CONTINUOUS-TIME GAMBLING PROBLEMS

by

David C. Heath[1] and William D. Sudderth[2]

Technical Report No. 214

1. Introduction.

Suppose a particle drifts according to a Brownian motion process starting from some point of an annulus in N-space. You can control the particle by adding a right-continuous function of the past whose total variation does not exceed some constant (representing your fuel). How should you control the particle to maximize the probability of hitting the inner boundary of the annulus before hitting the outer boundary? This problem, whose solution is given in Example 1 of Section 7, and other similar problems motivated us to consider continuous-time gambling problems.

An elaborate theory of discrete-time gambling (or stochastic control) was developed by Dubins and Savage in [2]. They presented their results in a finitely additive framework devoid of measurability assumptions and σ - fields. A continuous-time formulation is given here which uses the conventional countably additive machinery, but most of our notation, terminology, and theorems are adapted from [2].

Section 2 below gives a simple but fairly general formulation of the continuous-time gambler's problem and enough theory to handle the example mentioned above. Sections 3 and 4 are extensions of the theory which prove useful in other examples. Sections 5 and 6 are devoted to some of the more obvious theoretical questions.

2. Continuous-time gambling problems.

A discrete-time gambling problem, as formulated in [2], has three basic ingredients: a set F of possible fortunes or states, a real-valued function u on F which measures the utility of each fortune to the gambler, and a gambling house Γ which assigns the gambles available at each fortune in F. The continuous-time formulation used here has three quite similar ingredients.

Let F be a topological space and let $\mathfrak{F}$ be its σ - field of Borel sets. Let u be a bounded, $\mathfrak{F}$-measurable function from F to the real line. Think of F as the set of fortunes and u as the utility function for a gambler. (The assumption that u is bounded is convenient but not essential.)

In the discrete-time formulation of [2], the gambling house Γ specifies for any fortune x the set of possible distributions for the next fortune. Since there need not be a "next fortune" in continuous-time, there cannot be a complete analogy here. However, a gambling house in the sense of [2] also specifies for each x, at least implicitly, the set of possible distributions for the stochastic process corresponding to the gambler's sequence of future fortunes. This can be done as well in a continuous-time setting. Some additional machinery is needed for the precise definition.

Let Ω be the collection of all functions from $R^+ = [0, \infty)$ into F which are continuous from the right at each point in R^+ and have left limits at every strictly positive number. For $t \in R^+$ and $\omega \in \Omega$, let $X_t(\omega) = \omega(t)$. Take $\mathfrak{B}$ to be the smallest σ - field of subsets of Ω with respect to which every X_t, $t \geq 0$, is measurable from $(\Omega, \mathfrak{B})$ to $(F, \mathfrak{F})$, and, for $t \geq 0$, $\mathfrak{B}_t$ to be the σ - field generated by the collection $\{X_s, 0 \leq s \leq t\}$. Let $\mathfrak{B}_{t^+} = \bigcap_{s > t} \mathfrak{B}_s$. The function X_t is regarded as the gambler's fortune at time t.

Let Γ be a function defined on F which assigns to each $x \in F$ a non-empty set $\Gamma(x)$ of countably additive probability measures on (Ω, β). For $x \in F$, a member σ of $\Gamma(x)$ is said to be a strategy available at x in Γ. The function Γ is called a gambling house if the following are true:

(i) If $x \in F$ and $\sigma \in \Gamma(x)$, then $\sigma[X_0 = x] = 1$. (A gambler's fortune at time 0 must under every available strategy almost surely be his initial fortune.)

(ii) For this assumption, we need a bit of notation. For $\omega \in \Omega$ and $s, t \in R^+$, let $\omega_t(s) = \omega(s + t)$. Then, for $A \in \beta$, let $A_t = \{\omega : \omega_t \in A\}$. Now for every $x \in F$, $\sigma \in \Gamma(x)$, and $t \geq 0$, assume there is a map σ_t from Ω to the collection of probability measures on β such that $\sigma_t(\omega)(A)$ is a version of $\sigma(A_t | \beta_t)(\omega)$ for each $A \in \beta$. Thus σ_t is a regular conditional distribution under σ for the process $\{X_s, s \geq t\}$ of future fortunes given the past β_t. Finally, assume that, for $x \in F$, $\sigma \in \Gamma(x)$, and $t \geq 0$,

$$\sigma \{\omega : \sigma_t(\omega) \in \Gamma(X_t(\omega))\} = 1.$$

(Intuitively, (ii) says that, for any available strategy, a gambler's conditional strategy, given the past up to some time, must almost surely be a strategy available at the current fortune.)

By a stopping time τ is meant a Borel measurable map from (Ω, β) to $R^+ \cup \{\infty\}$ such that, for every $t \in R^+$, $[\tau \leq t] \in \beta_{t+}$. For $x \in F$, a policy available at x in Γ is a pair (σ, τ) where $\sigma \in \Gamma(x)$ and τ

is a stopping time such that $\sigma[\tau < \infty] = 1$.

