

A CONTINUITY QUESTION OF DUBINS AND SAVAGE

R. LARAKI,* *Université Paris-Dauphine*

W. SUDDERTH,** *University of Minnesota*

Abstract

Lester Dubins and Leonard Savage posed the question as to what extent the optimal reward function U of a leavable gambling problem varies continuously in the gambling house Γ , which specifies the stochastic processes available to a player, and the utility function u , which determines the payoff for each process. Here a distance is defined for measurable houses with a Borel state space and a bounded Borel measurable utility. A trivial example shows that the mapping $\Gamma \mapsto U$ is not always continuous for fixed u . However, it is lower semicontinuous in the sense that, if Γ_n converges to Γ , then $\liminf U_n \geq U$. The mapping $u \mapsto U$ is continuous in the supnorm topology for fixed Γ , but is not always continuous in the topology of uniform convergence on compact sets.

Dubins and Savage observed that a failure of continuity occurs when a sequence of superfair casinos converges to a fair casino, and queried whether this is the only source of discontinuity for the special gambling problems called casinos. For the distance used here, an example shows that there can be discontinuity even when all the casinos are subfair.

Keywords: Gambling theory; Markov decision theory convergence of value functions

2010 Mathematics Subject Classification: Primary 60G40, 90C40, 93E20

1. Introduction

A basic question concerning any mathematical problem is how the solution depends on the conditions. For a stochastic control problem, it is thus natural to ask how the optimal reward varies as a function of the stochastic processes available to the controller and of the reward structure. In the Dubins–Savage formulation [2], the processes available are determined by a gambling house Γ which specifies for each state x the set $\Gamma(x)$ of possible distributions for the next state. The worth of each state x to a player is the value $u(x)$ of the utility function at the state. In a leavable gambling problem a player chooses, in addition to one of the processes determined by Γ , a time to stop the game, and receives in reward the expected utility at the time of stopping. The optimal reward $U(x)$ is the supremum of the possible rewards starting from state x . (Precise definitions are in the next section.)

Dubins and Savage [2, p. 76] suggested that a notion of convergence be defined for gambling houses in order to study the extent to which U varies continuously in Γ and u . For the notion of convergence introduced in Section 3, a trivial example in Section 4 shows that the

Received 12 April 2016; revision received 13 October 2016.

* Postal address: Director of Research at CNRS, Université Paris-Dauphine, PSL Research University, Lamsade, 75016 Paris, France.

Supported by grants administered by the French National Research Agency as part of the Investissements d’Avenir program (Idex Grant agreement number ANR-11-IDEX-0003-02/Labex ECODEC No. ANR11-LABEX-0047 and ANR-14-CE24-0007-01 CoCoRiCo-CoDec).

** Postal address: School of Statistics, University of Minnesota, Minneapolis, MN 55455, USA.

Email address: bill@stat.umn.edu

mapping $\Gamma \mapsto U$ is not continuous in general. However, by Theorem 1, it is lower semicontinuous in the sense that, for Γ_n converging to Γ , $\liminf U_n \geq U$. Also, by Theorem 2, the mapping is continuous from below in the sense that, when the Γ_n increase to Γ then $\lim U_n = U$. By Corollary 1 in Section 5, the mapping $u \mapsto U$ is continuous in the supnorm topology. A simple example shows that the mapping is not always continuous for the topology of uniform convergence on compact subsets of X . Nonleavable gambling problems are discussed briefly in Section 6, where examples are given to show that the analogues to Theorems 1 and 2 do not hold for these problems. However, the analogue to Corollary 1 remains true.

The interesting special class of gambling problems called casinos are introduced in Section 7. Dubins and Savage [2, p. 76] observed that a discontinuity occurs when a sequence of superfair casinos converges to a fair casino (cf. Example 6 in Section 8). They surmised that this might be the only source of discontinuity for casinos with a fixed goal. For the definition of convergence used here, Example 8 shows that a discontinuity can occur even when all the casinos are subfair. However, Dubins and Meilijson [1] proved a continuity theorem for subfair casinos using a quite different notion of distance. A brief discussion of their work is in Section 9. The final section suggests the possibility of analogous results for continuous-time stochastic control problems.

There is related work available for control problems formulated as Markov decision processes including some very general results for finite-horizon and discounted models given by Langen [7]. There is little overlap with the main results here, which concern infinite horizon problems with no discounting.

The next section presents the necessary definitions and some general background material on the Dubins–Savage theory.

2. Preliminaries

A Dubins–Savage gambling problem comprises a *state space* or *fortune space* X , a *gambling house* Γ , and a *utility function* u . The gambling problems of this paper are assumed to be *measurable* in the sense of Strauch [12]. This means that X is assumed to be a nonempty Borel subset of a complete separable metric space. So, in particular, X is a separable metric space. The gambling house Γ is a function that assigns to each $x \in X$ a nonempty set $\Gamma(x)$ of probability measures defined on the Borel subsets $\mathcal{B}(X)$ of X . Let $\mathcal{P}(X)$ be the set of all probability measures defined on $\mathcal{B}(X)$ and give $\mathcal{P}(X)$ the usual weak* topology. The set $\{(x, \gamma) : \gamma \in \Gamma(x)\}$ is assumed to be a Borel subset of the product space $X \times \mathcal{P}(X)$. The utility function is a mapping from X to the real numbers with the usual interpretation that $u(x)$ represents the value to a player of each state $x \in X$. In this paper we assume that u is bounded and Borel measurable.

