ECONSTOR
Make Your Publications Visible.

A Service of
ZBW
Leibniz-Informationszentrum Wirtschaft
Leibniz Information Centre for Economics

Quah, John K.-H.; Strulovici, Bruno

Working Paper

Discounting and patience in optimal stopping and control problems

Discussion Paper, No. 1480

Provided in Cooperation with:
Kellogg School of Management - Center for Mathematical Studies in Economics and Management Science, Northwestern University

Suggested Citation: Quah, John K.-H.; Strulovici, Bruno (2009) : Discounting and patience in optimal stopping and control problems, Discussion Paper, No. 1480, Northwestern University, Kellogg School of Management, Center for Mathematical Studies in Economics and Management Science, Evanston, IL

This Version is available at:
https://hdl.handle.net/10419/59672

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

WWW.ECONSTOR.EU

Mitglied der Leibniz-Gemeinschaft

Northwestern University
2001 Sheridan Road 580 Leverone Hall Evanston, IL 60208-2014 USA

Discussion Paper #1480
October 14, 2009

"Discounting and Patience in Optimal Stopping and Control Problems"

Key words: capital growth, comparative statics, discounting, internal rate of return, optimal control, optimal stopping, patience, present value, project valuation

JEL classification: C61, D81, D90 E2, G11, G31

John K.-H Quah
Oxford University

Bruno Strulovici
Northwestern University

www.kellogg.northwestern.edu/research/math

CMS-EMS
The Center for Mathematical Studies in Economics & Management Sciences

Discounting and Patience in Optimal Stopping and Control Problems

John K.-H. Quah Bruno Strulovici*

October 14, 2009

Abstract

This paper establishes that the optimal stopping time of virtually any optimal stopping problem is increasing in "patience," understood as a particular partial order on discount rate functions. With Markov dynamics, the result holds in a continuation-domain sense even if stopping is combined with an optimal control problem. Under intuitive additional assumptions, we obtain comparative statics on both the optimal control and optimal stopping time for one-dimensional diffusions. We provide a simple example where, without these assumptions, increased patience can precipitate stopping. We also show that, with optimal stopping and control, a project's expected value is decreasing in the interest rate, generalizing analogous results in a deterministic context. All our results are robust to the presence of a salvage value. As an application we show that the internal rate of return of any endogenously-interrupted project is essentially unique, even if the project also involves a management problem until its interruption. We also apply our results to the theory of optimal growth and capital deepening and to optimal bankruptcy decisions.

*Email addresses: john.quah@economics.ox.ac.uk b-strulovici@northwestern.edu

1 Introduction

One of the most frequent comparative statics results in economic theory concerns the impact of discounting on optimal decisions. Consider for example the case of an entrepreneur deciding on the length of a project, or of a household deciding on a refinancing decision. How does the timing of such decision vary with the decision maker's patience, or discount factor? The answer would seem a priori to depend on the distribution of future cash flows. For example, what if positive cash flows repeatedly alternate with negative ones? What if there is negative intertemporal correlation across cash flow shocks?

As it turns out, these issues are not relevant: optimal stopping *always* increases with patience, for any cash flow dynamics, including non-Markov processes, and even in the presence of an arbitrary salvage value process. The result applies for example to experimentation, where the state variable represents one's accumulated knowledge about the underlying payoff distribution of a process. Indeed, the result holds not only for exponential discounting, but for any discount functions which can be compared as follows: let $t \mapsto \alpha(t)$ denote a positive discount function (thus, $\alpha(t) = e^{-rt}$ for exponential discounting) and β denote another discount function. Say that a decision maker is more patient under β than under α if β/α is increasing in t, so that payoffs decreases are a relatively slower rate under β than under α. More patience, in that sense, results in a later optimal stopping time. This result generalizes Quah and Strulovici (2009), who prove it for the case of deterministic cash flows.

It is well-known that if there are no restrictions on the flow of costs and benefits, then the value of a project is not necessarily a decreasing function of the interest rate. This observation has given rise to a significant literature that examines conditions under which monotonicity holds (see Arrow and Levhari (1969), Sen (1975), Ross, Spatt, and Dybvig (1979) and, on a closely related issue, Foster and Mitra (2003)), which are in part motivated by the need to provide microeconomic motivation for an investment curve that is downward sloping (relative to the interest rate). In particular, Arrow and Levhari (1969) showed, in a deterministic model, that a project's value *does* decrease with the interest rate if it may be optimally stopped. We generalize their result to a stochastic context; specifically, if the discount functions are normalized by setting $\beta(0) = \alpha(0)$, then the expected value of a project (with optimal stopping) under β is larger than its expected value under α.

In realistic problems, stopping is not the only decision. For example, an entrepreneur needs to manage a project on a continuous basis before eventually deciding when to interrupt it. How sensitive are our results to the introduction of such a control problem? The question is subtle, because decision makers with varying degrees of patience may apply different controls and hence drive an initially common state to widely different situations. Nonetheless, it is possible to show the following: first, the value function is *always* increasing with patience, the previous caveat notwithstanding, and second, when the underlying state has Markov dynamics, the continuation domain (i.e. set of states where not stopping is optimal) is increasing with patience. These results must not be confused with the claim that the optimal stopping time increases, because a more patient decision maker may drive the state faster to the stopping boundary. Indeed, the paper contains a simple example where more patience precipitates stopping.

