Optimal stopping and the sufficiency of randomized threshold strategies[*]

Vicky Henderson[†]   David Hobson[†]   Matthew Zeng[†]

Abstract

In a classical optimal stopping problem the aim is to maximize the expected value of a functional of a diffusion evaluated at a stopping time. This note considers optimal stopping problems beyond this paradigm. We study problems in which the value associated to a stopping rule depends on the law of the stopped process. If this value is quasi-convex on the space of attainable laws then it is well known that it is sufficient to restrict attention to the class of threshold strategies. However, if the objective function is not quasi-convex, this may not be the case. We show that, nonetheless, it is sufficient to restrict attention to mixtures of threshold strategies.

1 Introduction and main results

Let $Y = (Y_t)_{t \geq 0}$ be a time-homogeneous, continuous strong-Markov process. Let $\mathcal{T}$ be the set of all stopping times, and let $\mathcal{T}_T$ be the set of all (one- and two-sided) threshold stopping times, ie. stopping rules based on the first crossing of upper or lower thresholds. Let $V = V(\tau)$ be the value associated with a stopping rule τ. Consider the optimal stopping problem associated with V, ie. the problem of finding

$$V_*(\mathcal{S}) = \sup_{\tau \in \mathcal{S}} V(\tau) \tag{1.1}$$

where $\mathcal{S}$ is some set of stopping times (for example $\mathcal{S} = \mathcal{T}$ or $\mathcal{S} = \mathcal{T}_T$), and especially the problem of finding an optimizer for (1.1). We say the $V = V(\tau)$ is law invariant if, whenever σ, τ are stopping times, $\mathcal{L}(Y_\sigma) = \mathcal{L}(Y_\tau)$ implies that $V(\sigma) = V(\tau)$, where $\mathcal{L}(Z)$ is the law of Z. It follows that $V(\tau) = H(\mathcal{L}(Y_\tau))$ for some map H.

The following result is well-known, but we include it as a contrast to our result on the sufficiency of randomized threshold rules.

Main Result 1 (See Theorem 2 below). Suppose H is quasi-convex and lower semi-continuous. Then $V_*(\mathcal{T}_T) = V_*(\mathcal{T})$.

[†]University of Warwick, UK.
E-mail: vicky.henderson@warwick.ac.uk,d.hobson@warwick.ac.uk,m.zeng@warwick.ac.uk

Corollary 1. In the setting of Theorem 2, in solving the optimal stopping problem (1.1) over the set of all stopping times it is sufficient to restrict attention to threshold rules.

As the canonical example, consider expected utility where the utility is represented by a continuous, increasing function u. Then, $V(\tau) = \mathbb{E}[u(Y_\tau)]$, assuming that the expectation is well defined. It follows that V is law invariant. Indeed $V(\tau) = H(\mathcal{L}(Y_\tau))$ where $H(\zeta) = \int u(z)\zeta(dz)$. H is quasi-convex and lower semi-continuous. In this example it is well known that there is an optimal stopping rule which is of threshold form, see for example, Dayanik and Karatzas [5]. The fact that quasi-convexity means that there is no benefit from following randomized strategies is well understood in the economics literature, see Machina [14] Camerer and Ho [3], Wakker [21] and He et al [11].

Recently there has been a surge of interest in problems which, whilst they have the law invariance property, do not satisfy the quasi-convex criterion. Two examples are optimal stopping under prospect theory (Xu and Zhou [22]), and optimal stopping under cautious stochastic choice (Henderson et al [9]).

Introduce the set $\mathcal{T}_R$ of mixed or randomized threshold rules (ie stopping rules which are based on the first exit from a randomly chosen interval).

Main Result 2 (See Theorem 1 below)**.** Suppose law invariance holds for V, but not quasi-convexity for H. Then $V_*(\mathcal{T}_T) \leq V_*(\mathcal{T}_R) = V_*(\mathcal{T})$.

We will show by example that the first inequality may be strict.

Corollary 2. In the setting of Theorem 1, in solving the optimal stopping problem (1.1) over the set of all stopping rules it is sufficient to restrict attention to randomized threshold rules, but it may not be sufficient to restrict attention to (pure) threshold rules.

It should be noted that we do not include discounting in our analysis since a problem involving discounting does not satisfy the law invariance property. Nonetheless, as is well known, the conclusion of Corollary 1 remains true for the problem of maximizing discounted expected utility of the stopped process $V(\tau) = \mathbb{E}[e^{-\beta\tau}u(Y_\tau)]$. However, in problems which go beyond the expected utility paradigm, there are often modelling issues which mitigate against the inclusion of discounting. For this reason, historically the literature has concentrated on problems with no discounting. Finding the optimal stopping rule is often already challenging in these models.

The significance of Corollary 2 is as follows. In many classical models optimal stopping behavior involves stopping on first exit from an interval. If decision makers are observed to stop at levels which have already been visited by the process, then this behavior is inconsistent with the classical optimal stopping model. However, our result implies that the converse is not true: if decision makers are observed to stop only when the process is reaching new maxima or minima, then it does not necessarily mean that they are maximizers of expected payoffs. Instead the decision criteria may be more complicated, and they may be utilizing a randomized threshold rule.