A *strategy* σ is a sequence $\sigma_0, \sigma_1, \dots$ such that $\sigma_0 \in \mathcal{P}(X)$, and, for $n \geq 1$, σ_n is a universally measurable mapping from X^n into $\mathcal{P}(X)$. A *strategy* σ is *available* in Γ at x if $\sigma_0 \in \Gamma(x)$ and $\sigma_n(x_1, \dots, x_n) \in \Gamma(x_n)$ for every $n \geq 1$ and $(x_1, \dots, x_n) \in X^n$.

Every strategy σ determines a probability measure, also denoted by σ , on the Borel subsets of the *infinite history space* $H = X \times X \times \dots$ with its product topology. Let X_1, X_2, \dots be the coordinate process on H . Then, under σ , X_1 has distribution σ_0 and, for $n \geq 1$, X_{n+1} has conditional distribution $\sigma_n(x_1, \dots, x_n)$ given $X_1 = x_1, \dots, X_n = x_n$.

We concentrate on *leavable* gambling problems in which a player chooses a time to stop play as well as a strategy. A *stop rule* is a universally measurable function from H into $\{0, 1, \dots\}$ such that whenever $t(h) = n$ and h' agrees with h in the first n coordinates, then $t(h') = n$. It is convenient to assume, as we now do, that, for all x , the point mass measure $\delta(x) \in \Gamma(x)$.

This does not affect the value of the optimal reward function defined below, but does simplify some algebraic expressions in the sequel.

A player, who begins with fortune x , selects a strategy σ available at x and a stop rule t . The player's expected reward is then

$$\int u(X_t) d\sigma,$$

where $X_0 = x$. The *optimal reward function* is defined for $x \in X$ to be

$$U(x) = \sup \int u(X_t) d\sigma,$$

where the supremum is over all σ at x and all stop rules t . The *n-day optimal reward function* U^n is defined, for $n \geq 1$, in the same way except that stop rules are restricted to satisfy $t \leq n$.

The *one-day operator* $G = G_\Gamma$ is defined on the collection $\mathcal{M}(X)$ of bounded universally measurable functions g by

$$Gg(x) = \sup \left\{ \int g d\gamma : \gamma \in \Gamma(x) \right\}, \quad x \in X.$$

By [2, Theorem 2.15.1], the *n-day optimal rewards* U^n can be calculated by backward induction using G :

$$U^1 = Gu, \quad U^{n+1} = GU^n. \tag{1}$$

Because the universal measurability of the U^n was shown in [13], the operator G is well-defined on these *n-day optimal reward functions*. Note that

$$U^n = G^n u,$$

where G^n is the composition of G with itself n times. Furthermore, it follows easily from the definitions of U and U^n that

$$U^n \leq U^{n+1} \leq U \quad \text{and} \quad U = \lim_n U^n. \tag{2}$$

3. Convergence of gambling houses

To define a notion of convergence for gambling houses on X , first let d_V be the *total variation distance* defined for probability measures $\gamma, \lambda \in \mathcal{P}(X)$ by

$$d_V(\gamma, \lambda) = \sup \left\{ \left| \int g d\gamma - \int g d\lambda \right| : g \in \mathcal{M}(X), \|g\| \leq 1 \right\},$$

where $\|g\| = \sup\{|g(x)| : x \in X\}$ is the supremum norm.

Next let d_H be the *Hausdorff distance* on subsets of $\mathcal{P}(X)$ associated with d_V ; that is, for subsets C and D of $\mathcal{P}(X)$, let

$$d_H(C, D) = \inf\{\varepsilon \geq 0 : C \subseteq D_\varepsilon, D \subseteq C_\varepsilon\},$$

where D_ε (respectively, C_ε) is the set of all $\gamma \in \mathcal{P}(X)$ such that $d_V(\gamma, D) \leq \varepsilon$ (respectively, $d_V(\gamma, C) \leq \varepsilon$). Finally, for gambling houses Γ and Λ on X , let

$$D(\Gamma, \Lambda) = \sup_{x \in X} d_H(\Gamma(x), \Lambda(x)).$$

A sequence of houses Γ_n *converges* to Γ if $D(\Gamma_n, \Gamma) \rightarrow 0$; we write $\Gamma_n \rightarrow \Gamma$ if this holds. Note that $\Gamma_n \rightarrow \Gamma$ means that $d_H(\Gamma_n(x), \Gamma(x)) \rightarrow 0$ uniformly in x .

Remark 1. Other measures of distance for gambling houses can be obtained by following the procedure above starting from a different measure of distance on $\mathcal{P}(X)$. For example, suppose that the topology on the state space X is given by a bounded metric, say $\rho: X \times X \mapsto [0, 1]$ and define the space of 1-Lipschitz functions:

$$\mathcal{L}(X) = \{g: X \mapsto \mathbb{R} \text{ such that } |g(x) - g(y)| \leq \rho(x, y) \text{ for all } x, y\}.$$

The well-known Kantorovich metric on $\mathcal{P}(X)$ is

$$d_K(\gamma, \lambda) = \sup \left\{ \int g \, d\gamma - \int g \, d\lambda : g \in \mathcal{L}(X) \right\} = \sup \left\{ \left| \int g \, d\gamma - \int g \, d\lambda \right| : g \in \mathcal{L}(X) \right\}.$$

The corresponding Hausdorff distance d_{HK} on subsets of $\mathcal{P}(X)$ and the distance D_K on gambling houses can be defined by analogy with d_H and D above. It is easy to see (and probably well-known) that d_K is dominated by d_V . It follows that D_K is dominated by D .

4. Continuity with respect to Γ

The following trivial example shows that the mapping $\Gamma \mapsto U$ is not continuous, in general, for the distance D defined above. Some more interesting examples will be given in Section 8.