A gambler with initial fortune $x \in F$ selects a policy (σ, τ) available at x in Γ and receives $\sigma u(X_\tau) = \int u(X_\tau) d\sigma$, the expected utility of his fortune at the time of stopping. For $x \in F$, let

$$U(x) = \sup \sigma u(X_\tau),$$

where the supremum is over all policies (σ, τ) available at x. If (σ, τ) is a policy available at x for which $\sigma u(X_\tau) = U(x)$, it is called optimal. The problems of major interest are to determine U and to find optimal or nearly optimal policies.

If $\Gamma(x)$ contains only one strategy, then a gambler with initial fortune x is faced with an optimal stopping problem of the type considered by Fakeev in [4]. Consequently several of our results are related to those in [4].

A result which has proved to be extremely useful in the solution of discrete-time gambling problems is Theorem 2.12.1 of [2]. Our first theorem is a continuous-time version of that result.

Let Q be a bounded measurable function from F to the real line. The function Q is said to be excessive for Γ if, given $x \in F$, $\sigma \in \Gamma(x)$, and $t \in R^+$, then $\sigma Q(X_t) \leq Q(x)$.

Theorem 1: Suppose Q is a bounded Borel measurable map from F to the real numbers. Assume Q is excessive for Γ, $Q \geq u$, and, that under every strategy σ available in Γ, the process $\{Q(X_t), t \geq 0\}$ has sample paths which are right-continuous σ-almost surely. Then $Q \geq U$.

Proof: Let $x \in F$, $\sigma \in \Gamma(x)$, and $0 \leq s \leq t$. Then, by assumption (ii) and the excessiveness of Q,

$$\sigma(Q(X_t) | \mathcal{B}_s)(\omega) = \sigma_s(\omega)(Q(X_{t-s})) \leq Q(X_s(\omega))$$

σ - almost surely.

Thus the adapted process $\{Q(X_t), \mathcal{B}_t, t \geq 0\}$ is a bounded, right-continuous supermartingale with respect to σ. So, given $x \in F$ and a policy (σ, τ) available at x in Γ, it follows from the optional sampling theorem (VI T13, p. 98 of [8]) and our hypothesis $Q \geq u$ that

$$\sigma u (X_\tau) \leq \sigma Q(X_\tau)$$

$$\leq \sigma Q(X_0)$$

$$= Q(x) .$$

Take the supremum over policies (σ, τ) available at x to complete the proof. $\square$

<u>Corollary 1:</u> Let Q be a bounded, continuous map from F to the reals. If Q is excessive for Γ and $Q \geq u$, then $Q \geq U$.

The next result gives local criteria for a function Q to be excessive. Its proof uses the following elementary lemma.

<u>Lemma 1:</u> Let $f : R^+ \rightarrow R^+$ and suppose f is right-continuous at every point of R^+ and has a left limit $f(t -)$ at each strictly positive t. Then necessary and sufficient conditions for f to be nonincreasing are

(i) $f(t -) \geq f(t)$ for all $t > 0$,

and

(ii) $\lim_{h \downarrow 0} \sup \dfrac{f(t + h) - f(t)}{h} \leq 0$

for all $t \geq 0$.

<u>Theorem 2:</u> Let Q be a bounded, Borel measurable function from F to the real numbers. Suppose that for every $x \in F$ and $\sigma \in \Gamma(x)$ the process

$\{Q(X_t), \ t \geq 0\}$ has right-continuous sample paths with left limits σ-almost surely and, for $t > 0$, $\lim_{s \uparrow t} \sigma \ Q(X_s) \geq \sigma \ Q(X_t)$. Assume also that

(1) there exist h_0 and M such that

$$\sup_{0 < h \leq h_0} \ \sup_{x \in F, \sigma \in \Gamma(x)} h^{-1}(\sigma \ Q(X_h) - Q(x)) \leq M \ ,$$

and

(2) for every $x \in F$ and $\sigma \in \Gamma(x)$,

$$\limsup_{h \downarrow 0} h^{-1}(\sigma \ Q(X_h) - Q(x)) \leq 0.$$

Then Q is excessive for Γ.

<u>Proof:</u> Let $x \in F$ and $\sigma \in \Gamma(x)$. Define $f(t) = \sigma \ Q(X_t)$ for $t \geq 0$. It suffices to show that f is nonincreasing.

By hypothesis, $f(t-) \geq f(t)$ for $t > 0$. It remains to check condition (ii) of the lemma.

For $t \geq 0$,

$$\limsup_{h \downarrow 0} h^{-1}\{f(t + h) - f(t)\}$$

$$= \limsup_{h \downarrow 0} \int h^{-1}\{\sigma_t(\omega) \ Q(X_h) - Q(X_t(\omega))\} \ \sigma(d\omega)$$

$$\leq \int \limsup_{h \downarrow 0} h^{-1}\{\sigma_t(\omega) \ Q(X_h) - Q(X_t(\omega))\} \ \sigma(d\omega)$$

$$\leq 0.$$

The first inequality is by (1) and Fatou's lemma. The second follows from (2).

The theorem now follows from Lemma 1. □

The assumptions of right-continuity and left limits for the $Q(X_t)$ processes in Theorem 2 are, of course, satisfied if Q is continuous.

The conditions (1) and (2) for particular Q's are closely related to those required for a function to belong to the domain of the infinitesimal operator of a Markov process (cf. Chapter II of [3]).