2 Problem specification and the problem in natural scale

We work on a filtered probability space $(\Omega, \mathcal{F}, \mathbb{F} = \{\mathcal{F}_t\}_{t \geq 0}, \mathbb{P})$. Let $Y = (Y_t)_{t \geq 0}$ be a $(\mathbb{F}, \mathbb{P})$-stochastic process on this probability space with state space I which is an interval. Let $\bar{I}$ be the closure of I. We suppose that Y is a regular, time-homogeneous diffusion with initial value $Y_0 = y$ such that y lies in the interior of I. For Γ an interval in $\mathbb{R}$ or a rectangular set in $\mathbb{R}^2$, let $\mathcal{B}(\Gamma)$ denote the Borel σ-algebra on Γ, and let $\mathcal{P}(\Gamma)$ denote the set of probability measures on $(\Gamma, \mathcal{B}(\Gamma))$.

Let $\mathcal{T}$ be the class of all stopping times τ such that $\lim_{t \uparrow \infty} Y_{t \wedge \tau}$ exists (almost surely). We introduce two subclasses of stopping times

- $\mathcal{T}_T$, the subclass of (pure) threshold stopping times;
- $\mathcal{T}_R$, the subclass of randomized threshold stopping times.

Note that $\mathcal{T}_T \subset \mathcal{T}_R \subset \mathcal{T}$. The set of pure threshold stopping times can be written as

$$\mathcal{T}_T = \mathcal{T} \cap \left(\cup_{\beta \leq y \leq \gamma; \, \beta, \gamma \in \bar{I}} \{ \tau_{\beta, \gamma} \} \right), \tag{2.1}$$

where $\tau_{a,b} = \inf_{u \geq 0} \{ u : Y_u \notin (a, b) \}$. Note that if $a = y = b$ then $\tau_{a,b} \equiv 0$. Hence, $\tau \equiv 0$, the strategy of stopping immediately, lies in $\mathcal{T}_T$.

In order to be able to define a sufficiently rich class of randomized stopping times we need to assume that $\mathbb{F}$ is larger than the filtration generated by Y.

Assumption 1. $\mathcal{F}_0$ is sufficiently rich as to include a $U[0, 1]$ random variable, and the stochastic process Y is independent of this random variable.

Let $\mathcal{D} = ([-\infty, y] \cap \bar{I}) \times ([y, \infty] \cap \bar{I})$. It follows from the assumption that for any probability measure ζ on $(\mathcal{D}, \mathcal{B}(\mathcal{D}))$ there exists an $\mathcal{F}_0$-measurable random variable $\Theta = \Theta_\zeta = (A_\zeta, B_\zeta)$ such that (A_ζ, B_ζ) has law ζ. Then for any $\zeta \in \mathcal{P}(\mathcal{D})$ we can define the randomized stopping time τ_ζ as the first time Y leaves a random interval, where the interval is chosen at time 0 with law ζ. Then $\tau_\zeta = \tau_{A_\zeta, B_\zeta} = \inf \{ u : Y_u \notin (A_\zeta, B_\zeta) \}$. The set of randomized threshold rules $\mathcal{T}_R$ is given by

$$\mathcal{T}_R = \mathcal{T} \cap \left(\{ \tau_\zeta : \zeta \in \mathcal{P}(\mathcal{D}) \} \right). \tag{2.2}$$

Our analysis is focussed on problems in which the value associated with a stopping rule depends only on the law of the stopped process. Let $Q(\mathcal{S}) = \{ \mu : \mu = \mathcal{L}(Y_\tau), \tau \in \mathcal{S} \}$.

Assumption 2 (Law invariance). V is law invariant, ie $V(\tau) = H(\mathcal{L}(Y_\tau))$ for some function $H : Q(\mathcal{T}) \mapsto \mathbb{R}$.

Given that the value associated with a stopping rule is law invariant, one natural approach to finding the optimal stopping time is to try to characterize $Q(\mathcal{S})$. Often, the best way to do this is via a change of scale. Let s be a strictly increasing function such that $X = s(Y)$ is a local martingale. (Such a function s exists under very mild conditions on Y see, for example Rogers and Williams [16], and is called a scale function. For example, if Y solves the SDE $dY_t = \sigma(Y_t)dB_t + \xi(Y_t)dt$ for smooth functions σ and ξ with $\sigma > 0$ then $s = s(z)$ is a solution to $\frac{1}{2}\sigma(z)^2 s'' + \xi(z)s' = 0$. Note that if s is a scale function then so is any affine transformation of s and so we may choose any convenient normalization for s.) Let $I^X = s(I)$ and let $\bar{I}^X$ be the closure of I^X. Then X is a regular, time-homogenous local-martingale diffusion on I^X with initial value $x = s(y)$.

Set $Q^X(\mathcal{S}) = \{ \nu : \nu = \mathcal{L}(X_\tau), \tau \in \mathcal{S} \}$. Then if $\mathcal{L}(X_\tau) = \nu$ we have $\mathcal{L}(Y_\tau) = \nu \sharp s$ where $(\nu \sharp s)(D) = \nu(s(D))$. It follows that $\nu \in Q^X(\mathcal{S})$ if and only if $\nu \sharp s \in Q(\mathcal{S})$ and hence

$$Q(\mathcal{S}) = \{ \nu \sharp s; \nu \in Q^X(\mathcal{S}) \}. \tag{2.3}$$

Thus, if we can characterize $Q^X(\mathcal{S})$ then we can also characterize $Q(\mathcal{S})$. Moreover, defining $H^X : Q^X(\mathcal{T}) \mapsto \mathbb{R}$ by $H^X(\nu) = H(\nu \sharp s)$ we have $V_*(\mathcal{S}) = \sup_{\mu \in Q(\mathcal{S})} H(\mu) = \sup_{\nu \in Q^X(\mathcal{S})} H^X(\nu)$. The problem of optimizing over stopping laws for the problem with Y becomes a problem of optimizing over the possible laws of the stopped process X in natural scale.