Notation. When a sequence $\{\Gamma_n\}$ is considered below, the notation U_n is used for the optimal reward function of the house Γ_n , for each n , in order to avoid confusing it with the n -day optimal reward U^n of a given house Γ . Similarly, $U_n^k = G_{\Gamma_n}^k u$ is written for the k -day optimal reward function for Γ_n .

Example 1. Let $X = \{0, 1\}$ and $u(0) = 0, u(1) = 1$. Suppose that $\Gamma(0) = \{\delta(0)\}, \Gamma(1) = \{\delta(1)\}$, and, for $n \geq 1, \Gamma_n(0) = \{\delta(0), (1 - 1/n)\delta(0) + (1/n)\delta(1)\}$ and $\Gamma_n(1) = \{\delta(1)\}$. Then $\Gamma_n \rightarrow \Gamma$, but $U_n(0) = 1$ for all $n \geq 1$ and $U(0) = 0$.

Continuity does hold for finite-horizon problems and there is a form of lower semicontinuity in general.

Theorem 1. *Suppose that $\Gamma_n \rightarrow \Gamma$. Then,*

- (i) *for all $k \geq 1, \|U_n^k - U^k\| \rightarrow 0$ as $n \rightarrow \infty$; and*
- (ii) *for all $x \in X, \liminf_n U_n(x) \geq U(x)$.*

Our proof is based on the following lemma (but note the remarks after the proof).

Lemma 1. *Let $u, v \in \mathcal{M}(X), \gamma, \lambda \in \mathcal{P}(X), C, D$ be nonempty subsets of $\mathcal{P}(X)$, and Γ and Λ be gambling houses on X . Then the following hold:*

- (i) $|\int u \, d\gamma - \int u \, d\lambda| \leq \|u\|d_V(\gamma, \lambda);$
- (ii) $|\sup_{\gamma \in C} \int u \, d\gamma - \sup_{\lambda \in D} \int u \, d\lambda| \leq \|u\|d_H(C, D);$
- (iii) $|G_\Gamma u(x) - G_\Lambda u(x)| \leq \|u\|d_H(\Gamma(x), \Lambda(x)) \leq \|u\|D(\Gamma, \Lambda), x \in X;$
- (iv) $|\sup_{\gamma \in C} \int u \, d\gamma - \sup_{\gamma \in C} \int v \, d\gamma| \leq \|u - v\|;$
- (v) $|G_\Gamma u(x) - G_\Gamma v(x)| \leq \|u - v\|, x \in X;$
- (vi) $\|G_\Gamma^k u - G_\Lambda^k u\| \leq k\|u\|D(\Gamma, \Lambda).$

Proof. (i) This is clear if $\|u\| = 0$. If not then

$$\left| \int u \, d\gamma - \int u \, d\lambda \right| = \|u\| \left| \int \frac{u}{\|u\|} \, d\gamma - \int \frac{u}{\|u\|} \, d\lambda \right| \leq \|u\| d_V(\gamma, \lambda),$$

where the inequality holds by the definition of d_V .

(ii) Let $\varepsilon > 0$ and choose $\gamma^* \in C$ such that $\int u \, d\gamma^* \geq \sup_{\gamma \in C} \int u \, d\gamma - \varepsilon$. Then

$$\begin{aligned} \sup_{\gamma \in C} \int u \, d\gamma - \sup_{\lambda \in D} \int u \, d\lambda &\leq \int u \, d\gamma^* - \sup_{\lambda \in D} \int u \, d\lambda + \varepsilon \\ &= \inf_{\lambda \in D} \left[\int u \, d\gamma^* - \int u \, d\lambda \right] + \varepsilon \\ &\leq \|u\| \inf_{\lambda \in D} d_V(\gamma^*, \lambda) + \varepsilon \\ &= \|u\| d_V(\gamma^*, D) + \varepsilon \\ &\leq \|u\| d_H(C, D) + \varepsilon, \end{aligned}$$

where the second inequality above is justified by (i). Since ε is arbitrary, it follows that

$$\sup_{\gamma \in C} \int u \, d\gamma - \sup_{\lambda \in D} \int u \, d\lambda \leq \|u\| d_H(C, D).$$

By symmetry, the same inequality holds when the left-hand side is replaced by its negative. So (ii) follows.

(iii) The first inequality of (iii) is the special case of (ii) when $C = \Gamma(x)$ and $D = \Lambda(x)$. The second inequality follows from the definition of the distance D .

(iv) Calculate as follows:

$$\begin{aligned} \sup_{\gamma \in C} \int u \, d\gamma &= \sup_{\gamma \in C} \int ((u - v) + v) \, d\gamma \\ &\leq \sup_{\gamma \in C} \int (u - v) \, d\gamma + \sup_{\gamma \in C} \int v \, d\gamma \\ &\leq \|u - v\| + \sup_{\gamma \in C} \int v \, d\gamma. \end{aligned}$$

By symmetry, the same inequality holds with u and v interchanged, and (iv) follows.

(v) This is the special case of (iv) when $C = \Gamma(x)$.