3. A continuous-time Fatou equation.

In this section, continuous-time analogues of results in [11] are presented.

Let $\{Y_t, \mathcal{F}_t, t \geq 0\}$ be an adapted process on a probability space $(S, \mathcal{F}, P)$. A stopping time τ is a map from S to $[0, +\infty]$ such that, for $t \in R^+$, $[\tau \leq t] \in \mathcal{F}_t$ and $P[\tau < +\infty] = 1$. The stopping times form a directed set under the natural partial ordering $\tau \leq \tau'$ if $\tau(s) \leq \tau'(s)$ for all $s \in S$. Assume there is an integrable random variable Z such that $Y_t \leq Z$ for all $t \in R^+$. Then $\{EY_\tau\}$ is a net. Here we consider the limit of this net, written $\lim_{\tau \to \infty} EY_\tau$, and its lim sup which is defined by

$$\limsup_{\tau \to \infty} E Y_\tau = \inf_{\tau'} \sup_{\tau \geq \tau'} E Y_\tau .$$

(Parallel definitions and results for the 'lim inf' are left to the reader.)

Let $Y^* = \limsup_{t \to \infty} Y_t$. Here and below the symbols 'τ' and 'τ'' vary over the set of stopping times and the symbol 't' over R^+.

Theorem 3: $\displaystyle\limsup_{\tau \to \infty} E Y_\tau \leq E Y^*$.

Proof: Easy and essentially the same as the proof of Theorem 2 in [11]. □

Theorem 4: Assume that the process $\{Y_t, t \geq 0\}$ has right-continuous sample paths, that the σ-fields $\{\mathcal{F}_t\}$ are right-continuous, and that Y^* is integrable. Then

$$\limsup_{\tau \to \infty} E Y_\tau = E Y^*.$$

Proof: Theorem 3 gives one inequality. The other is proved as in Theorem 1 of [11] except that the appropriate continuous-time results must be quoted

(cf. the remark at the top of p. 97 and 44a, p. 67 in [8]). $\square$

Theorem 5: Suppose there is an integrable random variable Z such that $|Y_t| \leq Z$ for all $t \geq 0$. Then $\limsup_{\tau \to \infty} E Y_\tau = \inf_t \sup_{\tau \geq t} E Y_\tau$. If $Y_t \to Y^*$ a.s., then $\lim_{\tau \to \infty} E Y_\tau = E Y^*$.

Proof: The first assertion is easy to check. For the second, apply Theorem 3 to get

$$E Y^* = E \left(\liminf_{t \to \infty} Y_t \right) \leq \liminf_{\tau \to \infty} E Y_\tau$$

$$\leq \limsup_{\tau \to \infty} E Y_\tau \leq E \left(\limsup_{t \to \infty} Y_t \right) = E Y^*. \quad \square$$

4. The utility of a strategy.

Let F, u, and Γ be as in Section 2. Suppose the gambler is not allowed to stop at will, but must continue play indefinitely. If his strategy is σ, his payoff is then defined, as in Chapter 3 of [2], to be

$$u(\sigma) = \limsup_{\tau \to \infty} \int u(X_\tau) \, d\sigma .$$

Here the lim sup is over the directed set of stop rules τ such that $\sigma[\tau < \infty] = 1$.

For $\omega \in \Omega$, let $u^*(\omega) = \limsup_{t \to \infty} u(\omega(t))$. The next result is a corollary to Theorem's 4 and 5 of the previous section.

Theorem 6: If u is continuous or if $u(X_t)$ converges σ-almost surely as $t \to \infty$, then

$$u(\sigma) = \int u^* d\sigma .$$

Thus, in many interesting cases, the payoff $u(\sigma)$ is just the expected value of the gambler's limiting utility.

For $x \in F$, let $V(x) = \sup u(\sigma)$, where the supremum is over all strategies available at x in Γ. A strategy σ available at x is called optimal if $u(\sigma) = V(x)$. Again the problems of interest are to compute V and determine optimal or nearly optimal strategies. The next theorem and its proof are analogous to Theorem 1.

Theorem 7: Let Q be a bounded, Borel measurable map from F to the real numbers. Assume Q is excessive for the gambling house Γ and, for every strategy σ available in Γ, $Q(\sigma) \geq u(\sigma)$ and the process $\{Q(X_t), t \geq 0\}$ has right-continuous sample paths σ - almost surely. Then $Q \geq V$.

5. When does U = V?

In this section we identify a class of gambling problems for which $U = V$. Notice that $V \leq U$ is always true. For the opposite inequality two additional assumptions are needed.

For $x \in F$, let ϵ_x be the probability on $\mathfrak{B}$ such that $\epsilon_x[X_t = x] = 1$ for all $t \geq 0$. A gambling house Γ is called leavable if $\epsilon_x \in \Gamma(x)$ for all $x \in F$. Thus, in a leavable house, a gambler can effectively stop immediately whatever his initial fortune.

The next assumption is that a gambler is permitted to modify his conditional strategy in a measurable way at any time so long as his new strategy is almost surely available at the current fortune. The precise formulation of this is called assumption A for future reference.