In order to compare the actual duration of a project under control, we consider the case of a one-dimensional diffusion and show that under two intuitive assumptions (payoff increasing in state and state drift increasing in control), the optimal control, state path, and duration, are all increasing with patience. In particular, these results on the relationship between the state and control variables and the discount rate hold even when (as in problems of optimal growth) the optimization problem does not typically involve a stopping decision. These results have a natural extension to the case of multiple control variables which we also consider in the paper. Lastly, our results extend to the case where stopping entails a salvage value, and so in principle to switching problems, i.e., where sequential stopping decisions are made, such as in entry–exit problems (see, e.g., Dixit, 1987).

We provide several specific applications. (1) The flip side of the observation that the value of a project need not be a decreasing function of the interest rate is that the internal rate of return of a project is not necessarily unique. Arrow and Levhari (1969) adopted a definition of the internal rate of return that involved an optimal stopping time, and showed that the internal rate of return (so defined) is unique. We extend Arrow and Levhari's result to our context, in which we allow for (i) stochastic cash flows, (ii) project management up until termination, and (iii) the presence of a salvage value. (2) We also consider an optimal growth problem with stochastic shocks and multidimensional control, and show that more patience results in "capital deepening," i.e., almost surely, the economy with the more patient representative agent will have a higher capital stock at any time. This

provides a stochastic complement to Quah and Strulovici (2009)'s result for a deterministic environment. (3) Finally, we apply our theory to optimal bankruptcy decisions. We show that when bankruptcy is endogenously chosen, lower interest results in later bankruptcy decision, remarkably even when shareholders or households face asset management problems on top of their bankruptcy decision.

2 Pure Optimal Stopping Problem

Let $\{x_t\}$ be a stochastic process taking values in some topological space $\mathcal{X}$ adapted to some filtered probability space $(\Omega, P, \mathcal{F} = \{F_t\})$. An agent with the instantaneous utility $u : (x,t) \to u(x,t) \in R$ solves the following optimal stopping problem:

$$\text{maximize } E\left[\int_0^{\hat{\tau}} \alpha(s)u(x(s),s)ds\right] \text{ subject to } \hat{\tau} \in \mathcal{T}, \tag{1}$$

where $\mathcal{T}$ denotes the set of stopping times adapted[1] to $\mathcal{F}$ and taking values in the (possibly infinite) time interval $T = [0, \bar{t}]$. E denotes the expectation operator and we assume that the *discount function* $\alpha : T \to R$ is deterministic and strictly positive. We assume throughout that $E[(\int_0^{\bar{t}} \alpha(s)u(x(s),s)ds)^2]$ is finite, which guarantees that all expectations (including conditional expectations) are well-defined, and that Fubini's Theorem can be used.

Before proceeding with the analysis, we state the following lemma.

Lemma 1 *Let τ and $\hat{\tau}$ denote two stopping times, and consider the events*

$$\begin{aligned} A &= \left\{\omega \in \Omega : \hat{\tau} \leq \tau \text{ and } E\left[\int_{\hat{\tau}}^{\tau} \alpha(s)u(x(s),s)ds|F_{\hat{\tau}}\right] < 0\right\} \text{ and} \\ B &= \left\{\omega \in \Omega : \hat{\tau} \geq \tau \text{ and } E\left[\int_{\hat{\tau}}^{\tau} \alpha(s)u(x(s),s)ds|F_{\hat{\tau}}\right] > 0\right\}. \end{aligned}$$

If τ solves (1)*, then $P(A) = P(B) = 0$.*

Proof. Let $A_\varepsilon = \{\omega \in \Omega : \hat{\tau} \leq \tau \text{ and } E\left[\int_{\hat{\tau}}^{\tau} \alpha(s)u(x(s),s)ds|F_{\hat{\tau}}\right] \leq -\varepsilon\}$ for $\varepsilon > 0$. If we prove that $P(A_\varepsilon) = 0$ for all ε, continuity of P and monotonicity of A_ε will imply that

[1] When the process x has jumps, the decision maker may stop immediately after a jump, in effect interrupting the utility flow at the jump. The addition of a salvage value is studied in Section 4.

$P(A) = P(\cap A_\varepsilon) = 0$. Suppose on the contrary that $P(A_\varepsilon) > 0$ for some $\varepsilon > 0$, and let $\tau^* = \hat{\tau} 1_{\omega \in A_\varepsilon} + \tau 1_{\omega \notin A_\varepsilon}$. Since A_ε is $F_{\hat{\tau}}$ measurable, τ^* is a stopping time. Moreover, letting $g(s) = \alpha(s)u(x(s), s)$ and $A_\varepsilon^c = \Omega \setminus A_\varepsilon$,

$$\begin{aligned} E\left[\int_0^{\tau^*} \alpha(s)u(x(s),s)ds\right] &= P(A_\varepsilon^c)E\left[\int_0^{\tau^*} g(s)ds|A_\varepsilon^c\right] + P(A_\varepsilon)E\left[\int_0^{\tau^*} g(s)ds|A_\varepsilon\right] \\ &= P(A_\varepsilon^c)E\left[\int_0^{\tau} g(s)ds|A_\varepsilon^c\right] + P(A_\varepsilon)E\left[\int_0^{\tau^*} g(s)ds|A_\varepsilon\right] \\ &\geq P(A_\varepsilon^c)E\left[\int_0^{\tau} g(s)ds|A_\varepsilon^c\right] + P(A_\varepsilon)(E\left[\int_0^{\tau} g(s)ds|A_\varepsilon\right] + \varepsilon) \\ &= E\left[\int_0^{\tau} g(s)ds\right] + P(A_\varepsilon)\varepsilon, \end{aligned}$$

which contradicts optimality of τ. The equality $P(B) = 0$ is proved similarly. ■

We define the function $v : R_+ \to R$ by

$$v(s) = E[u(x(s), s)1_{s<\tau}]. \tag{2}$$

$v(s)$ is the expected payoff rate at time s, where the payoff is zero in the event that one has stopped before s (i.e. $s \geq \tau$). Denoting $W(\alpha) = E\left[\int_0^\tau \alpha(s)u(x(s), s)ds\right]$, Fubini's theorem implies that

$$W(\alpha) = \int_0^{\bar{t}} \alpha(s)v(s)ds. \tag{3}$$

Our first result is a simple consequence of the fact that, at every point in time, the expected payoff of an optimizing agent looking forward must be non-negative.