(vi) This is proved by induction on k . The $k = 1$ case is just (iii). Assume the desired inequality holds for k , and calculate as follows:

$$\begin{aligned} \|G_\Gamma^{k+1} u - G_\Lambda^{k+1} u\| &= \|G_\Gamma(G_\Gamma^k u) - G_\Lambda(G_\Lambda^k u)\| \\ &\leq \|G_\Gamma(G_\Gamma^k u) - G_\Lambda(G_\Gamma^k u)\| + \|G_\Lambda(G_\Gamma^k u) - G_\Lambda(G_\Lambda^k u)\| \\ &\leq \|G_\Gamma^k u\| D(\Gamma, \Lambda) + \|G_\Gamma^k u - G_\Lambda^k u\| \\ &\leq \|u\| D(\Gamma, \Lambda) + k \|u\| D(\Gamma, \Lambda). \end{aligned}$$

The penultimate inequality uses (iii) and (v); the final inequality uses the easily checked fact that $\|G_\Gamma^k u\| \leq \|u\|$ and the inductive assumption. □

Proof of Theorem 1. (i) Apply Lemma 1(vi) to see that

$$\|U_n^k - U^k\| = \|G_{\Gamma_n}^k u - G_{\Gamma}^k u\| \leq k\|u\|D(\Gamma_n, \Gamma),$$

which converges to 0 as $n \rightarrow \infty$ by hypothesis.

(ii) Let $\varepsilon > 0$ and $x \in X$. By (2), there exists k so that $U^k(x) = G_{\Gamma}^k u(x) \geq U(x) - \varepsilon$. By (i),

$$|U_n^k(x) - U^k(x)| \rightarrow 0 \text{ as } n \rightarrow \infty.$$

Hence,

$$\liminf_n U_n(x) \geq \liminf_n U_n^k(x) = U^k(x) \geq U(x) - \varepsilon.$$

Because ε is arbitrary, the proof of (ii) is complete. □

Remark 2. A version of Theorem 1 can be proved for the distance D_K , which arises from the Kantorovich distance d_K on $\mathcal{P}(X)$ as explained in Remark 2. For the proof of the analogue of Lemma 1(vi), one needs to know that if u is 1-Lipschitz then the same is true of $G_{\Gamma}u$ and $G_{\Lambda}u$. A condition on a gambling house Γ , called $\Lambda(1)$, is given in [8] that guarantees that G_{Γ} preserves the space $\mathcal{L}(X)$ of 1-Lipschitz functions. Using this result, one can show that if Γ_n converges to Γ in the D_K distance and if Γ and all the Γ_n satisfy $\Lambda(1)$, then Theorem 1(i) and 1(ii) hold as before.

Remark 3. As a referee observed, another proof of Theorem 1(i) can be based on a coupling of strategies that are close together in the total variation distance. Another referee has pointed out that Theorem 1(ii) follows from Theorem 1(i). Thus, the lower semicontinuity property will be valid for any topology on gambling houses for which Theorem 1(i) holds.

Suppose now that the houses Γ_n approach Γ from below so that, in particular, $U_n \leq U$ for all n . Thus, if $\Gamma_n \rightarrow \Gamma$ then, by Theorem 1, $U_n \rightarrow U$. However, the convergence condition is not needed in this case.

Theorem 2. *Suppose that, for all $x \in X$ and all n , $\Gamma_n(x) \subseteq \Gamma_{n+1}(x) \subseteq \Gamma(x)$, and $\bigcup_n \Gamma_n(x) = \Gamma(x)$. Then $\lim_n U_n(x) = U(x)$ for all x .*

Proof. Let $Q = \lim_n U_n$. The limit is well defined since $U_n \leq U_{n+1}$ for all n . These inequalities hold because all strategies available in each Γ_n are also available in Γ_{n+1} . Also $u \leq Q \leq U$ because $u \leq U_n \leq U$ for all n . To show $Q \geq U$, it suffices to verify that Q is excessive for Γ ; see [2, Theorem 2.12.1] or [9, Lemma 3.1.2]. That is, it suffices to show that, for $x \in X$ and $\gamma \in \Gamma(x)$, $\int Q \, d\gamma \leq Q(x)$. Now $\gamma \in \Gamma(x)$ implies that $\gamma \in \Gamma_n(x)$ for sufficiently large n . Also U_n is excessive for Γ_n (see [2, Theorem 2.14.1] or [9, Lemma 3.1.4]), so $\int U_n \, d\gamma \leq U_n(x)$ for sufficiently large n . Hence, for $\gamma \in \Gamma(x)$,

$$\int Q \, d\gamma = \int \lim_n U_n \, d\gamma = \lim_n \int U_n \, d\gamma \leq \lim_n U_n(x) = Q(x). \quad \square$$

There is no result analogous to Theorem 2 for the case that the Γ_n approach Γ from above. This is illustrated by the following example.

Example 2. Let X, u, Γ be as in Example 1. For $n \geq 1$, define

$$\Gamma_n(1) = \{\delta(1)\}, \quad \Gamma_n(0) = \{\delta(0)\} \cup \left\{ \left(1 - \frac{1}{k}\right)\delta(0) + \left(\frac{1}{k}\right)\delta(1) : k \geq n \right\}.$$

Then $\Gamma_{n+1}(x) \subseteq \Gamma_n(x)$, and $\bigcap_n \Gamma_n(x) = \Gamma(x)$ for all n and $x = 0, 1$. However, $U(0) = 0$ and $U_n(0) = 1$ for all n .

5. Continuity with respect to u

In this section the state space X and gambling house Γ are held constant, and the optimal reward function U is considered as a function of the utility u .

Lemma 2. *Let (X, Γ, u) and (X, Γ, w) be gambling problems with optimal reward functions U and W , respectively. Then $\|U - W\| \leq \|u - w\|$.*

Proof. Each strategy σ and stop rule t determine a distribution for the random state X_t . Fix x and let C be the collection of all such distributions that can be obtained by choosing a strategy σ available in Γ at x and a stop rule t . Then $U(x) = \sup_{\gamma \in C} \int u \, d\gamma$ and $W(x) = \sup_{\gamma \in C} \int w \, d\gamma$. Now apply Lemma 1(iv) to see that $|U(x) - W(x)| \leq \|u - w\|$. \square

An immediate corollary is the continuity of the optimal reward as a function of the utility in the supnorm topology.