(A) Let $x \in F$ and let (σ, τ) be a policy available at x in Γ. Suppose σ_y is a probability on $\mathfrak{B}$ for each $y \in F$ and that $y \to \sigma_y(A)$ is Borel measurable for each $A \in \mathfrak{B}$. Suppose also that

$$\sigma \{\omega : \sigma_{X_\tau(\omega)} \in \Gamma(X_\tau(\omega))\} = 1.$$

Let μ be the probability on $\mathfrak{B}$ defined by the requirements that μ agrees with σ on $\mathfrak{B}_\tau$ and $\mu_\tau(\omega) = \sigma_{X_\tau(\omega)}$ μ - almost surely. Then $\mu \in \Gamma(x)$.

__Theorem 8:__ For leavable gambling houses which satisfy assumption A, $U = V$.

__Proof:__ Let (σ, τ) be a policy available at x in Γ. Let $\sigma_y = \epsilon_y$ for $y \in F$ and let μ be defined as in (A). Then $V(x) \geq u(\mu) = \int u(X_\tau) \, d\sigma$. Take the supremum over (σ, τ) to get $V(x) \geq U(x)$. $\square$

6. Measurable gambling houses.

Additional regularity conditions are imposed in this section, and U and V are then shown to be universally measurable (i.e., measurable with respect to the completion of $\mathcal{F}$ under any probability measure) and excessive.

Assume F is a complete, separable metric space. For I = [a, b] an interval of real numbers, let D(I) be the set of functions from I to F which are right-continuous on [a, b) and have left-limits on (a, b]. Give D(I) the Skorohod topology (cf. [5]). Then D(I) has the topological structure of a complete, separable metric space, and the σ-field of Borel subsets of D(I) is just the σ-field generated by the coordinate maps. (Proofs of these facts for the case when F = R are in [1]. The new difficulties introduced in the general case are minor.) The set Ω can be identified with D[0, ∞). Give it the topology introduced in [6]. Then Ω is complete, separable metric and $\mathcal{B}$ is its σ-field of Borel sets (Theorems 1 and 2 in [6]. Again only the case F = R is explicitly treated in this reference, but the general case is no harder.)

Let P be the collection of probability measures on $\mathcal{B}$. Give P the usual topology of weak convergence (cf. Chapter II of [9]). Then P is complete, separable metric. Denote the σ-field of Borel subsets of P by $\mathcal{P}$.

A gambling house Γ is called measurable if $\{(x, \sigma) : \sigma \in \Gamma(x)\}$ is a Borel subset of F x P. (Measurable gambling houses in the discrete-time case were introduced by Strauch in [10]).

The major results of this section are the following two theorems.

Theorem 9: If Γ is a measurable gambling house, then U and V are universally measurable.

Theorem 10: If Γ is a measurable gambling house which satisfies assumption A, then U and V are excessive.

Theorem 10 should be compared to Theorems 1 and 7. Some preliminary results are needed for the proofs.

Let T be the collection of stopping times τ which are everywhere finite on Ω. For $\sigma \in P$, define

$$\varphi(\sigma) = \sup_{\tau \, \in \, T} \sigma \, u \, (X_\tau).$$

Clearly, $U(x) = \sup_{\sigma \, \in \, \Gamma(x)} \varphi(\sigma).$

Lemma 2: The maps $\sigma \to \varphi(\sigma)$ and $\sigma \to u(\sigma)$ are Borel measurable from P to the real line.

Proof: For $n = 1, 2,\ldots,$ let T_n be the collection of maps τ from Ω to $\{k \, 2^{-n} : k = 0, 1,\ldots\}$ such that $[\tau \le k \, 2^{-n}] \in \mathcal{B}_{k2}{-n}$ for all k. Set $\varphi_n(\sigma) = \sup_{\tau \, \in \, T_n} \sigma \, u \, (X_\tau).$ Clearly, $\varphi_n(\sigma) \le \varphi(\sigma)$ for all n. Moreover, $\varphi_n(\sigma) \uparrow \varphi(\sigma)$ as $n \to \infty$. To see this, fix $\sigma \in P$ and $\tau \in T$. For each n, let τ_n be the smallest $k \, 2^{-n}$ greater than τ. Then $\tau_n \in T_n$, $\tau_n \downarrow \tau$ as $n \to \infty$, and, hence, $\sigma \, u \, (X_{\tau_n}) \to \sigma \, u \, (X_\tau).$

By the corollary to Theorem 2 in [10], each of the maps $\sigma \to \varphi_n(\sigma)$ is Borel measurable from P to the real line and, consequently, so is the map $\sigma \to \varphi(\sigma).$

For $t \ge 0$ and $\sigma \in P$, let $\Psi_t(\sigma) = \sup_{\tau \, \ge \, t} \sigma \, u(X_\tau).$ By an argument

similar to the above, the function $\sigma \rightarrow \Psi_t(\sigma)$ is Borel. By the first assertion in Theorem 5, $\Psi_t(\sigma) \downarrow u(\sigma)$ as $t \rightarrow \infty$. Hence, $\sigma \rightarrow u(\sigma)$ is Borel. $\square$

Lemma 3: Let $\sigma \in P$ and $t \geq 0$. then

$\qquad$ (i) $\quad \varphi(\sigma) \geq \int \varphi(\sigma_t(\omega)) \, \sigma(d\omega),$

and

$\qquad$ (ii) $\quad u(\sigma) \geq \int u(\sigma_t(\omega)) \, \sigma(d\omega).$

Proof: Let φ_n be as in the proof of Lemma 2. Since $\varphi_n \uparrow \varphi$, it suffices for (i) to show

$$\varphi(\sigma) \geq \int \varphi_n(\sigma_t(\omega)) \, \sigma(d\omega)$$

for each n.