Lemma 2 *For all t in $[0, \bar{t})$, $\int_t^{\bar{t}} \alpha(s)v(s)ds \geq 0$.*

Proof. By definition of v,

$$\begin{aligned} \int_t^{\bar{t}} \alpha(s)v(s)ds &= E\left[\int_t^{\bar{t}} \alpha(s)u(x(s),s)1_{s<\tau}ds\right] \\ &= E\left[E\left[\int_t^{\bar{t}} \alpha(s)u(x(s),s)1_{s<\tau}1_{t<\tau}ds|F_t\right]\right] \\ &= E\left[E\left[\int_t^{\tau} \alpha(s)u(x(s),s)ds|F_t\right]1_{t<\tau}\right]. \end{aligned}$$

Optimality of τ and Lemma 3 imply that the inner expectation is almost surely nonnegative if $t < \tau$. Therefore, the random variable $E[\int_t^{\bar{t}} \alpha(s)u(x(s), s)1_{s<\tau}ds|F_t]1_{t<\tau}$ is always nonnegative, and so is its expectation. ■

Lemma 2 leads to the following result.

THEOREM 1 *Let τ and τ' be solutions to the optimal stopping problem (1) when the discount functions are α and β respectively. If $\beta(s)/\alpha(s)$ is increasing*[2] *in s, then $\tau \vee \tau'$ is also an optimal stopping time for the discount function β.*

Theorem 1 relies on the following lemma, whose proof can be found in Quah and Strulovici (2009).

LEMMA 3 *Suppose $[x', x'']$ is a compact interval of R and $\hat{\alpha}$ and h are real-valued functions defined on $[x', x'']$, with h integrable and $\hat{\alpha}$ increasing (and thus integrable as well). If $\int_x^{x''} h(s)ds \geq 0$ for all x in $[x', x'']$, then*

$$\int_{x'}^{x''} \hat{\alpha}(s)h(s)ds \geq \hat{\alpha}(x') \int_{x'}^{x''} h(s)ds. \tag{4}$$

We now conclude the proof of Theorem 1.

Proof. Consider any outcome $\omega \in \Omega$, and $t < \tau(\omega)$. It is enough to show that

$$E\left[\int_t^{\tau} \beta(s)u(x(s), s)|F_t\right] \geq 0$$

as it implies that waiting until τ is at least weakly better than stopping immediately. We can set without loss of generality $t = 0$, since the problem could otherwise be restated with the origin of time at t. We wish to show that $\int_0^{\bar{t}} \beta(s)v(s)ds \geq 0$ with v as defined by (2). Lemma 2 guarantees that the hypothesis of Lemma 3 is satisfied, so we obtain

$$\int_0^{\bar{t}} \beta(s)v(s)ds = \int_0^{\bar{t}} \frac{\beta(s)}{\alpha(s)} \alpha(s)v(s)ds \geq \frac{\beta(0)}{\alpha(0)} \int_0^{\bar{t}} \alpha(s)v(s)ds = W(\alpha) \geq 0. \tag{5}$$

■

The proof also implies that the value $W(\cdot)$ of the stopping problem decreases with the interest rate or, equivalently, increases with patience.

COROLLARY 1 *[Value Monotonicity] Suppose that the ratio $s \mapsto \beta(s)/\alpha(s)$ is increasing and that $\alpha(0) = \beta(0)$. Then $W(\beta) \geq W(\alpha)$.*

[2] Throughout, "increasing" and "decreasing" must be understood in a weak sense.

Proof. Without loss of generality, suppose that $\alpha(0) = \beta(0) = 1$. By definition, $W(\beta)$ is the utility achieved at an optimal stopping time τ' for the discount function β. So $W(\beta) \geq E_x[\int_0^{\tau} \beta(s)u(x(s), s)ds] = \int_0^{\bar{t}} \beta(s)v(s)ds$. The result then follows (5). ■

Theorem 1 and Corollary 1 tell us that as future utility is discounted less, the optimal horizon gets longer and the optimal value gets larger, independently of the particular stochastic process and payoff function under consideration. Suppose that the discount functions are exponential, i.e., $\alpha(s) = \exp(-\bar{\alpha}s)$ and $\beta(s) = \exp(-\bar{\beta}s)$, where α and β are positive scalars. Then the ratio $\beta(s)/\alpha(s)$ is increasing in s if $\bar{\beta} < \bar{\alpha}$. More generally, if we allow for a nonconstant discount rate, we have $\alpha(s) = \exp(-\int_0^s r_\alpha(z)dz)$ where the function r_α is positive and deterministic. Writing a similar expression for $\beta(s)$, it is easy to check that $\beta(s)/\alpha(s)$ is increasing in s if $r_\alpha(z) > r_\beta(z)$ for all z in $T = [0, \bar{t}]$.

3 Optimal Stopping Combined with Optimal Control

In many situations, the decision maker can also control the state, possibly at some cost. For example, an entrepreneur makes multiple decisions concerning a project in addition to its interruption time. As it turns out, it is possible to extend several results to that case.