Note that $\tau_{a,b} = \inf_{u \geq 0} \{ u : Y_u \notin (a, b) \} = \inf_{u \geq 0} \{ u : X_u \notin (s(a), s(b)) \} =: \tau_{s(a), s(b)}^X$. Hence $\mathcal{T}_T$ has the alternative representation

$$\mathcal{T}_T = \mathcal{T} \cap \left(\cup_{\beta \leq x \leq \gamma; \, \beta, \gamma \in \bar{I}^X} \{ \tau_{\beta, \gamma}^X \} \right),$$

and the set of threshold stopping times for Y is the set of threshold stopping times for X. Similarly, $\mathcal{T}_R$ can be rewritten as $\mathcal{T}_R = \mathcal{T} \cap \left(\{ \tau_\eta^X : \eta \in \mathcal{P}(\mathcal{D}^X) \} \right)$ where $\mathcal{D}^X = ([-\infty, x] \cap \bar{I}^X) \times ([x, \infty) \cap I^X))$ and

$$\tau_\eta^X = \inf_{u \geq 0} \{ u : X_u \notin (A_\eta, B_\eta) \text{where } (A_\eta, B_\eta) \text{ has law } \eta \}.$$

3 Characterizing the possible laws of the stopped process in natural scale

If $X = s(Y)$ is in natural scale then the state space of X is an interval $I^X = s(I)$ and $X_0 = x := s(y)$. There are four cases:

1. I^X is bounded;
2. I^X is unbounded above but bounded below;
3. I^X is bounded above but unbounded below;
4. I^X is unbounded above and below.

The third case can be reduced to the second by reflection. The first case is generally similar to the second case, and typically the proofs are similar but simpler. The final case is degenerate and will be treated separately. In the main text we will mainly present arguments for the second case (with the other cases covered in an appendix), but results will be stated in a form which applies in all cases.

Henceforth, in the main text we suppose I^X is bounded below, but unbounded above. Without loss of generality we may assume $I^X = (0, \infty)$ or $[0, \infty)$. Then X is a non-negative local martingale and hence a super-martingale. Moreover, $\lim_{t\to\infty} X_t$ exists. Hence $\mathcal{T}$ includes stopping rules which take infinite values and on $\{\tau = \infty\}$ we set $X_\tau = \lim_{t\to\infty} X_t = 0$. In this case $\mathcal{T}$ is the set of all stopping times and the intersection with $\mathcal{T}$ in the definitions (2.1) and (2.2) is not necessary. By Fatou's lemma and the super-martingale property

$$\mathbb{E}[X_\tau] = \mathbb{E}[\lim_{t\to\infty} X_{t\wedge\tau}] \le \liminf_{t\to\infty} \mathbb{E}[X_{t\wedge\tau}] \le x.$$

In particular, if we set $\mathcal{P}_{\le x} = \{\nu \in \mathcal{P}([0, \infty)) : \int z\nu(dz) \le x\}$ then $Q^X(\mathcal{T}) \subseteq \mathcal{P}_{\le x}$.

Lemma 1. $Q^X(\mathcal{T}) = Q^X(\mathcal{T}_R)$.

Proof. Here we prove the lemma in the case where I^X is bounded below. We show that $Q^X(\mathcal{T}) = Q^X(\mathcal{T}_R) = \mathcal{P}_{\le x}$. Given $\nu \in \mathcal{P}_{\le x}$ the aim is to find a stopping time $\tau \in \mathcal{T}_R$ such that $\mathcal{L}(X_\tau) = \nu$. The task of finding general stopping times with $\mathcal{L}(X_\tau) = \xi$ for given $\xi \in \mathcal{P}(\overline{I}^X)$ is known as the Skorokhod embedding problem (Skorokhod [18]). In fact we use an extension of an embedding due to Hall [8], see also Durrett [6]. The extension relates to the fact that we allow for target laws which have a different mean to the initial value of X, whereas the Hall embedding assumes $\int z\nu(dz) = x$. The Hall embedding, and the extension we give, are mixtures of threshold strategies.

Suppose ν is an element of $\mathcal{P}_{\le x}$ (and ν is not a point mass at x). The case of $\nu = \delta_x$ corresponds to the (threshold) stopping time $\tau = 0$. Let G be the (right-continuous) quantile function of ν. We have $x \ge \int z\nu(dz) = \int_{(0,1)} G(u)du$. In particular, unless $\lim_{u\uparrow 1} G(u) \le x$ there exists a unique solution $v^* \in [0, 1)$ to $\int_v^1 [G(w) - x]dw = 0$. Let $z^* = G(v^*) \le x$. If $\lim_{u\uparrow 1} G(u) \le x$ then set $v^* = 1$ and $z^* = \lim_{u\uparrow 1} G(u)$.

Let ν_0 be the measure of size v^* such that $\nu_0([0, z)) = v^* \wedge \nu([0, z))$. Then ν_0 has support contained in $[0, z^*]$. Let ν_1 be the measure of size $1 - v^*$ such that $\nu_1([0, z)) = (\nu([0, z)) - v^*)^+$. Then ν_1 has support in $[z^*, \infty)$ and barycentre x. Moreover $\nu = \nu_0 + \nu_1$.