$\qquad$ Fix n and $\epsilon > 0$. By Theorem 2 of [10], there is a countable set $C_n \subseteq T_n$ such that

$$\varphi_n(\sigma') = \sup_{\tau \in C_n} \sigma' \, u(X_\tau)$$

for all $\sigma' \in P$. Suppose $C_n = \{\tau_1, \tau_2, \ldots\}$. Let $A_k = \{\omega : k \text{ is the least integer} \ni \sigma_t(\omega) \, u(X_{\tau_k}) \geq \varphi_n(\sigma_t(\omega)) - \epsilon\}$. Define $\tau(\omega) = t + \tau_k(\omega_t)$ for $\omega \in A_k$. Then τ is a stopping time and

$$\varphi(\sigma) \geq \sigma \, u(X_\tau)$$

$$= \sum_k \int_{A_k} \sigma_t(\omega) \, u(X_{\tau_k}) \, \sigma(d\omega)$$

$$\geq \int \varphi_n(\sigma_t(\omega)) \, \sigma(d\omega) - \epsilon \, .$$

This completes the proof of (i).

To prove (ii), recall that, for every σ, $\Psi_t(\sigma) = \sup_{\tau \geq t} \sigma \, u(X_\tau) \to u(\sigma)$

as $t \to \infty$. So it suffices to prove (ii) with 'Ψ_t' in place of 'u'. Let $T_{n,t}$ be the collection of functions $\tau : \Omega \to \{t + k2^{-n} : k = 0, 1, 2, \ldots\}$ such that $[\tau \leq s] \in \beta_s$ for all $s \geq 0$. Set $\Psi_{n,t}(\sigma) = \sup_{\tau \in T_{n,t}} \sigma \, u(X_\tau)$ for each σ. Then $\Psi_{n,t}(\sigma) \to \Psi_t(\sigma)$ as $n \to \infty$ for each σ (see the proof that $\varphi_n \to \varphi$ in Lemma 2). So it suffices to show

$$\Psi_t(\sigma) \geq \int \Psi_{n,t}(\sigma_t(\omega)) \, \sigma(d\omega)$$

for all n and t. This can be shown by imitating the proof of (i). $\square$

Now we are ready to prove the theorems.

Proof of Theorem 9: Let r be a real number. The set $\{x \in F : U(x) > r\}$ is the projection on F of the set $\{(x,\sigma) \in F \times P : \sigma \in \Gamma(x), \varphi(\sigma) > r\}$. The latter set is Borel by Lemma 1 and the hypothesis that Γ is a measurable gambling house. Hence, its projection is analytic and, in particular, universally measurable ([8], T19, p. 39).

The proof that V is universally measurable is similar. $\square$

Proof of Theorem 10: Let $x \in F$, $\sigma \in \Gamma(x)$, $t > 0$, and $\epsilon > 0$. Let p_t be the distribution of X_t under σ. Since U is measurable for the completion of $\mathcal{F}$ under p_t, there is a Borel measurable function Q from F to the reals such that $p_t[Q = U] = 1$. Also, Q can be chosen so that $Q \leq U$.

Let $A = \{(y, \mu) \in F \times P : \mu \in \Gamma(y), \varphi(\mu) > Q(y) - \epsilon\}$. Then A is a Borel subset of $F \times P$ and each y-section A_y is nonempty. By Theorem 6.3

- 15 -

of [7], there is a Borel map $\eta : F \to P$ such that $p_t\{y : (y, \eta(y)) \in A\} = 1$.

Define a strategy μ by requiring that μ agree on β_t with σ and $\mu_t(\omega) = \eta(X_t(\omega))$ σ - almost surely. By assumption A, $\mu \in \Gamma(x)$. Hence,

$$U(x) \geq \varphi(\mu)$$

$$\geq \int \varphi(\mu_t(\omega)) \; \mu(d\omega) \qquad \text{(by Lemma 2)}$$

$$= \int \varphi(\eta(X_t(\omega))) \; \sigma(d\omega)$$

$$\geq \int U(X_t) \; d\sigma - \epsilon \; .$$

The proof that U is excessive is now complete. The proof for V is similar. $\square$

Notice that for Theorem 10, the full strength of assumption A is not used since only the constant stopping time t is used in the proof.

7. Examples.

Our first example is the situation mentioned in the introduction.