Consider the following optimization problem:

$$\text{maximize } E\left[\int_0^{\hat{\tau}} \alpha(s)u(x(s), \lambda(s), s)ds\right], \tag{6}$$

where the process λ is adapted to the filtration of x, and the law of the process $\{x_t\}_{t\geq 0}$ is controlled by λ in the sense that for each t the law of $\{x_s\}_{s\leq t}$ depends not only on exogenous uncertainty but also on the path $\{\lambda_s\}_{s\leq t}$. Later we will focus on the case where x is a Markov process, but such assumption is not needed for now. The process $\lambda = \{\lambda_s\}_{s\geq 0}$ is an *admissible* control, i.e. an adapted process taking values in some set[3] Λ, such that $E\left[(\int_0^{\bar{t}} \alpha(s)u(x(s), \lambda(s), s)ds)^2\right]$ is finite.

The first result is that the value function is increasing with patience, which generalizes Corollary 1. Let $W(\alpha)$ denote the value function resulting from optimization of (6).

[3]One could allow Λ to deterministically vary with time.

Theorem 2 (General Value Monotonicity) *Suppose that $\beta(s)/\alpha(s)$ is increasing in s and that $\alpha(0) = \beta(0)$. Then $W(\beta) \geq W(\alpha)$.*

Proof. Without loss of generality, suppose that $\alpha(0) = \beta(0) = 1$. Consider an optimal pair (τ, λ), assuming it exists,[4] under the discount function α, and let $y_t = (x_t, \lambda_t)$ denote the state-control pair, at any time. Now suppose that the same pair (τ, λ) is used under the discount function β. In that case, we are brought back to the previous analysis where the new state is y. We then have $W(\beta) \geq E_x[\int_0^\tau \beta(s)u(x(s), \lambda(s), s)ds] = \int_0^{\bar{t}} \beta(s)v(s)ds$. The result then follows from (5), applied to the state y_t. ■

It would be tempting to replicate this argument to conclude that the optimal stopping is also increasing with patience. However, this approach cannot be applied because, when optimal controls differ under α and β, the state processes will follow different paths which may yield non-comparable stopping times. Such an example is provided in Section 3.3.

When the state has Markov dynamics, one can show in full generality that the *continuation domain*, i.e. the set of states for which stopping is suboptimal, increases with patience. We show this result first, and then use it to show, under some key assumptions, that the optimal stopping time does increase with patience.

3.1 Domain-Based Comparative Statics for Markov Processes

Suppose now that x is a Markov process, i.e. that for any t and control $\{\lambda_s\}_{s\geq t}$, the distribution of $\{x_s\}_{s\geq t}$ only depends on past history through x_t.

In that case, it is well known[5] that there exists a Markov optimal control, i.e. deterministic functions $L : (x, t) \mapsto L(x, t) \in \Lambda$ and $C : t \mapsto C(t) \subset \mathcal{X}$ such that the pair $\lambda_t = L(x_t, t)$ and $\tau = \inf\{t : x_t \notin C(t)\}$ defines an optimal policy.

C is the (optimal) *continuation domain* of x. In general, C varies in time. If x has time-homogeneous dynamics, C is a constant subset of $\mathcal{X}$. For example, if x is a one-dimensional controlled diffusion, then C is an interval $[a, b]$ containing the initial state x_0, such that it is

[4] Otherwise, an approximation argument could be made.

[5] See e.g. Øksendal (2002) for the case of diffusion processes.

optimal to stop exactly when x_t hits one boundary of that interval.

Let $C(t, \alpha)$ denote the continuation domain at time t under discount function α

THEOREM 3 *Suppose that $\beta(s)/\alpha(s)$ is increasing in s. Then, for all t,*

$$C(t, \alpha) \subset C(t, \beta).$$

Proof. Suppose that $x_t \in C(t, \alpha)$. We need to show that $x_t \in C(t, \beta)$. Suppose that the discount function is β. Stopping at t yields a value of zero. Following the same control and stopping time as with discount function α yields a value function $\tilde{W}(\beta)$ greater than $W(\alpha)$, from Theorem 2 and its proof. Since $W(\alpha)$ is nonnegative, this shows that this control yields a nonnegative value under discount function β, and hence that continuing is optimal. ■

Theorem 3 applies, for example, to experimentation problems where the state x is the belief a decision maker has about one arm (see, for example, Bolton (1999), Keller et al. (2005), or Strulovici (2009)). Theorem 3 states the experimentation *domain*, i.e. the range of beliefs for which experimentation takes place, increases with the decision maker's patience. It does *not* generally imply, however, that the actual time spent experimenting increases with patience, because a more patient decision maker may experiment at a faster rate than a less patient one, resulting in the experimentation boundary being reached faster. This general distinction is developed below in an abstract example, which could be easily adapted to an experimentation context.

3.2 Time-Based Comparative Statics for One-Dimensional Markov Processes

Consider the following control problem:[6]

$$V(x, r) = \max_{\lambda, \tau} E\left[\int_0^\tau e^{-rt} v(x_t, \lambda_t, t) dt\right]$$

subject to

$$dx_t = \mu(x_t, \lambda_t, t)dt + \sigma(x_t, t)dB_t \qquad x_0 = x,$$

[6]We focus on exponential discounting, although any parameterized family $\{s \mapsto \alpha(s, r)\}_r$ of discount functions ranked according to order of Theorem 1 and differentiable with respect to r would yield a similar result.

where x is one dimensional and B is the standard Brownian motion and we assume for simplicity that σ is uniformly bounded above zero and that the value function is well defined and smooth.[7]

We make the following assumptions:

1. μ is increasing in λ for all x, t.
2. v is increasing in x for all λ, t.

