

CAHIER DU LAMSADE

374

Avril 2016

A continuity Question of Dubins and Savage

R. Laraki, W.Sudderth

A Continuity Question of Dubins and Savage

R. Laraki* and W. Sudderth†

April 9, 2016

Abstract

Lester Dubins and Leonard Savage posed the question as to what extent the optimal reward function U of a gambling problem varies continuously in the gambling house Γ and utility function u . 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. However, it is lower semicontinuous in the sense that, if Γ_n converges to Γ , then $\liminf U_n \geq U$.

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: 60G40, 90C40, 93E20.

*CNRS, Lamsade University of Paris 9 Dauphine and Department of Economics, École Polytechnique, Paris, France, rida.laraki@dauphine.fr. Rida Laraki's work was supported in part by a grant administered by the French National Research Agency as part of the Investissements d'Avenir program (Idex [Grant Agreement No. ANR-11-IDEX-0003-02/Labex ECODEC No. ANR11- LABEX-0047]).

†School of Statistics, University of Minnesota, Minneapolis, MN, USA, bill@stat.umn.edu

1 Introduction

A basic question about any problem of mathematics 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 (1965) formulation, the processes available are specified by a gambling house Γ and the reward is in terms of a utility function u . (Definitions are in the next section.) Dubins and Savage ([2], page 76) suggest that a notion of convergence be defined for gambling houses so that the continuity properties of the mapping from Γ to the optimal reward function U can be studied. For the notion of convergence introduced in section 3 below, a trivial example in section 4 shows that the mapping $\Gamma \mapsto U$ is not continuous in general. However, by Theorem 1 below, 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 $\lim U_n = U$.

Dubins and Savage studied in detail the interesting special class of gambling problems they called casinos. They observed ([2], page 76) that a discontinuity occurs when a sequence of superfair casinos converges to a fair casino (cf. Example 3 below). 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 5 shows that a discontinuity can occur even when all the casinos are subfair. However, Dubins and Meilijson (1974) proved a continuity theorem for subfair casinos using a quite different notion of distance.

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 (1981). 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; section 3 defines the notion of convergence to be used; section 4 has the main convergence results; and sections 5 and 6 are about the special case of casinos.

2 Preliminaries

A Dubins-Savage gambling problem is composed of 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 (1967). This means that X is assumed to be a nonempty Borel subset of a complete separable metric space. So, in particular, X is separable metric. 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 will 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 Theorem 2.15.1 of [2], the n-day optimal rewards U_n can be calculated by backward induction using G :

$$U_1 = Gu, \quad U_{n+1} = GU_n. \quad (2.1)$$

Because the universal measurability of the U_n was shown in [8], the operator G is well-defined on these n-day optimal reward functions. Notice that

$$U_n = G^n u \quad (2.2)$$

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

$$U_n \leq U_{n+1} \leq U \quad \text{and} \quad U = \lim_n U_n. \quad (2.3)$$

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, D of $\mathcal{P}(X)$ let

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

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

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

A sequence of houses Γ_n is now said to converge to Γ if $D(\Gamma_n, \Gamma) \rightarrow 0$ and 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:*

$$\mathcal{L}(X) = \{g : g : X \mapsto \mathbb{R}, (\forall x, y)(|g(x) - g(y)| \leq \rho(x, y))\}.$$

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

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

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

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 the final section.

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, $U_k^n = G_{\Gamma_n}^k u$ 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 $n \geq 1$, $\Gamma_n(0) = \{\delta(0), (1 - 1/n)\delta(0) + (1/n)\delta(1)\}$, $\Gamma_n(1) = \{\delta(1)\}$. 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

- (a) $\|U_k^n - U_k\| \rightarrow 0$ as $n \rightarrow \infty$, for all $k \geq 1$,
- (b) $\liminf_n U^n(x) \geq U(x)$, for all $x \in X$.