Example 1. Let $N \geq 3$ (similar results hold for $N = 1$ and 2, but the formulae are different), set $A = \{x \in R^N : 1 \leq |x| \leq R\}$ for some fixed $R > 1$, and $F = A \times [0, \infty)$. If $(x, b) \in F$, we refer to x as the position and b as the amount of fuel remaining. For each $(x_0, b_0) \in F$ we define $\Gamma(x_0, b_0)$ to be the collection of all measures on Ω which can be obtained as follows: Let $(\Omega', X_t', \mathcal{B}_t', P_x', x \in R^N)$ be standard Brownian motion on R^N and let $f: [0, \infty) \times \Omega' \to R^N$ be a right continuous non-anticipative function with $f(0, \omega') = 0$ and total variation (for each fixed ω') not exceeding b_0. Let $b(t, \omega')$ be the total variation of $f(s, \omega')$ on $[0, t]$. ($b(t, \omega')$ denotes the fuel used by the control f up to time t on the path ω'.) Let $\tau(\omega') = \inf \{t \geq 0 : X_t'(\omega') + f(t, \omega') \notin \text{int}(A)\}$. We require that $X_\tau'(\omega') + f(\tau, \omega') \in A$ P_{x_0}' - a.s. Then the map $T : \Omega' \to \Omega$ defined by

$$(T(\omega'))(t) = (X_{t \wedge \tau}'(\omega') + f(t \wedge \tau, \omega'), b_0 - b(t \wedge \tau, \omega'))$$

carries P_{x_0}' to a measure on Ω; this measure is an element of $\Gamma(x_0, b_0)$.

$$\text{Let} \quad u(x, b) = \begin{cases} 1 \text{ if } |x| = 1 \\ \\ 0 \text{ otherwise.} \end{cases}$$

We now describe a sequence σ_n of (ϵ-optimal) strategies starting at $(x_0, b_0) \in F$. We assume $1 + b_0 < R$ and $b_0 > 0$; the other cases are trivial. We also assume $|x_0| \neq 1 + b_0$; this case is handled in almost

exactly the same fashion as those handled below, but the technical assumption that $f(0, \omega') = 0$ produces a slight, but easily overcome, difficulty. Moreover, we give only a description of the strategy; the reader can (tediously) fill in the rigorous details.

To construct the nth strategy, we use recursion. Let the 0th strategy be to use no control at all; i.e., use $f(t, \omega') \equiv 0$. For the nth strategy, drift until you're about to hit the set $\{(x, b) : |x| = 1 + b_0\}$ or $\{(x, b) : |x| = R\}$. If you're about to hit the first set (i.e., at the hitting time for that set) use all of the fuel to jump distance b_0 toward the origin, thus hitting the inner surface of the annulus, which causes you to stop. If you're about to hit the second set, use $1/n$ of the fuel to jump a distance b_0/n toward the origin, thus (temporarily, at least) avoiding the outer surface of the annulus. From this new place, continue, using the $(n - 1)$th strategy. (This results in using the same amount of fuel, $\frac{1}{n-1} \cdot \frac{n-1}{n} \cdot b_0 = b_0/n$, each time you hit the outer surface of the annulus.)

Let $Q_n(x_0, b_0)$ be the return from this strategy. Since the only way to "lose" is to drift to the outer surface before drifting to radius $(1 + b_0)$, then drift from radius $R - (b_0/n)$ to radius R before radius $b_0(1 - (j/n))$ for $j = 1, 2, \ldots, n,$ it is easy to see that

$$
Q_n(x_0, b_0) = \begin{cases} 1 & \text{if} \quad |x_0| < 1 + b_0 \\[2mm] 1 - \zeta(1+b_0, |x_0|, R) \prod_{j=1}^{n} \zeta(1+b_0(1 - \tfrac{j}{n}), R - \tfrac{b_0}{n}, R) \\[2mm] & \qquad\qquad \text{if} \quad 1 + b_0 < |x_0| < R \\[2mm] 0 & \text{if} \quad |x_0| = R \end{cases}
$$

where $\zeta(a, b, c) = (a^{2-N} - b^{2-N})/(a^{2-N} - c^{2-N})$; if $|x_0| = 1 + b_0$ we can choose a strategy with payoff $1 - \frac{1}{n}$.

Now set $Q(x_0, b_0) = \lim\limits_{n \to \infty} Q_n(x_0, b_0)$. Then Q is given by

$$Q(x_0, b_0) = \begin{cases} 1 & \text{if } |x_0| \le 1 + b_0 \\[2ex] 1 - \zeta(1 + b_0, |x_0|, R) \exp\left[-\frac{N-2}{R} \int_0^{b_0} \frac{R^{2-N} \, d\beta}{(1+\beta)^{2-N} - R^{2-N}} \right] \\[1ex] \qquad\qquad\qquad\qquad\qquad \text{if } 1 + b_0 < |x_0| < R \\[2ex] 0 & \text{if } |x_0| = R. \end{cases}$$

Clearly $Q \le U$. We now apply Theorem 1 to show that equality holds. Obviously $Q \ge u$. Since Q is continuous at all points (x, b) except if $|x| = R$ (and once paths reach $\{(x, b) : |x| = R\}$ they stay there), and paths of the controlled processes are right continuous, we have that $Q(X_t(\omega))$ is right continuous for σ - a.e. ω.