Define $c = \int_x^\infty (y - x)\nu(dy)$. By construction, $c = \int_x^\infty (y - x)\nu_1(dy)$ and we have from the fact that ν_1 has barycentre x that $\int_{z^*}^\infty (y - x)\nu_1(dy) = 0$ and hence

$$c = \int_{z^*}^x (x - y)\nu_1(dy). \tag{3.1}$$

Let $\eta \in \mathcal{P}([0, x] \times (x, \infty])$ be given by

$$\eta(da, db) = \nu_0(da)I_{\{0\le a\le z^*\}}I_{\{b=\infty\}} + \nu_1(da)\nu_1(db)\frac{(b-a)}{c}I_{\{z^*\le a\le x<b<\infty\}}.$$

Note first that η is a probability measure:

$$
\int_{0 \leq a \leq x} \int_{x < b \leq \infty} \eta(da, db)
$$
$$
= v^* + \int_{z^* \leq a \leq x} \nu_1(da) \int_{x < b < \infty} \frac{b-x}{c} \nu_1(db) + \int_{z^* \leq a \leq x} \frac{x-a}{c} \nu_1(da) \int_{x < b < \infty} \nu_1(db)
$$
$$
= v^* + \int_{z^* \leq a \leq x} \nu_1(da) + \int_{x < b < \infty} \nu_1(db) = v^* + \nu_1([z^*, \infty)) = 1
$$

where we use the definition of c and (3.1) in going from the second line to the third.

It remains to show that $\mathcal{L}(X_{\tau_\eta^X}) = \nu$. Let f be a bounded test function. Then, using the fact that if $b = \infty$ then $X_{\tau_{a,\infty}^X} = a$, and the definition of c and (3.1) for the penultimate line,

$$
\mathbb{E}[f(X_{\tau_\eta^X})] = \int \int \eta(da, db) \mathbb{E}[f(X_{\tau_{a,b}^X})]
$$
$$
= \int \nu_0(da) f(a)
$$
$$
+ \int_{z^* \leq a \leq x} \int_{x < b < \infty} \nu_1(da) \nu_1(db) \frac{b-a}{c} \left[f(a) \frac{(b-x)}{b-a} + f(b) \frac{(x-a)}{b-a} \right]
$$
$$
= \int \nu_0(da) f(a) + \int_{z^* \leq a \leq x} \nu_1(da) f(a) \int_{x < b < \infty} \nu_1(db) \frac{(b-x)}{c}
$$
$$
+ \int_{z^* \leq a \leq x} \frac{(x-a)}{c} \nu_1(da) \int_{x < b < \infty} f(b) \nu_1(db)
$$
$$
= \int_{0 \leq z \leq z^*} f(z) \nu_0(dz) + \int_{z^* \leq z \leq x} f(z) \nu_1(dz) + \int_{x < z} f(z) \nu_1(dz)
$$
$$
= \int f(z) \nu(dz).
$$

Hence $\mathcal{L}(X_{\tau_\eta^X}) = \nu$ as required. $\qquad \square$

Let $\chi_{a,b} = \frac{b-x}{b-a} \delta_a + \frac{x-a}{b-a} \delta_b$. Then $\chi_{a,b}$ is the law of $X_{\tau_{a,b}^X}$. Moreover, $\mathcal{L}(X_{\tau_{a,\infty}^X}) = \delta_a$. Then,

$$
Q^X(\mathcal{T}_T) = (\cup_{0 \leq a \leq x} \delta_x) \cup (\cup_{0 \leq a < x < b < \infty} \chi_{a,b}).
$$

4 Sufficiency of mixed threshold rules

Our main result is that in a large class of problems it is sufficient to search over the class of mixed threshold rules.

Theorem 1. Suppose Y is a regular, time-homogeneous diffusion. Suppose the law invariance property holds (Assumption 2) and that the filtration is sufficiently rich (Assumption 1). Then $V_*(\mathcal{T}) = V_*(\mathcal{T}_R)$.

Proof. Since $Q^X(\mathcal{T}) = Q^X(\mathcal{T}_R)$ (Lemma 1) we have $Q(\mathcal{T}) = Q(\mathcal{T}_R)$. Then

$$
V_*(\mathcal{T}) = \sup_{\mu \in Q(\mathcal{T})} H(\mu) = \sup_{\mu \in Q(\mathcal{T}_R)} H(\mu) = V_*(\mathcal{T}_R). \qquad \square
$$

Note that it is not our claim that every optimal stopping rule is a mixed threshold rule. Typically, at least in the case where $V(\mathcal{T}_T) < V(\mathcal{T})$, there will be other optimal stopping rules which are not of threshold type.

4.1 Examples

4.1.1 Rank dependent utility and optimal stopping

Let Z be a non-negative random variable. Let $v : [0, \infty) \mapsto [0, \infty)$ be an increasing, differentiable function with $v(0) = 0$. Then the expected value of $v(Z)$ can be expressed as $\mathbb{E}[v(Z)] = \int_0^\infty v'(z)\bar{F}_Z(z)dz$. Under rank-dependent utility (Quiggin [15]) or probability weighting (Tversky and Kahneman [20]) the prospect value $\mathcal{E}_v(Z)$ of Z is

$$\mathcal{E}_v(Z) = \int_0^\infty v'(z)w(\bar{F}_Z(z))dz$$

where $w : [0, 1] \mapsto [0, 1]$ is an increasing, differentiable probability weighting function. Writing $G_Z = F_Z^{-1}$ for the quantile function of Z, then after a change of variable and integration by parts we have (see Xu and Zhou [22, Lemma 3.1]) the alternative representation

$$\mathcal{E}_v(Z) = \int_0^1 w'(1-u)G_Z(u)du.$$

Now let $Y = (Y_t)_{t \geq 0}$ be a non-negative diffusion and consider the problem of maximizing over stopping times the prospect value of the stopped process Y, ie of finding

$$\sup_{\tau \in \mathcal{T}} \mathcal{E}_v(Y_\tau). \tag{4.1}$$

Clearly the prospect value depends on the stopping time only through the law of the stopped process. Hence it is sufficient to characterize the optimal target distribution, for example via its quantile function. Xu and Zhou [22] solve for the optimal quantile function in several cases. One relevant case is the following:

Proposition 1 (Xu and Zhou [22]). Suppose Y is in natural scale and has state space $[0, \infty)$ and initial value y. Suppose v and w are concave. Suppose there exists $\lambda^* \in (0, \infty)$ which solves

$$\int_0^1 (v')^{-1}\left(\frac{\lambda^*}{w'(1-u)}\right)du = y.$$

Then the quantile function of the optimal stopping distribution is given by $G^*(u) = (v')^{-1}\left(\frac{\lambda^*}{w'(1-u)}\right)$.

