Brownian motion, concave hull.

AMS 1991 Subject Classification: Primary 60G40 Secondary 60J45, 31C10

Abbreviated title: Concave Hull of a Function

1 Introduction.

Let D be a compact, convex d-dimensional domain in d-dimensional Euclidean space and let f be a nonnegative, bounded, upper semi-continuous, real-valued function defined on D. The classical optimal stopping problem is to find a stopping time τ^* that attains the following supremum:

$$v(x) = \sup_{\tau} E_x f(B(\tau)).$$

Here, B is a d-dimensional Brownian motion with absorption on the boundary of D and the supremum is over all stopping times. Dynkin [2] was the first to study this problem. He showed that (for almost-Borel, finely-continuous-from-below functions f) v is the smallest superharmonic majorant of f and that τ^* is the first hitting time of the support set $\{x : v(x) = f(x)\}$.

More recently, Mandelbaum et.al. [7] studied a related problem in which D is a rectangle in $\mathbb{R}^2$, B is replaced with a 1-dimensional Brownian motion diffusing in either the horizontal or vertical direction, and the supremum is over all strategies for switching between these two diffusion directions. In cases where the payoff function f is continuous on the boundary and the values on the boundary dominate the values in the interior, they obtained explicit and surprisingly nontrivial solutions. This problem is related to a certain nonlinear Dirichlet problem. Such problems, and their probabilistic interpretation, were first studied by Walsh [12].

Vanderbei [11] extended the results in [7] to higher dimensional domains but only when the payoff function f has a certain special form. For these problems, he showed that the optimal strategy is of so-called Gittins type (i.e., there is a merit function associated with each coordinate direction and the optimal strategy is to diffuse in the direction associated with the highest merit function).

Recently, Cairoli and Dalang [1] studied the same problem as in [7] but on discrete state spaces using simple random walk instead of Brownian motion. They have distilled the essential ideas (which are mostly geometric) from the problem and in so doing they provide a simple and general analysis of the problem and characterization of the optimal solution.

In this paper, we modify the classical optimal stopping problem by allowing B to be essentially any drift-free diffusion. For example, it could be a Brownian motion diffusing on some lower dimensional affine set. As in [7], one is allowed to switch among these diffusions at any time. The problem is to find a stopping time and a switching strategy that together attain the supremum over all stopping times and all switching strategies. For this problem, we show that v is characterized as the smallest concave majorant of f, i.e., the concave hull of f.

2 Problem Formulation.

We begin by defining the class of controlled drift-free processes over which we shall optimize. Let $(\Omega, \mathcal{F})$ denote a sample space and, for each $x \in \mathbb{R}^d$, let P_x be a probability measure on this space. We assume that $(\Omega, \mathcal{F})$ is complete with respect to each of the measures P_x. For each u in the unit sphere S^{d-1} in $\mathbb{R}^d$, we assume that there is a 1-dimensional Brownian motion

$$B_u = \{B_u(t) : t \geq 0\}$$

defined on this probability space. Associated with each B_u is a filtration

$$\mathcal{F}_u = \{\mathcal{F}_u(t) : t \geq 0\}$$

with respect to which B_u is adapted. Furthermore, we assume that the Brownian motions are mutually independent and that they each start at zero.

In addition to the Brownian motions, we assume that there is an $\mathbb{R}^d$ valued random variable ξ. This random variable will be the starting point for the controlled diffusion. Hence, we assume that

$$\xi = x \text{ a.s. } P_x \text{ for all } x \in \mathbb{R}^d$$

and that ξ is measurable with respect to $\mathcal{F}_u(0)$ for every $u \in S^{d-1}$.

A *switching strategy* T is a collection of random time allocations

$$T = \{T_u(t) \in \mathbb{R}_+ : u \in S^{d-1}, t \geq 0\}$$

with the following properties:

1. for each $u \in S^{d-1}$, $T_u(0) = 0$ and $T_u(\cdot)$ is nondecreasing;
2. the set of "used" directions, $\{u \in S^{d-1} : \lim_{t\to\infty} T_u(t) > 0\}$, is finite;

3. for each $t \geq 0$,

$$\sum_{\substack{u \in S^{d-1} \\ \text{s.t. } T_u(t)>0}} T_u(t) = t;$$

4. for each $t \geq 0$ and each function $s : S^{d-1} \to \mathbb{R}_+$,

$$\cap_{u \in S^{d-1}} \{T_u(t) \leq s_u\} \in \vee_{u \in S^{d-1}} \mathcal{F}_u(s_u).$$

Switching strategies, essentially as defined here, have proved to be a useful tool in the study of optimal control problems (see [4, 5, 6, 7, 11, 9]). They first appeared in [8] (in a discrete-time setting) and independently in [13] where they were called optional increasing paths. It is important to note the uncoutable sum in (3) is actually finite because of property (2). The entire model could be constructed on a much smaller sample space (having only one Brownian motion!), but it seems that introducing a Brownian motion for each possible diffusion direction makes the subsequent construction more transparent.