Corollary 1. *Let (X, Γ, u) and (X, Γ, u_n) , $n = 1, 2, \dots$, be gambling problems with optimal reward functions U and U_n , $n = 1, 2, \dots$, respectively. If $\|u_n - u\| \rightarrow 0$ then $\|U_n - U\| \rightarrow 0$.*

The optimal reward is not a continuous function of the utility for the topology of pointwise convergence, or the topology of uniform convergence on compact subsets. The latter topology corresponds on metric spaces to the topology of ‘continuous convergence’ used by Langen [7] in his study of related questions for dynamic programming models. Here is an example.

Example 3. Let $X = \mathbb{N}$ be the set of positive integers, and, for each $n \in \mathbb{N}$, let $\Gamma(n) = \{\delta(n), \delta(n+1)\}$. Then there is a strategy at each state under which the sequence of states moves deterministically up in steps of size 1. Now let u_n be the indicator function of $\{n, n+1, \dots\}$ so that u_n converges pointwise to the function u which is identically zero. It is trivial to check that, for each n , the optimal reward function for (X, Γ, u_n) is identically equal to 1, and that for (X, Γ, u) is identically 0.

6. Nonleavable gambling problems

A *nonleavable* gambling problem has the same three ingredients (X, Γ, u) as a leavable problem. However, in a nonleavable problem, the player is not allowed to stop the game. (The assumption that $\delta(x) \in \Gamma(x)$ for all x is not made in this section.) A player at an initial state x chooses a strategy σ available at x and is assigned as reward the quantity $u(\sigma) = \int [\limsup_n u(X_n)] \, d\sigma$. (This definition of $u(\sigma)$ is equivalent to that of Dubins and Savage as is explained in [9, Chapter 4].) The optimal reward $V(x)$ is defined to be the supremum over all σ at x of $u(\sigma)$.

The optimal reward function V is, in general, more difficult to calculate than U . There is an algorithm for V , but unlike the backward induction algorithm (1) for U , the algorithm for V is transfinite (cf. [3] or [9, Section 4.7]). It is not surprising that results like Theorems 1 and 2 fail to hold in the nonleavable case. The two examples below illustrate the failure of the analogues to the two theorems. In both examples there are gambling problems (X, Γ, u) and (X, Γ_n, u) , $n \in \mathbb{N}$, with associated optimal reward functions V and V_n , $n \in \mathbb{N}$. Also both examples have state space $X = \{0, 1\}$ and utility function $u(0) = 0$, $u(1) = 1$.

Example 4. Let $\Gamma(0) = \Gamma_n(0) = \{\delta(0)\}$, $\Gamma(1) = \{\delta(1)\}$, and $\Gamma_n(1) = \{(1 - 1/n)\delta(1) + 1/n\delta(0)\}$ for all $n = 1, 2, \dots$. Clearly, $\Gamma_n \rightarrow \Gamma$ and $V(0) = V_n(0) = 0$ for all n . It is also clear that $V(1) = 1$. However, $V_n(1) = 0$ for all n since under the unique strategy available at 1 in Γ_n , the process of states is eventually absorbed at 0 with probability 1.

Example 5. Let $\Gamma(0) = \Gamma_n(0) = \{\delta(0)\}$ for all n . Set $\gamma_n = (1 - 1/n)\delta(1) + 1/n\delta(0)$, $n = 1, 2, \dots$. Then let $\Gamma_n(1) = \{\gamma_1, \gamma_2, \dots, \gamma_n\}$ for each n and let $\Gamma(1) = \cup_n \Gamma_n(1) = \{\gamma_1, \gamma_2, \dots\}$. The hypotheses of Theorem 2 are satisfied and clearly $V(0) = V_n(0) = 0$ for all n . Also $V_n(1) = 0$ for each n since every gamble in $\Gamma_n(1)$ assigns probability of at least $1/n$ to state 0 so that the process of states must be absorbed at 0 with probability 1. However, $V(1) = 1$ because the player starting from state 1 in Γ can choose to play a sequence $\gamma_{n_1}, \gamma_{n_2}, \dots$ such that the product $\prod_k (1 - 1/n_k)$ is arbitrarily close to 1.

Unlike Theorems 1 and 2, the analogues to Lemma 2 and Corollary 1 do hold for nonleavable problems. Indeed, let (X, Γ, u) and (X, Γ, u') be gambling problems with optimal reward functions V and V' , respectively. Let σ be a strategy. One can check that $|u(\sigma) - u'(\sigma)| \leq \|u - u'\|$ and it follows that $\|V - V'\| \leq \|u - u'\|$. The exact analogue to Corollary 1 is immediate.

7. Red and black casinos

Dubins and Savage [2, p. 76] expressed particular interest in the continuity properties of the special class of gambling problems they called *casinos with a fixed goal*. These problems have the fortune space $X = [0, \infty)$ and the utility function u equal to the indicator of $[1, \infty)$. So the objective of a gambler is to reach a fortune of at least 1. The gambling house must satisfy two conditions expressed colorfully in [2] as ‘a rich gambler can do whatever a poor one can do’ and ‘a poor gambler can, on a small scale, imitate a rich one.’ For the formal definition, see [2, p. 64].

The next section has three examples to illustrate how discontinuities can occur in the special case of casinos with a fixed goal, and to answer, in part, the question raised by Dubins and Savage about such discontinuities. A different approach to the same question due to Dubins and Meilijson [1] is sketched in Section 9.