To show that Q is excessive we use Theorem 2. By arguments similar to those above, $Q(X_t(\omega))$ has left limits. We must show that

$$\lim\limits_{s \uparrow t} \sigma Q(X_s) \ge \sigma Q(X_t) .$$

This will follow from the dominated convergence theorem if we show

$$\lim\limits_{s \uparrow t} Q(X_s) \ge Q(X_t) \qquad \sigma - \text{a.s.}$$

We consider two cases: suppose $X_t = (x, b)$ with $|x| = R$. Then $Q(X_t) = 0$

so the inequality holds.

If $X_t = (x, b)$ with $|x| < R$ and $\lim Q(X_s) \neq Q(X_t)$, then X_s is not left-continuous at $s = t$ (since Q is continuous at X_t.) Modify Q on $\{(x, b) : |x| = R\}$ in the obvious way (call the modified version $\widetilde{Q}$) to make it continuous everywhere on F. We must then show

$$\widetilde{Q}(X_{t-}) \geq \widetilde{Q}(X_t).$$

Let $X_{t-} = (x_0, b_0)$ and $X_t = (x_1, b_1)$. Clearly we have $|x_0| - |x_1| \geq b_0 - b_1$. Now $\widetilde{Q}$ depends on (x, b) only through $|x|$ and b; set $q(|x|, b) = \widetilde{Q}(x, b)$. Differentiation and some algebra show that

$$\frac{\partial q(\rho, b)}{\partial \rho} + \frac{\partial q(\rho, b)}{\partial b} \geq 0 \qquad 1 + b \leq \rho \leq R$$

(the derivatives are one-sided at the boundaries; strict inequality holds in (the interior.) Thus from the mean value theorem we obtain

$$q(|x_0|, b_0) \geq q(|x_1|, b_0 - (|x_0| - |x_1|))$$

From the monotonicity of q in b we have

$$q(|x_1|, b_0 - (|x_0| - |x_1|)) \geq q(|x_1|, b_0 - (b_0 - b_1)) = q(|x_1|, b_1)$$

which completes the argument.

It remains to check conditions (1) and (2) of Theorem 2. First we extend the function q defined above as follows:

$$q(\rho, b) = \begin{cases} 1 & \text{if } \rho \leq 1 + b \\ 1 - \zeta(1+b, \rho, R) \exp\left[- \frac{N-2}{R} \int_0^b \frac{R^{N-2} d\beta}{(1+\beta)^{N-2} - R^{N-2}}\right] & \text{if } 1+b < \rho \leq R+b. \end{cases}$$

We must use more explicitly the form of strategies: for any $\sigma \in \Gamma(x_0, b_0)$ we let

$$X_t(\omega) = (X'_{t \wedge \tau}(\omega') + f(t \wedge \tau, \omega'), b_0 - b(t \wedge \tau, \omega')), \quad \text{where } \omega = T(\omega').$$

We now check (1) with $h_0 = 1$: Define τ' by

$$\tau'(\omega') = \inf \{t : |X'_t(\omega')| = 1 + b_0\} \wedge \tau(\omega').$$

Clearly $1 + b_0 \leq |X'_{t \wedge \tau'}(\omega')| \leq R + b_0$ for all $t \geq 0$. Since for each fixed b, $q(|x|, b)$ is bounded and harmonic in x we have

$$E_{x_0}(Q(X'_{h \wedge \tau'}(\omega'), b_0)) = Q(x_0, b_0).$$

We now show that

$$(*) \qquad Q(X'_{h \wedge \tau}(\omega') + f(h \wedge \tau, \omega'), b_0 - b(h \wedge \tau, \omega')) \leq$$

$$Q(X'_{h \wedge \tau}(\omega'), b_0) + M((|X'_{h \wedge \tau}(\omega')| - R) \vee 0)^2.$$

This suffices since, taking expectations, we get

$$\sigma Q(X_h) - Q(x_0, b_0) \leq M E((|X'_{h \wedge \tau'}(\omega')| - R) \vee 0)^2$$

$$\leq M E((X'_{h \wedge \tau'}(\omega'))^2) \leq M h.$$

To check $(*)$, note that $h \wedge \tau'(\omega') < h \wedge \tau(\omega') \Rightarrow |X'_{h \wedge \tau'}(\omega')| = 1 + b_0$ so the right

- 21 -

hand side exceeds 1; hence the inequality holds. Thus we need only consider ω'

for which $h \wedge \tau'(\omega') = h \wedge \tau(\omega')$. The following lemma is easily checked:

__Lemma 4:__ For fixed $y_0 \in R^N$ and $c_0 \in R^+$ with $1 + c_0 \le |y_0| \le R + c_0$,

the $z \in R^N$ satisfying $|z| \le c_0$ and $|y_0 + z| \le R + c_0$ which maximizes

$q(|y_0 + z|, b - |z|)$ is given by

$$z = \begin{cases} 0 & \text{if} \quad |y_0| \le R \\ \\ (R - |y_0|) \, y_0/|y_0| & \text{if} \quad |y_0| > R. \end{cases}$$