Let $U(t) = \{u \in S^{d-1} : T_u(t) > 0\}$. Associated with each switching strategy T, there is a controlled drift-free process in $\mathbb{R}^d$ given by

$$Y^T(t) = \xi + \sum_{u \in U(t)} B_u(T_u(t))u, \quad t \geq 0.$$

We now proceed to give a precise formulation of our optimal control problem. It involves the restriction of Y^T to the compact, convex domain D, given by

$$X^T(t) = Y^T(t \wedge \sigma),$$

where $\sigma = \inf\{t > 0 : Y^T(t) \notin D\}$ is the first exit time from the domain D. Given a nonnegative, bounded, upper semi-continuous function f defined on D, the problem is to find a switching strategy T^* and a stopping time τ^* that together attain the following supremum:

$$v(x) = \sup_{T,\tau} \mathbf{E}_x f(X^T(\tau)), \quad x \in D. \tag{2.1}$$

As usual, the function v is called the *value function.*

3 The Supermartingale Connection.

The basic tool in our analysis is the identification of which functions when composed with X^T form supermartingales for every choice of switching strategy T. To make this identification, we need to specify the filtration with respect to which these compositions should be supermartingales. Indeed, associated with any switching strategy T, there is a one-parameter filtration $\mathcal{F}^T = \{\mathcal{F}^{T(t)} : t \geq 0\}$ where $\mathcal{F}^{T(t)}$ is defined as the σ-algebra consisting of all sets C for which $C \cap \{T_u(t) \leq s_u : u \in S^{d-1}\} \in \vee_{u \in S^{d-1}} \mathcal{F}_u(s_u)$ for every function $s : S^{d-1} \to \mathbb{R}_+$. It is easy to check that the process X^T is adapted to $\mathcal{F}^T$.

Theorem 1 *Let w be a function defined on D.*

1. *If w is affine, then for any switching strategy T, $w(X^T(t))$, $t \geq 0$, is an $\mathcal{F}^T$-martingale.*
2. *If w is concave, then for any switching strategy T, $w(X^T(t))$, $t \geq 0$, is an $\mathcal{F}^T$-supermartingale.*

Proof. The notation is different and the situation is slightly modified, but the idea behind the proof of this Theorem is the same as in the proof of Theorems 2.4 and 3.1 in [13] to which we refer the reader. □

A function w defined on D is called a *concave majorant* of f if it is concave and $w(x) \geq f(x)$, for all $x \in D$.

Theorem 2 *Let w be a concave majorant of f. Suppose that there exists a switching strategy T for which $w(X^T(t \wedge \tau))$, $t \geq 0$, is an $\mathcal{F}^T$-martingale where τ denotes the first hitting time of the set $\{x : w(x) = f(x)\}$ by the process X^T. If $\mathbf{E}_x \tau < \infty$ for all x, then w is the value function defined by (2.1), T is an optimal switching strategy, and τ is an optimal stopping time.*

Proof. Appealing to (2) of Theorem 1 and the optional sampling theorem, we conclude that

$$w(x) \geq \mathbf{E}_x w(X^T(\tau)) \geq \mathbf{E}_x f(X^T(\tau)) \tag{3.1}$$

for any switching strategy T and any stopping time τ. For the specific strategy T and stopping time τ described in the Theorem, we can apply the martingale optional sampling theorem to get that

$$w(x) = \mathbf{E}_x w(X^T(\tau)) = \mathbf{E}_x f(X^T(\tau)). \tag{3.2}$$

From (3.1) and (3.2), we conclude that w is the value function and that T and τ form an optimal strategy and stopping time, respectively. $\square$

4 The Concave Hull.

In this section, we shall show that the value function is precisely the concave hull of f. To this end, we shall need several definitions and theorems from convex analysis. We will give all definitions and, unless specified otherwise, parenthetical references to theorems shall refer to theorems in Rockafellar's classic text [10].