For convenience, the examples that follow are based on the red and black casinos of Dubins and Savage [2, Chapter 5]. For each $w \in [0, 1]$, the *red and black casino with parameter w* is the gambling house Γ_w defined by

$$\Gamma_w(x) = \{\gamma_w(s, x) : 0 \leq s \leq x\}, \quad x \in [0, \infty),$$

where $\gamma_w(s, x) = w\delta(x + s) + \bar{w}\delta(x - s)$ and $\bar{w} = 1 - w$. The optimal reward function for Γ_w is denoted by U_w .

Here are a few facts from [2].

- (1) For $\frac{1}{2} < w \leq 1$, Γ_w is *superfair* and $U_w(x) = 1$ for all $x > 0$.
- (2) For $w = \frac{1}{2}$, Γ_w is *fair* and $U_w(x) = x$ for $0 \leq x \leq 1$.
- (3) If $0 < w < \frac{1}{2}$, Γ_w is *subfair* and U_w is continuous, strictly increasing on $[0, 1]$ with $0 < U_w(x) < x$ for $0 < x < 1$. An optimal strategy for Γ_w in the subfair case is *bold play* which *stakes* $s(x) = \min(x, 1 - x)$ whenever the current state is $x \in [0, 1]$; that is, bold play uses the gamble $\gamma_w(s(x), x)$ at x .
- (4) If $0 < w < w' < \frac{1}{2}$ then $U_w(x) < U_{w'}(x)$ for $0 < x < 1$. (This follows from (3) since it is easily seen that bold play in Γ_w is less likely to reach 1 than bold play in $\Gamma_{w'}$ from an $x \in (0, 1)$.)
- (5) For $w = 0$, Γ_w is *trivial* and $U_w(x) = 0$ for $0 \leq x < 1$.

Another trivial casino is Γ_T defined by $\Gamma_T(x) = \{\delta(x)\}$ for all x . Obviously, the optimal reward function U_T of Γ_T satisfies $V_T(x) = 0$ for $0 \leq x < 1$.

8. Three examples

The first example highlights the phenomenon mentioned by Dubins and Savage [2, p. 76].

Example 6. A sequence of superfair casinos converging to a fair casino.

Let $\frac{1}{2} < w_n < 1$ for all n and suppose that $w_n \rightarrow \frac{1}{2}$ as $n \rightarrow \infty$. A simple calculation shows that, for all $x \geq 0$ and $0 \leq s \leq x$, $d_V(\gamma_{w_n}(s, x), \gamma_{1/2}(s, x)) \leq 2(w_n - \frac{1}{2})$. Consequently, $d_H(\Gamma_{w_n}(x), \Gamma_{1/2}(x)) \leq 2(w_n - \frac{1}{2})$ for all x so that $\Gamma_{w_n} \rightarrow \Gamma_{1/2}$. However, by points (1) and (2) of Section 7, $U_{w_n}(x) = 1$ and $U_{1/2}(x) = x$ for $0 < x < 1$. Hence, U_{w_n} does not converge to $U_{1/2}$.

The next two examples use modifications of red and black defined, for $0 \leq w \leq 1$, $x \geq 0$, $n \geq 1$, by

$$\Gamma_{w,n}(x) = \{\gamma_w(s, x, n) : 0 \leq s \leq x\},$$

where

$$\gamma_w(s, x, n) = \frac{w}{n}\delta(x + s) + \left(1 - \frac{1}{n}\right)\delta(x) + \frac{\bar{w}}{n}\delta(x - s).$$

Note that a gambler playing at position x in the casino $\Gamma_{w,n}$, $n > 1$, can, by repeatedly using $\gamma_w(s, x, n)$, eventually achieve the same outcome as a gambler playing at position x in $\Gamma_w = \Gamma_{w,1}$ who uses $\gamma_w(s, x)$.

Bold play in the house $\Gamma_{w,n}$ refers to the strategy that uses the gamble $\gamma_w(s(x), x, n)$ whenever the current state is $x \in [0, 1]$. As before $s(x) = \min(x, 1 - x)$.

Lemma 3. Assume that $0 < w \leq \frac{1}{2}$. Then, for all $n \geq 1$, bold play is optimal in the house $\Gamma_{w,n}$ and the optimal reward function $U_{w,n}$ for $\Gamma_{w,n}$ is equal to the optimal reward function U_w for Γ_w .

Proof. Let x, X_1, X_2, \dots be the process of fortunes of a gambler who begins with x and plays boldly in the house $\Gamma_{w,n}$. Let Y_1 be the first X_n that differs from x . Clearly, the distribution of Y_1 is $\gamma_w(s(x), x)$. If Y_1 is equal to 0 or 1, let $Y_2 = Y_1$. If $0 < Y_1 < 1$, let Y_2 be the next X_n different from Y_1 . Then the conditional distribution of Y_2 given that $Y_1 = y_1$ is $\gamma_w(s(y_1), y_1)$. Continue in this fashion to define x, Y_1, Y_2, \dots and note that this process has the same distribution as the process of fortunes for a gambler who begins with x and plays boldly in the house Γ_w . Now the probability that the process x, X_1, X_2, \dots reaches 1 is the same as that for the process x, Y_1, Y_2, \dots , and this probability is equal to $U_w(x)$ by point (3) of Section 7. So the gambler playing in $\Gamma_{w,n}$ can reach 1 from x with probability at least $U_w(x)$ and, hence, $U_{w,n}(x) \geq U_w(x)$.