Xu and Zhou [22] point out that although there is a unique optimal prospect there are infinitely many stopping rules which attain this prospect. They advocate the use of the stopping rule based on the Azéma-Yor stopping time [1], in which case the stopping rule has a drawdown feature, and involves stopping the first time the process falls below some function of the maximum. Our main result says that there is also a randomized threshold rule which is optimal.

4.1.2 Cautious stochastic choice

Given a process Y and a utility function u the certainty equivalent associated with a stopping time τ is $\mathcal{C}_u(\tau) = u^{-1}(\mathbb{E}[u(Y_\tau)])$. The idea in Cautious stochastic choice (Cerreia-Vioglio et al [4]) is that agents use multiple utility functions and evaluate an outcome in a robust manner as the least favorable of the individual certainty equivalents. If the set of utility functions is $\{u_\alpha\}_{\alpha \in \mathcal{A}}$, and if we write $\mathcal{C}_\alpha$ as shorthand for $\mathcal{C}_{u_\alpha}$ then the CSC value of a stopping rule is

$$CSC(\tau) = \inf_{\alpha \in \mathcal{A}} \mathcal{C}_\alpha(\tau) = \inf_{\alpha \in \mathcal{A}} u_\alpha^{-1}(\mathbb{E}[u_\alpha(Y_\tau)]), \tag{4.2}$$

and an optimal stopping rule is the one which maximizes the CSC value.

Clearly the CSC value of a stopping rule depends only on the law of Y_τ. Moreover, suppose $\mathcal{A} = \{\alpha, \beta\}$ and suppose u_α and u_β are strictly increasing and continuous with strictly increasing and continuous inverses. Suppose further that there exist τ_1 and τ_2 and $\tilde{y}$ such that $u_\alpha^{-1}(\mathbb{E}[u_\alpha(Y_{\tau_1})]) > \tilde{y} > u_\beta^{-1}(\mathbb{E}[u_\beta(Y_{\tau_1})])$ and $u_\alpha^{-1}(\mathbb{E}[u_\alpha(Y_{\tau_2})]) < \tilde{y} < u_\beta^{-1}(\mathbb{E}[u_\beta(Y_{\tau_2})])$. Let τ^θ be a mixture of τ_1 and τ_2, defined such that if Z is a $\mathcal{F}_0$-measurable random variable taking values in $\{1, 2\}$ with $\mathbb{P}(Z = 1) = \theta$ then $\tau^\theta = \tau_Z$. Then for $\gamma \in \mathcal{A}$, $\mathcal{C}_\gamma(\tau^\theta) = u_\gamma^{-1}(\theta \mathbb{E}[u_\gamma(Y_{\tau_1})] + (1 - \theta)\mathbb{E}[u_\gamma(Y_{\tau_2})])$ is a continuous function of θ. Moreover, $\mathcal{C}_\alpha(\tau^\theta)$ is strictly increasing in θ and $\mathcal{C}_\beta(\tau^\theta)$ is strictly decreasing. By our assumptions it follows that the best choice θ^* of θ is such that $\mathcal{C}_\alpha(\tau^{\theta^*}) = \mathcal{C}_\beta(\tau^{\theta^*})$ and then $\theta^* \in (0, 1)$ and $CSC(\tau^{\theta^*}) > \max\{CSC(\tau_1), CSC(\tau_2)\}$.

In particular, the value associated with a stopping rule is not quasi-convex. By the analysis of this section, in searching for an optimal stopping rule it is sufficient to restrict attention to randomized threshold rules, but we cannot expect in general that there is a pure threshold rule which is optimal. For a deeper study of optimal stopping in the context of Cautious stochastic choice see Henderson et al [9].

5 Sufficient conditions for the optimality of pure threshold rules

In this section we argue that if the value associated with a stopping rule is law invariant, and if H is quasi-convex and lower semi-continuous then pure threshold rules are optimal.

Recall that H is quasi-convex if $H(\lambda \mu_1 + (1 - \lambda)\mu_2) \leq \max\{H(\mu_1), H(\mu_2)\}$ for $\lambda \in (0, 1)$. It follows by induction that if $\mu = \sum_{i=1}^{N} \lambda_i \mu_i$ where $\lambda_i \geq 0$, $\sum_{i=1}^{N} \lambda_i = 1$ and $\mu_i \in Q(\mathcal{T})$ then

$$H(\mu) \leq \max_{1 \leq i \leq N} H(\mu_i) \leq \sup_{\tilde{\mu} \in Q(\mathcal{T})} H(\tilde{\mu}). \tag{5.1}$$

Recall also that if H is lower semi-continuous and $\mu_n \Rightarrow \mu$ then $H(\mu) \leq \liminf H(\mu_n)$. In fact we do not require $H(\mu) \leq \liminf H(\mu_n)$, but rather the weaker condition $H(\mu) \leq \limsup H(\mu_n)$.