Given the Markov nature of the problem, there exist boundaries $a(t) \leq b(t)$ such that the optimal stopping time is given by $\tau = \inf\{x_t \notin (a_t, b_t)\}$. Given the monotonicity assumption on v, the value function $V(x, r)$ is increasing in x since for $y > x$, one can always replicate the control applied when starting from x and get a higher payoff flow at all times before stopping. This implies in particular, that $b_t = \infty$ for all t, since the value function cannot be positive somewhere below b and equal to zero above b. Thus, stopping is determined by a possibly time–varying lower boundary $a(t)$. Moreover, the optimal control only depends on the current state and time (see Øksendal, 2002). Let $\lambda : (x, t, r) \mapsto \lambda(x, t, r)$ denote the control at time t and state x that is optimal under discount r (dependence on r or t will be dropped from the notation when the context is unambiguous). Let $\tau(x, r)$ denote the (stochastic) optimal stopping time starting from x when the discount rate is r. And let $x_t(r)$ denote the state at time t when using the control that is optimal for r (although we drop it in the notation, we only compare states across discount rates when starting from the same initial condition x).

Theorem 4 *The function $(x, t, r) \mapsto \lambda(x, t, r)$ is decreasing in r. Consider any initial state x and discount rates $r < r'$. Almost surely, we have $x_t(r) \geq x_t(r')$ for all t and $\tau(x, r) \geq \tau(x, r')$.*

We start the proof with a few observations, fixing for now the discount rate. Starting with $y > x$ and any Markov control, we have $y_t \geq x_t$ for all t, path by path, by continuity of the paths. Moreover, for each t, $V(y_t, t) \geq V(x_t, t)$ whenever $\tau(x) \geq t$, where $V(x, t)$ is denotes the value function at time t when current state is x. Indeed, one can always mimic, when

[7] The assumption on σ guarantees that V is smooth enough to solve almost everywhere the HJB equation used in Lemma 6.

starting from y_t, the control used when starting from x_t, and stop after the same time. Since v is monotonic in x, this yields a higher value.

Let

$$h_s = E\left[v(y_s, \lambda(y_s, s), s)1_{\tau(y)\geq s} - v(x_s, \lambda(x_s, s), s)1_{\tau(x)\geq s}\right], \tag{7}$$

where the expectation is taken with respect to the Wiener measure on the Brownian noise (and, therefore, independent from the initial condition).

LEMMA 4 *For all t, $\int_t^\infty h_s ds \geq 0$.*

Proof. Proceeding as in the proof of Lemma 2, one may show that

$$\int_t^\infty e^{-rs} h_s ds \geq 0 = E\left[V(y_t, t)1_{t\leq\tau(y)} - V(x_t, t)1_{t\leq\tau(x)}\right].$$

Since $V(y_t, t) \geq V(x_t, t) \geq 0$ and $1_{t\leq\tau(y)} \geq 1_{t\leq\tau(x)}$ (the latter because it takes more time to hit the lower boundary when starting from a higher level), the difference inside the expectation is nonnegative almost surely and, therefore, so is the expectation. ■

Using Lemma 3, this time with the function $\alpha(s) = s$, we conclude from Lemma 4 that

$$\int_0^\infty s e^{-rs} h_s ds \geq 0.$$

From the definition of h, this inequality may be rewritten as

$$E\left[\int_0^{\tau(y)} s e^{-rs} v(y_s, \lambda(y_s, s), s) ds\right] \geq E\left[\int_0^{\tau(x)} s e^{-rs} v(x_s, \lambda(x_s, s), s) ds\right] \tag{8}$$

For the rest of the proof we focus on the time-homogeneous case, dropping direct dependence on t for expositional simplicity. The general case is easily obtained from the proof below. Let $V(x, r)$ denote the value function starting with x with discount r.

LEMMA 5 *$V(x, r)$ is submodular in (x, r).*

Proof. By a generalized envelope theorem (see Milgrom and Segal, 2002),

$$V_r(x, r) = \frac{\partial}{\partial r} E\left[\int_0^\tau e^{-rt} v(x_t, \lambda_t) dt\right],$$

evaluated at the optimal controls λ and τ. Computing the derivative explicitly,

$$V_r(x,r) = E\left[\int_0^{\tau} (-t)e^{-rt} v(x_t, \lambda_t)dt\right].$$

This implies that for $y > x$,

$$V_r(y,r) - V_r(x,r) = -E\left[\int_0^{\tau(y)} se^{-rs} v(y_s, \lambda(y_s))ds\right] + E\left[\int_0^{\tau(x)} se^{-rs} v(x_s, \lambda(x_s))ds\right],$$

which is less than zero from (8). ■

LEMMA 6 *$\lambda(x,r)$ is decreasing in r.*

Proof. The HJB equation for the control part of the problem is

$$0 = \sup_{\lambda}\left\{v(x,\lambda) + \mu(x,\lambda)V_x(x,r) + \frac{1}{2}\sigma^2(x)V_{xx}(x,r) - rV(x,r)\right\}.$$

Therefore, $\lambda(x,r)$ maximizes the objective $v(x,\lambda) + \mu(x,\lambda)V_x(x,r)$, which is submodular in (λ, r). Indeed, its cross-derivative with respect to λ and r equals $\mu_\lambda V_{xr}$, which is less than zero since μ is increasing in λ and V is submodular in (x,r) from Lemma 5. Therefore, $\lambda(x,r)$ is decreasing in r for each x. ■

We can now conclude the proof of Theorem 4. Since $\lambda(x,r)$ is decreasing in r for all x, we have path by path, starting from a given x_0, $x_t(r) \geq x_t(r')$ for $r < r'$, where $x_t(r)$ is the path of the optimal state for discount rate r. Indeed, the paths of $x_t(r)$ and $x_t(r')$ cannot cross because, given that they are continuous, so they must first touch, and the moment they touch (i.e. have the same value of the state variable at the same time), Lemma 6 kicks in, guaranteeing that the one with the lower discount rate has the bigger control.[8] We also know from Theorem 3 that the boundary $a(r)$ is increasing in r. Combining the above implies that $x_t(r)$ hits $a(r)$ later than $x_t(r')$ hits $a(r')$, path by path, which concludes the proof.