For the opposite inequality, it suffices to show that U_w is excessive for $\Gamma_{w,n}$; see [2, Theorem 2.12.1] or [9, Theorem 3.1.1]. To see that this is so, let $0 < x < 1$, $0 \leq x \leq s$ and consider

$$\begin{aligned} \int U_w d\gamma_w(s, x, n) &= \frac{w}{n}U_w(x + s) + \left(1 - \frac{1}{n}\right)U_w(x) + \frac{\bar{w}}{n}U_w(x_s) \\ &= \frac{1}{n} \int U_w d\gamma_w(s, x) + \left(1 - \frac{1}{n}\right)U_w(x) \\ &\leq U_w(x). \end{aligned}$$

The last inequality holds because U_w is excessive for Γ_w ; see [2, Theorem 2.14.1] or [9, Theorem 3.1.1].

It now follows that bold play is optimal at x in the house $\Gamma_{w,n}$ because it reaches 1 with probability $U_w(x) = U_{w,n}(x)$. □

Example 7. (A sequence of subfair casinos converging to a trivial casino.) Let $0 < w < \frac{1}{2}$ and consider the sequence of casinos $\Gamma_{w,n}$. If $0 < x < 1$, $0 \leq s \leq x$ then $d_V(\gamma_w(s, x, n), \delta(x)) \leq 1/n$ and it follows that $d_H(\Gamma_{w,n}(x), \Gamma_T(x)) \leq 1/n$, where Γ_T is the trivial house from Section 7. Thus, $\Gamma_{w,n} \rightarrow \Gamma_T$. By Lemma 3 and point (3) of Section 7, $U_{w,n}(x) = U_w(x) > 0 = U_T(x)$ for $0 < x < 1$. So $U_{w,n}$ does not converge to U_T .

Example 8. (A sequence of subfair casinos converging to a subfair casino.) Let $0 < w < w' < \frac{1}{2}$ and define $\Gamma_n(x) = \Gamma_w(x) \cup \Gamma_{w',n}(x)$ for all $n \geq 1$ and $x \geq 0$. As in Example 7, $\Gamma_{w',n}$ converges to the trivial house Γ_T . Since $\delta(x) = \gamma_w(0, x) \in \Gamma_w(x)$ for all x , the trivial house is a subhouse of Γ_w . So it is easy to conclude that Γ_n converges to Γ_w . By point (4) of Section 7, $U_w(x) < U_{w'}(x)$ for $0 < x < 1$, and, by Lemma 4 below, the optimal reward function U_n of Γ_n is equal to $U_{w'}$ for all n . So U_n does not converge to U_w .

Lemma 4. For every $n \geq 1$, an optimal strategy in Γ_n is to play boldly in $\Gamma_{w',n}$. Hence, the optimal reward function of Γ_n is $U_n = U_{w'}$ for all n .

Proof. By Lemma 1, $U_w = U_{w',n}$ for all n and bold play is optimal for the house $\Gamma_{w',n}$. Clearly, $U_n \geq U_{w'}$ because every strategy available in $\Gamma_{w',n}$ is also available in the larger house Γ_n . To see that the reverse inequality $U_n \leq U_{w'}$ also holds, it suffices to show that $U_{w'}$ is excessive for Γ_n ; see [2, Theorem 2.12.1]. Now $U_{w'}$ is certainly excessive for $\Gamma_{w',n}$ since it is the optimal reward function for this house. So it suffices to show that $\gamma_w(s, x)U_{w'} \leq U_{w'}(x)$ for $x \geq 0$, $0 \leq s \leq x$. However,

$$\begin{aligned} \int U_{w'} d\gamma_w(s, x) &= wU_{w'}(x + s) + \bar{w}U_{w'}(x - s) \\ &\leq w'U_{w'}(x + s) + \bar{w}'U_{w'}(x - s) \\ &= \int U_{w'} d\gamma_{w'}(s, x) \\ &\leq U_{w'}(x). \end{aligned}$$

The first inequality above holds because $w < w'$ and $U_{w'}$ is nondecreasing; the final inequality holds because $U_{w'}$ is excessive for $\Gamma_{w'}$. □

Remark 4. It was proved in [8] that subfair casinos satisfy the condition $\Lambda(1)$ mentioned in Remark 2 and also that they are nonexpansive for the Kantorovitch metric, that is, that $d_K(\Gamma(x), \Gamma(y)) \leq d(x, y)$. Moreover, a subfair casino induces an acyclic law of motion (any monotone and strictly concave function decreases in expectation along the trajectories). Nevertheless, Example 8 shows that continuity fails even in that case.

9. A different approach to continuity

Dubins and Meilijson [1] defined measures of closeness for casinos that are different from that used above. For the purposes of comparison, one of these is described here. The definition begins with the notion of a *lottery* at a fortune x .

If γ is a gamble available at x in a casino Γ and Y is a random variable with distribution γ , then the lottery θ associated with γ is the distribution of $Y - x$. Suppose now that θ and θ' are lotteries with means μ and μ' , and distribution functions F and F' , respectively. A measure of distance used in [1] is

$$\rho(\theta, \theta') = \frac{\int |F(x) - F'(x)| dx}{-\mu - \mu'}.$$

(The application is to subfair casinos where the lotteries have negative means.) This distance is used to induce a measure of distance between subfair casinos for which there are interesting continuity results; see Theorem 1 and the Corollary to Theorem 2 in [1].