The maximum value is then

$$q(|y_0|, b_0) \qquad \text{if} \quad |y_0| \le R$$

$$q(R, b_0 - (|y_0| - R)) \quad \text{if} \quad |y_0| \ge R.$$

Applying this lemma (setting $y_0 = X'_{h \wedge \tau}(\omega')$ and $c_0 = b_0$), and using the

monotonicity of q in b on the set where $Q > 0$, we obtain

$$Q(X'_{h \wedge \tau}(\omega') + f(h \wedge \tau, \omega'), b_0 - b(h \wedge \tau, \omega')) \le$$

$$\begin{cases} Q(X'_{h \wedge \tau'}(\omega'), b_0) & \text{if} \quad |X'_{h \wedge \tau}(\omega')| \le R \\ \\ q(R, b_0 - (|X'_{h \wedge \tau}(\omega')| - R)) & \text{otherwise.} \end{cases}$$

Elementary calculus (using $\dfrac{\partial q}{\partial \rho} + \dfrac{\partial q}{\partial b} = 0$ when $\rho = R$) shows that for

$0 < \epsilon < b_0$,

$$q(R, b_0 - \epsilon) \le q(R + \epsilon, b_0) + M \epsilon^2.$$

This completes the proof of the next-to-last condition of Theorem 2. The same argument used above, together with the observation that if $1 + b_0 < |x_0| < R$, $P'_{x_0}\{|X'_s(\omega')| > R \text{ for some } s \in [0, h]\} = 0(h)$, verifies the last condition. Thus Q is excessive, so $Q \geq U$.

It is tedious but not difficult to prove that there is no optimal control available at (x_0, b_0) in Γ if $b_0 > 0$ and $1 + b_0 < |x_0| < R$: let $f(s) = f(s, \cdot)$ be the non anticipative function associated with a control, and let τ be the time at which the uncontrolled process (the Brownian motion) hits $\{x: |x| = 1 + b \text{ or } |x| = R\}$. Arguments similar to those above show that $|f(\tau-)| > 0$ implies that the control is sub-optimal; however, $|f(\tau)|$ must be positive for an optimal control. Thus f has a jump at τ. One can construct (using the ϵ-optimal strategies obtained above) a control f^* better than f with $f^*(s) = 0$ for $0 \leq s < \tau$ and $f^*(\tau) = \frac{1}{2}f(\tau)$. Thus f is sub-optimal.

Example 2. This example is in many ways similar to the preceeding example, but represents a non leavable situation. Let $N \geq 3$, and $A = \{x \in R^N : |x| \geq 1\}$. Let $F = A \times [0, \infty)$, and for $(x, b) \in F$, let $u(x, b) = 0$ if $|x| = 1$ and $= 1$ otherwise. The gambler is required to play indefinitely; his goal is to avoid the unit ball. The formal description of Γ is the same as that in example 1 with the A there replaced appropriately by the A defined above. Again a sequence of strategies can be given: starting at (x_0, b_0), divide the available fuel into n pieces, each of size b_0. Each time you are about to hit the unit ball, use one piece of fuel to move to a new position with distance $1 + b_0/n$ from the origin.

The utility of this strategy, computed according to Theorem 6 , is

$$Q_n(x_0, b_0) = \begin{cases} 0 & \text{if} \quad |x_0| = 1 \\ \\ 1 - |x_0|^{2-N}(1 + b/n)^{2-N} & \text{if} \quad |x_0| > 1 \end{cases}$$

Let $Q = \lim\limits_{n \to \infty} Q_n$; clearly $Q(x_0, b) = 1 - |x_0|^{2-N} e^{(2-N)b}$. Again using Theorem 2, we can show that Q is excessive. Using Theorem 6 again together with the fact that the radius of the position either has limit 1 or limit ∞, we see that $Q(\sigma) \geq u(\sigma)$ for all $\sigma \in \Gamma$. Hence $Q \geq V$; since obviously $Q \leq V$, we obtain equality.

References

[1] Billingsley, P. (1968). Convergence of Probability Measures. John Wiley, New York.

[2] Dubins, L. E. and Savage, L. J. (1965). How to Gamble If You Must. McGraw - Hill, New York.

[3] Dynkin, E. B. (1965). Markov Processes. Springer - Verlag, Berlin.

[4] Fakeev, A. G. (1970). "Optimal stopping rules for stochastic processes with continuous parameter". Theory Prob. Applications. 15, 324-331.

[5] Kolmogorov, A. N. (1956). "On Skorohod convergence", Theory of Prob. and Appl. 1, 215-222.

[6] Lindvall, F. (1973). "Weak convergence of probability measures and random functions in the function space $D[0, \infty)$. J. App. Prob. 10, 109-121.

[7] Mackey, G. W. (1957). "Borel structures in groups and their duals." Trans. A. M. S. 85, 134-165.

[8] Meyer, P. A. (1966). Probability and Potentials. Blaisdell Publishing, Waltham, Mass.

[9] Parthasarathy, K. R. (1967). Probability Measures on Metric Spaces. Academic Press, New York.

[10] Strauch, R. E. (1967). "Measurable gambling houses". Trans. A. M. S. 126, 64-72.

[11] Sudderth, W. D. (1971). "A 'Fatou equation' for randomly stopped variables", Ann. Math. Statist. 42, 2143-2146.