Lemma 2. Suppose $\nu \in Q^X(\mathcal{T})$ consists of finitely many atoms. Then there exists $\eta \in \mathcal{P}(\mathcal{D}^X)$ such that η consists of finitely many atoms and $\mathcal{L}(X_{\tau_\eta^X}) = \nu$.

Proof. It follows from the construction in the proof of Lemma 1 that if μ is purely atomic then so is η. □

Lemma 3. Let ν be an element of $Q^X(\mathcal{T})$. Then there exist $(\eta_n)_{n \geq 1}$ such that η_n has finite support for each n and such that $\mathcal{L}(X_{\tau_{\eta_n}^X}) \Rightarrow \nu$.

Proof. Since $\nu \in Q^X(\mathcal{T}) = Q^X(\mathcal{T}_R)$ there exists η such that $\mathcal{L}(X_{\tau_\eta^X}) = \nu$. Let $(\eta_n)_{n \geq 1}$ be a sequence of measures with finite support such that $\eta_n \Rightarrow \eta$. Then for $f : [0, \infty) \mapsto \mathbb{R}$ a bounded continuous test function define $\tilde{f} : [0, x] \times [x, \infty)$ by $\tilde{f}(a, b) = f(a)\frac{b-x}{b-a} + f(b)\frac{x-a}{b-a}$ for $a < b$ with $\tilde{f}(x, x) = f(x)$. Then, since $\tilde{f}$ is bounded and continuous

$$\mathbb{E}[f(X_{\tau_{\eta_n}^X})] = \int \int \eta_n(da, db)\tilde{f}(a, b) \to \int \int \eta(da, db)\tilde{f}(a, b) = \mathbb{E}[f(X_{\tau_\eta^X})]$$

and it follows that $\nu_n := \mathcal{L}(X_{\tau_{\eta_n}^X}) \Rightarrow \nu$. □

Theorem 2. Suppose Y is a regular, time-homogeneous diffusion. Suppose the law invariance property holds (Assumption 2). Suppose that H is quasi-convex and lower semi-continuous. Then $V_*(\mathcal{T}) = V_*(\mathcal{T}_T)$.

Proof. Clearly $V_*(\mathcal{T}) \geq V_*(\mathcal{T}_T)$.

For any μ_n with finite support we can define $\nu_n = \mu_n \sharp s^{-1}$. Then we can find a measure η_n with finite support such that $\mathcal{L}(X_{\tau^X_{\eta_n}}) = \nu_n$. Moreover ν_n can be decomposed as a convex combination

$$\nu_n = \sum_{i=1}^{N} \gamma_i \chi_{a_i,b_i} + \sum_{j=1}^{M} \lambda_j \delta_{a_j}.$$

Then, since H is quasi-convex,

$$
\begin{aligned}
H(\mu_n) &\leq \left(\max_{1 \leq i \leq N} H(\chi_{a_i,b_i} \sharp s^{-1}) \right) \vee \left(\max_{1 \leq j \leq M} H(\delta_{s^{-1}(a_j)}) \right) \\
&\leq \left(\sup_{0 \leq a \leq x \leq b < \infty} H(\chi_{a,b} \sharp s^{-1}) \right) \vee \left(\sup_{0 \leq a \leq x} H(\delta_{s^{-1}(a)}) \right) = V_*(\mathcal{T}_T).
\end{aligned}
$$

Then, for $\tau \in \mathcal{T}$, if $\mu = \mathcal{L}(Y_\tau)$ and if $\mu_n \Rightarrow \mu$

$$V_\tau = H(\mu) \leq \limsup H(\mu_n) \leq V_*(\mathcal{T}_T).$$

Hence $V_*(\mathcal{T}) \leq V_*(\mathcal{T}_T)$. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\square$

6 Discussion

In classical optimal stopping problems involving maximizing expected utility the optimal strategy is a threshold rule and involves stopping the first time that the process leaves an interval. However, in more general settings the optimal strategy may be more sophisticated. In some settings, for example those involving regret (Loomes and Sugden [13]) the optimal stopping rule may depend on some functional of the path (for example the maximum price to date). But, as argued here, for a large class of problems the payoff depends only on the distribution of the stopped process, and then there are many optimal stopping rules, some of which take the form of randomized threshold rules. In this article we have utilized (an extended version of) the Hall solution of the Skorokhod embedding problem (Hall [8]) to give our randomized threshold rule, but there are other solutions of the Skorokhod embedding which can also be viewed as mixed threshold rules, including the original solution of Skorokhod [18] and the solution of Hirsch et al [12].

The idea that if the objective is expressed in terms of a function which is not quasi-convex then agents may want to use randomized strategies is well appreciated in static settings. In a dynamic setting He et al [11] argue that in binomial-tree, probability-weighted model of a casino (Barberis [2]) gamblers may prefer path-dependent strategies over strategies which are defined via a partition of the set of nodes into those at which the gambler stops and those at which he continues. (See also Ebert and Strack [7] and Henderson et al [10] for discussion of a related optimal stopping problem with probability weighting based on a diffusion process.) He et al [11] argue further that the path-dependent strategy can be replaced by a randomized strategy under which the decision about whether to stop at a node depends not on the path history but rather the realization of an independent uniform random variable. This preference for randomization mirrors our result, but takes a different form. In our perpetual problem the agent chooses a randomized pair of levels and then follows a threshold strategy based on these levels. In He et al [11] a zero-one decision about whether to stop at a node is replaced by a probability of continuing, and the stopping rules which arise are not randomized threshold rules.