A lemma is needed for 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) $|\int u d\gamma - \int u d\lambda| \leq \|u\| \cdot d_V(\gamma, \lambda)$,
- (ii) $|\sup_{\gamma \in C} \int u d\gamma - \sup_{\lambda \in D} \int u d\lambda| \leq \|u\| \cdot d_H(C, D)$,
- (iii) $|G_\Gamma u(x) - G_\Lambda u(x)| \leq \|u\| \cdot d_H(\Gamma(x), \Lambda(x)) \leq \|u\| \cdot D(\Gamma, \Lambda)$, $x \in X$,
- (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) $\|G_\Gamma^k u - G_\Lambda^k u\| \leq k\|u\| \cdot D(\Gamma, \Lambda)$.

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

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

where the inequality is by definition of d_V .

For part (ii), let $\epsilon > 0$ and choose $\gamma^* \in C$ such that

$$\int u d\gamma^* \geq \sup_{\gamma \in C} \int u d\gamma - \epsilon.$$

Then

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

The second inequality in the calculation above is by part (i). Because ϵ is arbitrary, it follows that

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

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

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

For part (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 part (iv) follows.

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

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

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

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

Now, to prove part (a) of Theorem 1, apply part (vi) of the lemma to see that

$$\|U_k^n - U_k\| = \|G_{\Gamma_n}^k u - G_\Gamma^k u\| \leq k\|u\| \cdot D(\Gamma_n, \Gamma),$$

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

To prove part (b) of the theorem, let $\epsilon > 0$ and $x \in X$. By (2.3) there exists k so that $U_k(x) = G_\Gamma^k u(x) \geq U(x) - \epsilon$. By part (a),

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

Hence,

$$\liminf_n U^n(x) \geq \liminf_n U_k^n(x) = U_k(x) \geq U(x) - \epsilon.$$

Because ϵ is arbitrary, the proof of part (b) 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 1. For the proof of the analogue of part (vi) of Lemma 1, 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 [4] 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 G_Γ in D_K distance and if G_Γ and all the Γ_n satisfy $\Lambda(1)$, then parts (a) and (b) of Theorem 1 hold as before.*

Suppose now that the houses Γ_n approach Γ from below so that, in particular, $U^n \leq U$ for all n . Thus, if $D(\Gamma_n, \Gamma) \rightarrow 0$, 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 $\cup_n \Gamma_n(x) = \Gamma(x)$. 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 ([2], Theorem 2.12.1 or [5], Lemma 3.1.2). That is, it suffices to show that, for $x \in X$ and $\gamma \in \Gamma(x)$, that $\int Q d\gamma \leq Q(x)$. Now $\gamma \in \Gamma(x)$ implies that $\gamma \in \Gamma_n(x)$ for n sufficiently

large. Also U^n is excessive for Γ_n ([2], Theorem 2.14.1 or [5], Lemma 3.1.4), so $\int U^n d\gamma \leq U^n(x)$ for n sufficiently large. Hence, for $\gamma \in \Gamma(x)$,

$$\int Q d\gamma = \int \lim_n U^n d\gamma = \lim_n \int U^n d\gamma \leq \lim_n U^n(x) = Q(x).$$

□

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

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

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

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

5 Red-and-Black Casinos

Dubins and Savage ([2], page 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], page 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. See Dubins and Meilijson [1] for another approach to the same question.

The examples to follow will, for convenience, be 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

$$\gamma_w(s, x) = w\delta(x + s) + \bar{w}\delta(x - s).$$

(Here $\delta(y)$ is the point mass at y 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 $1/2 < w \leq 1$, Γ_w is *superfair* and $U_w(x) = 1$ for all $x > 0$.
2. For $w = 1/2$, Γ_w is *fair* and $U_w(x) = x$ for $0 \leq x \leq 1$.
3. If $0 < w < 1/2$, 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 $0 < w < w' < 1/2$, then $U_w(x) < U_{w'}(x)$ for $0 < x < 1$. (This follows from item 3 since it is easily seen that bold play in Γ_w is less likely to reach one 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 $0 \leq x < 1$.

6 Three Examples

The first example is an instance of the phenomenon mentioned by Dubins and Savage ([2], page 76).