By an easy modification of Lemma 6, we can extend the result to multidimensional control,

THEOREM 5 *Suppose that λ is K-dimensional, that $\mu(x,\lambda)$ and $v(x,\lambda)$ are supermodular in λ and that μ is increasing in λ. Then $(x,r) \mapsto \lambda(x,r)$ is decreasing in r. Given an initial state x and discount factors $r < r'$, $\tau(r) \geq \tau(r')$ almost surely.*

[8] Continuity plays an important role in this argument. In discrete time, or if jumps were allowed, then paths could cross each other.

Proof. It suffices to modify the proof of Lemma 6. The HJB equation is still

$$0 = \sup_{\lambda} \left\{ v(x,\lambda) + \mu(x,\lambda)V_x(x,r) + \frac{1}{2}\sigma^2(x,t)V_{xx}(x,r) - rV(x,r) \right\}$$

so the vector λ maximizes the objective $v(x,\lambda) + \mu(x,\lambda)V_x(x,r)$, which is supermodular in $(\lambda, -r)$ by assumptions on v and μ and the fact that V is increasing in x and submodular in x, r. The rest of the proof follows as before. ■

3.3 A deterministic counterexample for time-based comparative statics

Suppose that $x_0 = 0$ and that there are two control levels: $\Lambda = \{1, 2\}$. Suppose that utility flow is given by $u(x,\lambda) = M$ for $x \in [1, 10]$, $u(x,\lambda) = -M$ for all $x > 10$ where M is a large positive constant, and $u(x, 1) = 1$ and $u(x, 2) = -0.01$ for $x \in [0, 1)$. Finally, suppose that

$$\frac{dx}{dt} = \lambda_t.$$

In this problem, thus, the state can only go up. Moreover, it is clearly optimal to stop at $x = 10$, and not before, since there is always a control yielding positive utility before that level. Thus, the continuation domain is $C = [0, 10]$ for all discount functions. Finally, it is optimal to spend as much time as possible in the region with payout rate M, i.e. set $\lambda(x) = 1$ for $x \in [1, 10]$. The only question, therefore, is how fast to get to $x = 1$. A very impatient decision maker will never use the control $\lambda = 2$ on the domain $[0, 1)$, because that control yields negative instantaneous utility, the only utility a very impatient decision maker really cares about. By contrast, a patient decision maker puts more value on future cash flows, in particular the ones received once x_t reaches 1. The small negative cash flow 0.01 is only incurred for a short time, and brings high cash flows.

We now show this intuition formally. The HJB equation for this problem, for $x \in [0, 1]$, is, assuming a constant discount rate r,

$$0 = \max\{-rV(x) + 1 + V'(x); -rV(x) - 0.01 + 2V'(x)\},$$

where the first term corresponds to the control $\lambda = 1$ and the second term to $\lambda = 2$. Thus it is optimal to choose $\lambda = 2$ if and only if

$$V'(x) \geq 1.01.$$

We will check that for r small enough, the solution to the HJB equation is maximized by the second term, which corresponds to control $\lambda = 2$. The general solution to

$$-rV(x) - 0.01 + 2V'(x) = 0$$

is

$$V(x) = \frac{-0.01}{r} + c \exp\left(\frac{rx}{2}\right). \tag{9}$$

The boundary condition is $V(1) = M(1 - \exp -rT)/r$, where T is the time it takes for x to go from 1 to 10 under control $\lambda = 1$ (i.e., $T = 9$). This implies that

$$c = \frac{e^{-r/2}}{r}\left[M\left((1 - e^{-rT}\right) + 0.01\right], \tag{10}$$

and hence that

$$V'(x) = \frac{\exp\left(r(x-1)/2\right)}{2}\left[M\left(1 - e^{-rT}\right) + 0.01\right] \geq \frac{\exp\left(-r/2\right)}{2}\left[M\left(1 - e^{-rT}\right) + 0.01\right], \tag{11}$$

which is uniformly, arbitrarily large for M large and r fixed. This shows that the function V defined by 9, with c defined in 10, and the control $\lambda = 2$, solve the HJB equation, and hence that setting $\lambda = 2$ is optimal for patient enough decision makers, which precipitates the time of interruption, compared to an impatient decision maker.[9]

4 Salvage Value

So far we have assumed that the value received upon stopping is zero. Suppose instead that the decision maker receives a lump-sum G_t upon stopping,

$$\text{maximize } E\left[\int_0^{\hat{\tau}} \alpha(s)u(x(s), s)ds + \alpha(\tau)G_\tau\right] \text{ subject to } \hat{\tau} \in \mathcal{T}, \tag{12}$$

where G_t is an adapted process taking value in $\mathbb{R}$.

One may extend Theorem 1 as follows. Consider the stopping time τ that is optimal for discount function α. By definition, this means that the value function $W_t(\alpha)$ at any time $t < \tau$ is greater than or equal to G_t. As we have seen in Section 2, this implies that the

[9] From straightforward inspection of the righthand side of (11), one can show the following, sharper result: for all $M > 1$, there exists $\bar{r}$ such that for all $r < \bar{r}$, it is optimal to set $\lambda = 2$ for $x \leq 1$.

value obtained with stopping τ under discount function β (assuming $\beta(0) = \alpha(0)$) is greater than or equal to G_t as well, and hence that continuing until at least τ does at least as well as stopping immediately.