The *lower epigraph* (or just epigraph) of f is the set F defined as

$$F = \{(x, y) \in \mathbb{R}^{d+1} : x \in D, y \leq f(x)\}. \tag{4.1}$$

The upper semi-continuity of f implies that F is closed (Theorem 7.1). The boundedness of f implies that F is bounded and hence compact. Let W denote the *convex hull* of F defined as the intersection of all convex sets containing F. The closedness of F together with the commutativity of closure and convexification of bounded sets (Theorem 17.2) implies then that W is closed. Let w denote the real-valued function defined on D by

$$w(x) = \max\{y \in \mathbb{R} : (x, y) \in W\}, \quad x \in D. \tag{4.2}$$

The function w is called the *concave hull* of f. It is well-known (see, e.g. Section 2.5 of [3]) that w is the smallest concave majorant of f. The closedness of W implies that w is upper semi-continuous (Theorem 7.1). The concavity of w implies that it is continuous in the relative interior of D (Theorem 10.1).

A *face* of W is any convex subset V of W for which every closed line segment in W with relative interior contained in V has both endpoints in V as well. Since W is convex, it can be partitioned into the collection of relative interiors of its non-empty faces (Theorem 18.2). Let $\mathcal{B}$ denote this collection.

For each $x \in D$, let $B(x)$ denote that element of $\mathcal{B}$ that contains $(x, w(x))$. Clearly, for each x and $\tilde{x}$ in D, either $B(x) = B(\tilde{x})$ or $B(x) \cap B(\tilde{x}) = \emptyset$. Let $A(x)$ denote the projection of $B(x)$ onto D:

$$A(x) = \{\tilde{x} : (\tilde{x}, w(\tilde{x}) \in B(x)\}.$$

It is easy to see that the collection $\mathcal{A}$ of distinct elements of $\{A(x) : x \in D\}$ forms a partition of D into convex sets. We call it the *facial decomposition* of D associated with f.

We now have the tools necessary to construct a specific switching strategy that we shall prove is optimal. For each A in $\mathcal{A}$ with $\dim A \geq 1$, let $u^*(A)$ denote a direction vector from S^{d-1} "lying" in A. That is, if x is any point in A, $x + u^*(A)$ must belong to the affine hull of A. For each $x \in D$, let $A(x)$ denote the convex set in $\mathcal{A}$ that contains x and let

$$C = \{x \in D : \dim A(x) \geq 1\}.$$

Let T denote the switching strategy for which, until the first exit time from C, only $T_{u^*(A)}$ increases when $X^T(t) \in A$. After the first exit time from C, any fixed direction, call it $\bar{u}$, is chosen and is used from that time onward. Regarding the four defining properties of a switching strategy, it appears that only the second might be in doubt, so we focus on it.

Fix $x \in C$. Since $\dim A(x) \geq 1$, the direction $u^*(A(x))$ is chosen initially. Let τ_1 denote the first exit time from $A(x)$ obtained using this diffusion direction. At this time, the process lies on one of two relative boundary points of $A(x)$. Let ζ_1 denote this point. The set $A(\zeta_1)$ must have dimension strictly smaller than the dimension of $A(x)$. If it has dimension 0, then we switch to $\bar{u}$ and use this direction for all future time. If however, it has dimension greater than 0, then the direction $u^*(A(\zeta_1))$ is selected. Clearly, this process continues and after a finite number of steps, say N, the diffusion must arrive at a point ζ_N for which the dimension of $A(\zeta_N)$ is 0 after which only direction $\bar{u}$ is used. It is clear that the (random) number N is bounded by d and hence that no more than $2^d + 1$ directions are ever used.

In Theorem 4, we show that the strategy just constructed is an optimal strategy. For this theorem, the most crucial property of upper semi-continuous concave functions is the following result about their boundary behavior.

Theorem 3 *Let w be an upper semi-continuous concave function on D. For every x in the interior of D and every y on the boundary of D, the following limit holds:*

$$w(y) = \lim_{\lambda \to 0} w(\lambda x + (1-\lambda)y).$$

This theorem is exactly Theorem 7.5 in [10] and so we don't prove it here. While an upper semi-continuous concave function need not be continuous on all of D, it is continuous in the interior (as mentioned before) and this Theorem says that the continuity property also holds at boundary points as long as one approaches along a line segment from the interior. Stronger results are known, see e.g. Theorem 10.2 in [10], but shall not be needed here.