It may be helpful, as a referee suggested, to compare the distance ρ with the total variation distance d_V for lotteries from the casinos of Example 7 in Section 8. Let $0 < x < 1$, $0 < s \leq x$ and consider the gambles

$$\gamma = \delta(x) \in \Gamma_T(x), \quad \gamma_n = \frac{w}{n}\delta(x+s) + \left(1 - \frac{1}{n}\right)\delta(x) + \frac{\bar{w}}{n}\delta(x-s) \in \Gamma_{w,n}(x)$$

with associated lotteries

$$\theta = \delta(0), \quad \theta_n = \frac{w}{n}\delta(s) + \left(1 - \frac{1}{n}\right)\delta(0) + \frac{\bar{w}}{n}\delta(-s).$$

Then $d_V(\gamma, \gamma_n) = d_V(\theta, \theta_n) = 1/n \rightarrow 0$, but $\rho(\theta, \theta_n) = 1/(1 - 2w) \not\rightarrow 0$. Thus, the casinos $\Gamma_{w,n}$ do not approach Γ_T in the Dubins–Meilijson sense, and there is no violation of their continuity results when $U_{w,n}$ fails to converge to U_T .

10. Continuous-time problems

Consider the problem of controlling a continuous-time process $X = \{X_t, t \geq 0\}$ with state-space a Borel subset B of \mathbb{R}^n that satisfies a stochastic differential equation

$$dX_t = \mu(t) dt + \sigma(t) dW_t.$$

Here $\{W_t\}$ is a standard n -dimensional Brownian motion. The nonanticipative control processes $\mu(t)$ and $\sigma(t)$ take values in \mathbb{R}^n and the space \mathbb{M}^n of $n \times n$ matrices, respectively, and satisfy appropriate conditions to ensure the existence of a solution to the equation. There is given, for each $y \in B$, a nonempty control set $C(y) \subseteq \mathbb{R}^n \times \mathbb{M}^n$ from which the controller is required to choose the value of $(\mu(t), \sigma(t))$ whenever $X_t = y$. Assume also that the controller selects a stopping time τ for the controlled process and receives $\mathbb{E}[u(X_\tau)]$, where $u: I \mapsto \mathbb{R}$ is a bounded, Borel measurable utility function. Let $U(x)$ be the supremum of the controller's possible rewards starting from x .

Similar formulations of this 'continuous-time leavable gambling problem' were given in [6] and [5]. An explicit solution to the one-dimensional problem when B is an interval and u is continuous can be found in [4]. No such solution is likely in higher dimensions. However, it is straightforward to define a distance between problems starting from the Hausdorff distance on the control sets and proceeding by analogy with Section 3. Perhaps there are continuous-time versions of the discrete-time theorems above.

Suppose now that the controlled processes are one-dimensional with state-space the unit interval, and that the object of the controller is to reach 1. The problem was called a continuous-time casino problem by Pestien and Sudderth [11] if the control sets satisfy certain conditions

similar to those assumed by Dubins and Savage in the discrete-time case. Many of the properties from [2] have counterparts in continuous time. For example, the classification of casinos as being trivial, subfair, fair, or superfair still holds. Examples similar to those of Section 3 might be based on the continuous-time red and black model in [10].

There may also be a result for continuous-time subfair casinos analogous to those of Dubins and Meilijson [1]. In the continuous-time case, the optimal return is a function of the ratios μ/σ^2 , where μ and σ are the control variables [11, Theorem 4.1]. This suggests defining a notion of closeness based on these ratios.

Acknowledgements

We thank Roger Purves for reminding us of the article by Dubins and Meilijson. We also thank two anonymous referees for their comments that have improved both the exposition and the substance of the paper.

References

- [1] DUBINS, L. AND MEILIJSON, I. (1974). On stability for optimization problems. *Ann. Prob.* **2**, 243–255.
- [2] DUBINS, L. E. AND SAVAGE, L. J. (1965). *How to Gamble If You Must. Inequalities for Stochastic Processes*. McGraw-Hill, New York.
- [3] DUBINS, L., MAITRA, A., PURVES, R. AND SUDDERTH, W. (1989). Measurable, nonleavable gambling problems. *Israel J. Math.* **67**, 257–271.
- [4] KARATZAS, I. AND SUDDERTH, W. D. (1999). Control and stopping of a diffusion process on an interval. *Ann. Appl. Prob.* **9**, 188–196.
- [5] KARATZAS, I. AND WANG, H. (2000). Utility maximization with discretionary stopping. *SIAM J. Control Optimization* **39**, 306–329.
- [6] KARATZAS, I. AND ZAMFIRESCU, I.-M. (2006). Martingale approach to stochastic control with discretionary stopping. *Appl. Math. Optimization* **53**, 163–184.
- [7] LANGEN, H.-J. (1981). Convergence of dynamic programming models. *Math. Operat. Res.* **6**, 493–512.
- [8] LARAKI, R. AND SUDDERTH, W. D. (2004). The preservation of continuity and Lipschitz continuity by optimal reward operators. *Math. Operat. Res.* **29**, 672–685.
- [9] MAITRA, A. P. AND SUDDERTH, W. D. (1996). *Discrete Gambling and Stochastic Games*. Springer, New York.
- [10] PESTIEN, V. C. AND SUDDERTH, W. D. (1985). Continuous-time red and black: how to control a diffusion to a goal. *Math. Operat. Res.* **10**, 599–611.
- [11] PESTIEN, V. C. AND SUDDERTH, W. D. (1988). Continuous-time casino problems. *Math. Operat. Res.* **13**, 364–376.
- [12] STRAUCH, R. E. (1967). Measurable gambling houses. *Trans. Amer. Math. Soc.* **126**, 64–72.
- [13] SUDDERTH, W. D. (1969). On the existence of good stationary strategies. *Trans. Amer. Math. Soc.* **135**, 399–414.