This shows the following result.

THEOREM 6 (STOPPING WITH SALVAGE VALUE) *Let τ and τ' be solutions to the optimal stopping problem (1) when the discount functions are α and β respectively. If $\beta(s)/\alpha(s)$ is increasing in s and $\alpha(0) = \beta(0)$, then $\tau \vee \tau'$ is also an optimal stopping time for the discount function β.*

Similarly, we can extend several results for the optimal control problem with stopping. Consider the following problem:

$$\text{maximize } E\left[\int_0^{\hat{\tau}} \alpha(s)u(x(s), \lambda(s), s)ds + \alpha(\tau)G_\tau\right] \text{ subject to } \hat{\tau} \in \mathcal{T}, \tag{13}$$

where G_t is the salvage value process and x has controlled dynamics in the sense of Section 3. Theorem 2 extends to this setting. Let $W(\alpha)$ denote the value function of the optimization problem 13 with discount function α.

THEOREM 7 (GENERAL VALUE MONOTONICITY WITH SALVAGE VALUE) *Suppose that $\beta(s)/\alpha(s)$ is increasing in s and that $\alpha(0) = \beta(0)$. Then $W(\beta) \geq W(\alpha)$.*

Proof. Without loss of generality, suppose that $\alpha(0) = \beta(0) = 1$. Consider an optimal pair (τ, λ), assuming it exists,[10] under the discount function α, and let $y_t = (x_t, \lambda_t)$ denote the state-control pair, at any time. From that the same pair (τ, λ) is used under the discount function β. This implies that $W(\beta) \geq E_x[\int_0^\tau \beta(s)u(x(s), \lambda(s), s)ds + \beta(\tau)G_\tau] = \int_0^{\bar{t}} \beta(s)v(s)ds + E[\beta(\tau)G_\tau]$. The result then follows from (5) applied to the state y_t and from the fact that $\beta(\tau) \geq \alpha(\tau)$. ■

Theorem 3 can be extended similarly, where $G_t = G(x_t, t)$ for some function $G : \mathcal{X} \times \mathbb{R}_+ \to \mathbb{R}$. We omit the proof.

THEOREM 8 *Suppose that $\beta(s)/\alpha(s)$ is increasing in s. Then, for all t,*

$$C(t, \alpha) \subset C(t, \beta).$$

[10] Otherwise, an approximation argument could be made.

Finally, we return to time-based comparative statics of Section 3.2, under one assumption on G.

Assumption $G(x,t)$ is increasing in x for all t.

Under this assumption, it is still true that the value function $V(x,t)$ of the problem is increasing in x, replicating the argument of Section 3.2.

Now we modify the function h in (7) to account for the salvage value term. Let

$$h_s^x = E^x\left[v(x_s,\lambda_s,s)1_{s\leq\tau}\right] + f_\tau^x(s)E^x[G(x_s,s)|\tau=s],$$

where the superscript x indicates the initial condition, and where f_τ^x is the density of the stopping time starting from x.[11] and let

$$h_s = h_s^y - h_s^x.$$

With this new definition, it is easy[12] to check that one still has $\int_t^\infty e^{-rs}h_s ds = E[V(y_t,t)1_{\tau(y)>t} - V(x_t,t)1_{\tau(x)>t}] \geq 0$, and that $V_r(y,r) - V_r(x,r) = \int_0^\infty se^{-rs}h_s ds \geq 0$, by another application of Lemma 3. The rest of the proof then follows.

This shows the following extension.

THEOREM 9 *$\lambda(x,t,r)$ and $\tau(x,r)$ are decreasing in r.*

Monotonicity of G seems reasonable when x corresponds to capital which can be sold when a project is interrupted. The theorem is easily extended to a multidimensional control, as in Theorem 5.

[11] Since the stopping is a hitting in our context, and x_t follows a diffusion with a volatility coefficient uniformly bounded away from zero, the distribution of the stopping time does have a density.

[12] Formally, observe that $E^x\left[E_t[e^{-r\tau}G(x_\tau,\tau)]1_{\tau\geq t}\right] = E^x[e^{-r\tau}G(x_\tau,\tau)1_{\tau\geq t}] = \int_0^\infty f_\tau^x(s)e^{-rs}E[G(x_s,s)|\tau = s]1_{s\geq t}ds = \int_t^\infty e^{-rs}f_\tau^x(s)E[G(x_s,s)]ds$. The integral part is the same as before.

5 Applications

5.1 Internal Rate of Return

The internal rate of return of project is the discount rate at which one must discount future cash flows in order to set the present value of the project equal to zero. It is well known that, in many instances, a given cash-flow sequence can give rise to multiple internal rate of returns. Non-uniqueness has concerned eminent economists, including Samuelson (1937) and Arrow and Levhari (1969), as it potentially undermined the decreasing relation postulated by Keynes between aggregate investment and interest rate. Several attempts were made to restore uniqueness by endogenizing the project's life, a feature known as project "truncatability." Arrow and Levhari (1969) showed, using differential methods and an induction on the number of "roots" of the deterministic cash flow function, that the present value of a project is decreasing in the discount rate, implying that the internal rate of return is (essentially) unique. Results from the previous section extend Arrow and Levhari in multiple directions: i) stochastic cash flows (where the induction method fails) and ii) management of the project, and iii) salvage value. This shows a great robustness of the monotonicity studied in their and other papers.

THEOREM 10 *For a given decision problem of the type 1 or 6, let $A \subset \mathbb{R}$ denote the set of discount rates α for which $W(\alpha) = 0$. Then, A is an interval.*

The proof is a direct application of Theorem 7.