We are now ready for our main result:

Theorem 4 *The value function v is the concave hull of f. Let T denote the switching strategy defined above and let τ denote the first hitting time of the support set of v. Then T and τ together form an optimal strategy and stopping time.*

Proof. Let w denote the concave hull of f, let T denote the switching strategy defined above, and let τ be the first hitting time of $\{x : w(x) = f(x)\}$. It follows from part (1) of Theorem 1 and Theorem 3 that $w(X^T(t \wedge \tau))$ is an $\mathcal{F}^T$-martingale. All that remains then is to apply Theorem 2, which requires that $\mathbf{E}_x \tau < \infty$ for every starting point $x \in D$. To show this, we first note that the stop set $\{x \in D : w(x) = f(x)\}$ is not too small. Indeed, we shall show that its complement is contained in C. Hence, the first hitting time of the support set is bounded above by the first hitting time of a zero dimensional set A in the facial decomposition $\mathcal{A}$. But this latter hitting time is simply the sum of a finite number of exit times from finite intervals each of which clearly has finite expectation.

It only remains to show that $\{x \in D : w(x) > f(x)\} \subset C$. Fix an x for which $w(x) > f(x)$. Caratheodory's Theorem tells us that the concave hull w has the following representation:

$$w(x) = \sup\{\sum_{i=1}^{d+1} \lambda_i f(x_i) : \sum_{i=1}^{d+1} \lambda_i x_i = x; \lambda_i \geq 0, i = 1, 2, \ldots, d+1; \sum_{i=1}^{d+1} \lambda_i = 1\}$$

(see, e.g. Corollary 17.1.5 in [10]). The compactness of the set over which this optimization takes place implies that the supremum is attained. That is, there exist points $x_i \in D$, $i = 1, 2, \ldots, d+1$, such that $w(x)$ is a convex combination of the $f(x_i)$:

$$w(x) = \sum_{i=1}^{d+1} \lambda_i f(x_i).$$

Let k denote the number of λ_i's that are strictly positive. Without loss of generality, we may assume that $\lambda_i > 0$ for $i = 1, 2, \ldots, k$ and $\lambda_i = 0$ for $i = k+1, k+2, \ldots, d+1$. Since $w(x) > f(x)$, it follows that $k > 1$. Since w is concave and majorizes f, it follows that $f(x_i) = w(x_i)$ for $i = 1, 2, \ldots, k$ and so

$$w(x) = \sum_{i=1}^{k} \lambda_i w(x_i).$$

The fact that $k > 1$ now implies that $\dim(B(x)) > 0$, which in turn implies that $\dim(A(x)) > 0$. Hence $x \in C$. □

Acknowledgement: The author would like to thank for Michael Aizenman for suggesting this problem.

References

[1] R. Cairoli and R.C. Dalang. Optimal switching between two random walks. Technical report, Tufts University, 1993.

[2] E.B. Dynkin. The optimum choice of the instant for stopping a Markov process. *Dokl. Akad. Nauk SSSR*, 150:238–240, 1963.

[3] J.-B. Hiriart-Urruty and C. Lemarechal. *Convex Analysis and Minimization Algorithms I.* Springer-Verlag, 1993.

[4] A. Mandelbaum. Discrete multiarmed bandits and multiparameter processes. *Probability Theory and Related Fields*, 71:129–147, 1986.

[5] A. Mandelbaum. Continuous multi-armed bandits and multi-parameter processes. *Ann. Prob.*, 15:1527–1556, 1987.

[6] A. Mandelbaum. Navigating and stopping multi-parameter bandit processes. In W. Fleming and P.L. Lions, editors, *IMA Volume 10, Stochastic Differential Systems, Stochastic Control Theory and Applications.* Springer-Verlag, 1988.

[7] A. Mandelbaum, L.A. Shepp, and R.J. Vanderbei. Optimal switching between a pair of Brownian motions. *Ann. Prob.*, 18:1010–1033, 1990.

[8] A. Mandelbaum and R.J. Vanderbei. Optimal stopping and supermartingales over partially ordered sets. *Z. Warsch. verw. Gebiete*, 57:253–264, 1981.

[9] A. Mandelbaum and R.J. Vanderbei. Brownian bandits, 1994. To appear in special volume in honor of E.B. Dynkin's 70th birthday.

[10] R.T. Rockafellar. *Convex Analysis.* Princeton University Press, 1970.

[11] R.J. Vanderbei. Optimal switching among several brownian motions. *SIAM J. Control and Opt.*, 30:1150–1162, 1992.

[12] J.B. Walsh. Probability and Dirichlet problem for multiply superharmonic functions. *Ann. Inst. Fourier*, 18:221–279, 1968.

[13] J.B. Walsh. Optional increasing paths. In *Colloque ENST–CNET; Lecture Notes in Math 863*, pages 172–201. Springer–Verlag, 1981.

Robert J. Vanderbei, Dept. of Civil Engineering and Operations Research, Program in Statistics and Operations Research, Princeton University, Princeton NJ 08544, USA