Example 3. *A sequence of superfair casinos converging to a fair casino.*

Let $1/2 < w_n < 1$ for all n and suppose that $w_n \rightarrow 1/2$ as $n \rightarrow \infty$. A simple calculation shows, for all $x \geq 0$, $0 \leq s \leq x$, that $d_V(\gamma_{w_n}(s, x), \gamma_{1/2}(s, x)) \leq 2(w_n - 1/2)$. Consequently, $d_H(\Gamma_{w_n}(x), \Gamma_{1/2}(x)) \leq 2(w_n - 1/2)$ for all x so that $\Gamma_{w_n} \rightarrow \Gamma_{1/2}$. However, by items 1 and 2 of the previous section, $U_{w_n}(x) = 1$ 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

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

Notice 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 $\gamma_w(s, x)$.

By *bold play* in the house $\Gamma_{w,n}$ is meant 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 2. *Assume $0 < w \leq 1/2$. 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 $\Gamma_{w,n}$ equals 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 equals 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 equals $U_w(x)$ by item 3 of the previous section. 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 $\Gamma_{w,n}$ ([2], Theorem 2.12.1 or [5], Theorem 3.1.1). To see that this is so, let $0 < x < 1$, $0 \leq s \leq x$ and consider

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

The last inequality holds because U_w is excessive for Γ_w ([2], Theorem 2.14.1 or [5], 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 4. A sequence of subfair casinos converging to a trivial casino.

Let $0 < w < 1/2$ 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 the previous section. Thus $\Gamma_{w,n} \rightarrow \Gamma_T$. By Lemma 1 and item 3 of the previous section, $U_{w,n}(x) = U_w(x) > 0 = U_T(x)$ for $0 < x < 1$. So $U_{w,n}$ does not converge to U_T .

Example 5. A sequence of subfair casinos converging to a subfair casino.

Let $0 < w < w' < 1/2$ and define $\Gamma_n(x) = \Gamma_w(x) \cup \Gamma_{w',n}(x)$ for all $n \geq 1$ and $x \geq 0$.

Lemma 3. 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 ([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 $x \geq 0$, $0 \leq s \leq x$. But

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

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

As in the previous example, $\Gamma_{w',n}$ 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 item 4 of the previous section, $U_w(x) < U_{w'}(x)$ for $0 < x < 1$; so the optimal reward functions $U^n = U_{w'}$ do not converge to the value function U_w .

Remark 3. *It was proved in [4] that subfair casinos satisfy the condition $\Lambda(1)$ mentioned in Remark 2 and also that they are non-expansive for the Kantorovitch metric, that is $d_K(\Gamma(x), \Gamma(y)) \leq d(x, y)$. Moreover, a subfair casino induces an acyclic law of motion (any monotone and strictly concave function decreases in expectation along the trajectories). Nevertheless, example 5 shows that continuity fails even in that case.*

7 Continuous-time problems

Perhaps there is a version of Theorem 1 that holds for continuous-time stochastic control problems, and perhaps the results of Dubins and Meilijson [1] generalize to the case of continuous-time casinos as formulated by Pestien and Sudderth [6]. It does seem likely that the examples of section 6 can be adapted to continuous-time.

Acknowledgement We thank Roger Purves for reminding us of the article by Dubins and Meilijson.

References

- [1] Dubins, L. and Meilijson, I. (1974). On stability for gambling problems. *Ann. Probab.* 2 243-255.
- [2] Dubins, L. and Savage, L. (1965). *How to Gamble If You Must: Inequalities for Stochastic Processes*. McGraw-Hill, New York. (Dover editions in 1976 and 2014.)
- [3] Langen, H-J. (1981). Convergence of dynamic programming models. *Math. Oper. Res.* 6 493-512.
- [4] Laraki R. and W. Sudderth (2004). The preservation of continuity and Lipschitz continuity by optimal rewards operators. *Mathematics of Operations Research* 29, 672-685.
- [5] Maitra, A. and Sudderth, W. (1996). *Discrete Gambling and Stochastic Games*. Springer-Verlag, New York.

- [6] Pestien, V. and Sudderth, W. (1988). Continuous-time casino problems.
Mathematics of Operations Research 13, 364-376.
- [7] Strauch, R. (1967). Measurable gambling houses. *Trans. Amer.Math. Soc.* 126 64-72.
- [8] Sudderth, W. (1969). On the existence of good stationary strategies.
Trans. Amer.Math. Soc. 135 399-414.