Many optimal stopping models in the economics literature predict that the agent will stop on first exit from an interval, which necessarily involves stopping either at

the current maximum or the current minimum. If instead, observed behavior includes stopping at levels which are not equal to one of the running extrema of the process then this is evidence against the model. (Strack and Viefers [19] present experimental evidence from a laboratory game that players do not follow threshold strategies - instead players visit the same price three times on average before stopping.) But, our results imply that the converse is not true. Even if agents only ever take a decision to sell at a time when the process is at a new maximum or new minimum, this does not necessarily mean that agents are following a pure threshold rule. They could have any target distribution, as for example in Proposition 1, but be realizing this target distribution via a randomized threshold rule.

References

[1] Azéma J. and M. Yor, 1979, Une solution simple au problème de Skorokhod.*Sem. de Prob. XIII*, 90-115. MR-0544782

[2] Barberis N., 2012, A Model of Casino Gambling, *Management Science*, 58, 35-51.

[3] Camerer, C. F., and T. Ho, 1994, Violations of the betweenness axiom and nonlinearity in probabilities, *Journal of Risk and Uncertainty*, 8, 167-196.

[4] Cerreia-Vioglio, S., D.Dillenberger, P. Ortoleva, and G. Riella, 2017, Deliberately Stochastic, Working paper, Columbia University. *SSRN Working paper, id2977320*.

[5] Dayanik S. and I. Karatzas, 2003, On the optimal stopping problem for one-dimensional diffusions. *Stoc. Proc. & Appl*, 107, 2, 173-212. MR-1999788

[6] Durrett R., 1991, *Probability: Theory and Examples*. Wadsworth, Pacific Grove, California. MR-1068527

[7] Ebert S. and P. Strack, 2015, Until the Bitter End: On Prospect Theory in a Dynamic Context. *American Economic Review*, 105(4), 1618-1633.

[8] Hall W.J., 1968, On the Skorokhod embedding theorem. *Technical Report 33* Stanford University, Department of Statistics.

[9] Henderson V., D. Hobson and M. Zeng, 2017, Cautious Stochastic Choice, Optimal Stopping and Deliberate Randomization, *SSRN Working paper, id3118906*.

[10] Henderson V., Hobson D. and A.S.L.Tse, 2017, Randomized Strategies and Prospect Theory in a Dynamic Context, *Journal of Economic Theory*, 168, 287-300. MR-3603292

[11] He X., S. Hu, J. Obloj and X.Y. Zhou, 2017, Path dependent and randomized strategies in Barberis' Casino Gambling model, *Operations Research*, 65, 1, 97-103. MR-3613093

[12] Hirsch F., C. Profetta, B. Roynette and M. Yor, 2011, Constructing self-similar martingales via two Skorokhod embeddings. *Sem. de Prob. XLIII* 451-503 LNM 2006, Springer-Verlag, Berlin. MR-2790387

[13] Loomes G. and R. Sugden, 1982, Regret theory: An alternative theory of rational choice under uncertainty, *Economic Journal*, 92, 805-824.

[14] Machina M., 1985, Stochastic Choice Functions Generated from Deterministic Preferences over Lotteries, *Economic Journal*, 95, 379, 575-594.

[15] Quiggin J., 1982, A Theory of Anticipated Utility, *Journal of Economic Behaviour and Organisation*, 3, 323-343.

[16] Rogers L.C.G. and D. Williams, 2000, *Diffusions, Markov Processes and Martingales: Itô Calculus* Wiley, Chichester. MR-0921238

[17] Rogozin B.A., 1996, On the distribution of functionals related to boundary problems for processes with independent increments. *Th. Prob. Appl.*, 11, 580-591.

[18] Skorokhod A.V., 1965, *Studies in the theory of random processes*, Addison-Wesley, Reading, Mass.. MR-0185620

[19] Strack P. and P. Viefers, 2017, Too Proud to Stop: Regret in Dynamic Decisions, *SSRN Working paper, id2465840*.

[20] Tversky, A. and D. Kahneman, 1992, Advances in Prospect Theory: Cumulative Representation of Uncertainty, *Journal of Risk and Uncertainty*, 5, 297-323.

[21] Wakker P., 2010, Prospect Theory for Risk and Ambiguity, *Cambridge University Press.* MR-2829560

[22] Xu Z.Q. and X.Y. Zhou, 2013, Optimal stopping under probability distortion. *Ann. Appl. Prob.*, 23, 1, 251-282. MR-3059235

A Extension to other state spaces for the process in natural scale

A.1 The range is unbounded below but bounded above

In this case we may assume that $I^X = (-\infty, 0)$ or $(-\infty, 0]$. The analysis goes through almost unchanged except that now X is a convergent sub-martingale and $Q^X(\mathcal{T}) = Q(\mathcal{T}_R) = \mathcal{P}_{\geq x}$ where $\mathcal{P}_{\geq x} = \{\nu \in \mathcal{P}((-\infty, 0]) : \int z\nu(dz) \geq x\}$.

A.2 The range is bounded

Suppose X is bounded. In this case $Q(\mathcal{T}) = Q(\mathcal{T}_R) = \mathcal{P}_{=x}$ where $\mathcal{P}_{=x} = \{\nu \in \mathcal{P}(\bar{I}^X) : \int z\nu(dz) = x\}$. To see this note that X is a uniformly integrable martingale and not just a super-martingale. Therefore we must have $\mathbb{E}[X_\tau] = \lim \mathbb{E}[X_{\tau \wedge t}] = x$ and hence $Q(\mathcal{T}) \subseteq \mathcal{P}_{=x}$. Conversely, by the same argument as in Lemma 1, but this time with $v^* = 0$ and $\nu_1 \equiv \nu$, we deduce that for any $\nu \in \mathcal{P}_{=x}$ there exists a randomization η such that $\mathcal{L}(X_{\tau_\eta^X}) = \nu$. It follows that $Q(\mathcal{T}) = Q(\mathcal{T}_R) = \mathcal{P}_{=x}$.