5.2 Optimal Growth and Capital Deepening

Consider the following

$$\max_{c,\tau} E\left[\int_0^\tau e^{-rs}u(c_s, k_s, s)ds + e^{-r\tau}G(k_\tau, \tau)\right]$$

subject to

$$dk_t = H(k_t, c_t, t)dt + \sigma(k_t, t)dB_t,$$

where B is the standard Brownian motion, and c is a finite dimensional control. k represents the capital available at any time, and c's components could correspond to consumption (or,

more precisely, the opposite thereof) and labor input, and G is a retirement value (which one may set to an arbitrary negative number if one is not interested in the stopping part of the result). Suppose that H is increasing and supermodular in c and that u is supermodular in c and increasing in k. (Note that in the case where c is scalar, H and u are trivially supermodular in c.) Then one can apply Theorem 9, extended to the multidimensional control, to conclude that k and the stopping time is decreasing in r.

5.3 Bankruptcy Decisions

As an application of Theorem 1, consider the model of endogenous-default setting introduced by Leland (1994) and generalized by Manso et al. (2004). Equity holders of a firm must pay a coupon rate $c(x)$ to debtholders, where c is decreasing in some performance measure x, and receive a payout rate $\delta(x)$, with δ increasing in x.[13] The performance measure $\{x_t\}$ is a time-homogeneous diffusion (for example, geometric Brownian motion, or a mean-reverting) process. The shareholder problem is thus to solve

$$W(x,r) = \sup_{\hat{\tau}\in\mathcal{T}} E_x\left[\int_0^{\hat{\tau}} e^{-rt}(\delta(x_t) - c(x_t))dt\right].$$

Given the time-homogeneous, Markov structure of the problem, and since $\delta - c$ is increasing it is easy to show that optimal default takes the form of a hitting time $\tau_{A_B(r)} = \inf\{t : x_t \leq A_B(r)\}$; $A_B(r)$ is called the *default-triggering level* of the firm, and is independent of the initial asset level x. Theorem 2 says that $A_B(r)$ is increasing in r, and Corollary 1 says that $W(x,r)$ is decreasing in r. We can check this result directly when $\delta(x) = \delta x$, $c(x) = c$, and x is the geometric Brownian motion with drift μ and volatility σ. In this case, standard computations (see, Manso et al. (2009)) show that

$$A_B(r) = \frac{\gamma(r)}{\gamma(r)+1}\left(1 - \frac{\mu}{r}\right)\frac{c}{\delta},$$

[13]For standard debt, c is a constant. However, in many contracts such as performance-pricing loans or step-up bonds, c increases as some performance measure of the firm deteriorates. This measure maybe the credit rating, or directly related to the earnings (EBITDA, price-earning ratio, etc.) of the issuing firm. See Manso, et al. (2009) for examples. The model can easily be modified to account for tax and bankruptcy costs.

where $\gamma(r) = (m + \sqrt{m^2 + 2r\sigma^2})/\sigma^2$ and $m = \mu - \sigma^2/2$. Since A_B increases in γ and r,[14] and $\gamma(r)$ increases in r, necessarily A_B increases in r. In general, A_B cannot be computed explicitly. However, Corollary 1 ensures that monotonicity with respect to the interest rate holds for very general asset processes and coupon and payout profiles.

The result may also be interpreted as connecting household decisions to default on their mortgage with their financial perspective as evaluated with the discount rate. In that light, providing households with a lower interest rate suggests a lower incidence of default decisions.

Finally, the result holds also in the presence of an additional asset management problem and of a salvage value.

[14]More precisely, A_B is the product of two positive factors, each increasing in r.

References

Arrow, K., Levhari, D. (1969) "Uniqueness of the Internal Rate of Return with Variable Life of Investment," *Economic Journal*, Vol. 79, pp. 560-566.

Bolton, P., Harris, C. (1999) "Strategic Experimentation," *Econometrica*, Vol. 67, No. 2, pp. 349–374.

Foster, J., Mitra, T. (2003) "Ranking investment projects," *Economic Theory*, Vol. 22, pp. 469-494.

Keller, G., Rady, S., and Cripps, M. (2005) "Strategic Experimentation with Exponential Bandits," *Econometrica*, Vol. 73, No. 1, pp. 39–68.

Leland, H. (1994) "Corporate Debt Value, Bond Covenants, and Optimal Capital Structure," *Journal of Finance*, Vol. 49, pp. 1213-1252.

Manso, G., Strulovici, B., and Tchistyi, A. (2009) "Performance-Sensitive Debt," forthcoming in the *Review of Financial Studies*.

Milgrom, P., Segal, I. (2002) "Envelope Theorems for Arbitrary Choice Sets," *Econometrica*, Vol. 70, No. 2, pp. 583–601.

Øksendal, B. *Stochastic Differential Equations: An Introduction with Applications*, Fifth Edition. Springer–Verlag, New York, NY.

Quah, J., Strulovici, B. (2009) "Comparative Statics, Informativeness, and the Interval Dominance Order," forthcoming in *Econometrica*.

Ross, S., Spatt, C., and Dybvig, P. (1980) "Present values and internal rates of return," *Journal of Economic Theory*, No. 23, pp. 66-81.

Samuleson, P. (1937) "Some Aspects of the Pure Theory of Capital," *Quarterly Journal of Economics*, pp. 469–496.

Sen, A. (1975) "Minimal conditions for monotonicity of capital value," *Journal of Economic Theory*, No. 11, pp. 340-355.

Strulovici, B. (2009) "Learning While Voting: Determinants of Collective Experimentation," Working Paper, Norhtwestern UIniversity.