The proofs of Lemma 2, Lemma 3 and Theorem 1 go through unchanged.

A.3 The range is unbounded above and below

Now suppose I^X is unbounded above and below. By the Rogozin trichotomy (Rogozin [17]) $-\infty = \liminf_t X_t < x < \limsup_t X_t = \infty$ and $\lim_{t \uparrow \infty} X_t$ does not exist. In this case we must restrict $\mathcal{T}$ to the set of stopping times with $\mathbb{P}(\tau < \infty) = 1$. In the main text we set $\mathcal{T}_T = \mathcal{T} \cap \left(\cup_{\beta \leq y \leq \gamma, \beta, \gamma \in \bar{I}^Y} \{\tau_{\beta, \gamma}\} \right)$ but we could equivalently write $\mathcal{T}_T = \cup_{(\beta, \gamma) \in \mathcal{D}_0} \{\tau_{\beta, \gamma}\}$, where $\mathcal{D}_0 = ([-\infty, y] \cap \bar{I}^Y) \times ([y, \infty] \cap \bar{I}^Y) \setminus \{s^{-1}(-\infty), s^{-1}(\infty)\}$. We have to exclude the threshold rule $\tau_{s^{-1}(-\infty), s^{-1}(\infty)}$ since $\tau_{s^{-1}(-\infty), s^{-1}(\infty)} = \infty$ almost surely and Y_∞ is not defined. In terms of threshold rules $\tau_{a,b}^X$ for X we allow $a = -\infty$ or $b = \infty$ but not both. Then $\mathcal{T}_T = \{\tau_{\beta, \gamma} : (\beta, \gamma) \in \mathcal{D}_0^X)\}$ where $\mathcal{D}_0^X = \mathcal{D}^X \setminus \{-\infty, \infty\} = [\infty, x] \times [x, \infty] \setminus \{-\infty, \infty\}$.

In the definition of randomized threshold rules we can write $\mathcal{T}_R = \{\tau_\zeta : \zeta \in \mathcal{P}(\mathcal{D}_0)\}$ where $\mathcal{D}_0$ is as above and similarly $\mathcal{T}_R = \{\tau_\eta^X : \eta \in \mathcal{P}(\mathcal{D}_0^X)\}$.

When $I^X = \mathbb{R}$ we claim that we have $Q^X(\mathcal{T}) = Q^X(\mathcal{T}_R) = \mathcal{P}(\mathbb{R})$. Since stopping times are finite almost surely we must have $Q^X(\mathcal{T}) \subseteq \mathcal{P}(\mathbb{R})$ so it is sufficient to show that for any $\nu \in \mathcal{P}(\mathbb{R})$ we have $\nu \in Q^X(\mathcal{T}_R)$. Given $\nu \in \mathcal{P}(\mathbb{R})$ let A_ν be a $\mathcal{F}_0$-measurable random variable with law ν and set $\tau = \inf\{u : X_u = A_\nu\}$. Then $\mathcal{L}(X_\tau) = \mathcal{L}(A_\nu) = \nu$.

The proofs of Lemma 2, Lemma 3 and Theorem 1 go through unchanged.

A.4 Other results

Proof of Proposition 1. A proof is given in Xu and Zhou [22, Theorem 5.1], but since it is short, elegant and pertinent to our main results we include it here. From the characterization of $Q(\mathcal{T})$ we have that a quantile function must satisfy $\int_0^1 G(u)du \leq y$. By construction G^* has this property, and since v' and w' are decreasing, G^* is increasing. Hence G^* has the properties required of a quantile function of a distribution which can be obtained by stopping Y. On the other hand, for any non-negative function G with

$\int_0^1 G(u)du \leq y$,

$$
\begin{aligned}
\int_0^1 w'(1-u)v(G(u))du &= \int_0^1 [w'(1-u)v(G(u)) - \lambda^* G(u)]du + \lambda^* \int_0^1 G(u)du \\
&\leq \int_0^1 \sup_{g>0}[w'(1-u)v(g) - \lambda^* g]du + \lambda^* y \\
&= \int_0^1 [w'(1-u)v(G^*(u)) - \lambda^* G^*(u)]du + \lambda^* y \\
&= \int_0^1 w'(1-u)v(G^*(u))du. \quad \square
\end{aligned}
$$

Electronic Journal of Probability
Electronic Communications in Probability

Advantages of publishing in EJP-ECP

- Very high standards

- Free for authors, free for readers

- Quick publication (no backlog)

- Secure publication (LOCKSS[1])

- Easy interface (EJMS[2])

Economical model of EJP-ECP

- Non profit, sponsored by IMS[3], BS[4] , ProjectEuclid[5]

- Purely electronic

Help keep the journal free and vigorous

- Donate to the IMS open access fund[6] (click here to donate!)

- Submit your best articles to EJP-ECP

- Choose EJP-ECP over for-profit journals

---

[1] LOCKSS: Lots of Copies Keep Stuff Safe http://www.lockss.org/
[2] EJMS: Electronic Journal Management System http://www.vtex.lt/en/ejms.html
[3] IMS: Institute of Mathematical Statistics http://www.imstat.org/
[4] BS: Bernoulli Society http://www.bernoulli-society.org/
[5] Project Euclid: https://projecteuclid.org/
[6] IMS Open Access Fund: http://www.imstat.org/publications/open.htm