Discounting and Patience in Optimal Stopping and Control Problems

John K.-H. Quah Bruno Strulovici*

October 7, 2010

Abstract

The optimal stopping time of any pure stopping problem with nonnegative termination value is increasing in "patience," understood as a partial ordering of discount functions. When utility depends on some Markov state controlled by the agent, the result holds in a "continuation–domain" sense. When utility is increasing in the state, and the state is a one–dimensional diffusion, the result also holds in a "time-domain" sense, and the entire state path is increasing in patience. In all cases, the expected value of the agent is increasing in patience. We extend our concept of patience to study the impact of stochastic discounting and of time–inconsistency on the results and identify, as a by-product, new properties of these environments. For example, we find that a naive agent always experiments longer than a sophisticated one, and introduce an ordering on "partial sophistication." As an application, the internal rate of return of any endogenously-interrupted project is essentially unique, even when cash flows are stochastic and controlled by the entrepreneur. Other applications include growth theory, bankruptcy decisions, environmental protection, and experimentation.

*Email addresses: john.quah@economics.ox.ac.uk and b-strulovici@northwestern.edu. We are grateful for comments from Chiaki Hara, Paul Milgrom, Jean Tirole, and seminar participants at the World Congress of the Econometric Society (2010).

1 Introduction

The relationship between an agent's discounting function, or "patience," and his decisions, plays a major role in many economic problems. A household's decision to buy a house or refinance a mortgage, and when, crucially depends on the interest rate available to the household. Similarly, a firm's decision to declare bankruptcy, and when, depends on how it discounts future cash-flows. In search models, which are becoming ubiquitous in economic analysis, the time at which an agent stops searching crucially depends on how the agent discounts future periods. Similarly, the trade–off between exploration and exploitation, at the heart of experimentation models, is determined by the discount rate of the agents. Another striking illustration is the recent controversy regarding the discount rate that should be used to compute the economic costs and benefits of addressing climate change.[1]

While the impact of discounting on choice is central to the analysis of many economic problems, this impact is unclear even in the most basic settings. The principal concern of this paper is the seemingly simple relationship between the interest rate and the value of a project. Macroeconomic analyzes often postulate a negative relationship between investment and interest rate, but it is known (at least since Samuelson (1937)) that, when the duration of a project is fixed, the value of the project is not necessarily decreasing in the interest rate. As a consequence, an entrepreneur may drop such a project if the interest rate goes *down*. However, when an agent can optimally choose the interruption of a project, Arrow and Levhari (1969) showed, in a deterministic setting, that monotonicity is restored. A similar result was also established by Hicks in his book *Capital and Time* (1973), in which this result is referred to as the *Fundamental Theorem* (of capital). Sen (1975) emphasized that the monotonicity result does not hold when project interruption was followed by a stream of possibly negative "consequence" cash flows. Sen concludes that the negative relation between interest and investment is "not easy to derive from monotonicity results in a microeconomic framework (...) the fact that the monotonicity relation may well be violated once the environmental consequences of investments are considered does raise serious doubts on the negative relation between interest and investment."

[1] In particular, many economists have observed that the conclusions of the Stern Review on the Economics of Climate Change (2006) largely depended on the choice of a very low discount rate. See, e.g., Nordhaus (2007) and Weitzman (2007).

In reality, cash flows in all aforementioned applications are typically stochastic. Moreover, stopping is rarely the only decision faced by an agent. For example, an entrepreneur continually manages a project before deciding to interrupt it. A decision to drastically cut pollution may be preceded by a number of technological improvements to reduce the cost of such cut. Furthermore, agents often face stochastic discounting, either because their future consumption and, hence, intertemporal marginal rate of substitution, is stochastic, or because their access to capital and borrowing rate depends future market conditions. Finally, agents sometimes exhibit time inconsistency: how they perceive cash flow substitution between two given dates depends on the time at which they consider this trade-off. How does the introduction of cash flow uncertainty, the addition of a control problem in addition of an optimal stopping one, stochastic discounting, and time-inconsistency affect the answer to the previous questions? One difficulty for obtaining general results is that standard comparative statics techniques do not apply. First, the optimal stopping time of a cash flow stream can be discontinuous in the discount rate, so that differential techniques do not work. Second, as noted by Quah and Strulovici (2009), the single–crossing property, central to the recent comparative statics literature[2] is typically violated by the value function of an agent facing *any* cash flow stream that, at some point, goes from being positive to being negative. This paper introduces new techniques to analyze comparative statics of dynamic cash flows.

Our analysis starts with an agent who faces a pure stopping problem (without control).[3] We show that, provided that the problem is unrestricted (no exogenous constraints on the stopping time), and that the termination value is nonnegative,[4] the optimal stopping time *always* increases[5] with patience, defined as follows: An agent with discount function $t \mapsto \beta(t)$ (e.g., $\beta(t) = \exp(-rt)$) is *more patient* than an agent with discount function α if the ratio $\beta(t)/\alpha(t)$ is increasing in t. This means that the discount factor decreases at a relatively slower rate under β than under α.[6] More patience, in that sense, results in a later optimal

[2]Milgrom and Shannon (1994).

[3]We first focus on a time–consistent agent with a deterministic discount function. These assumptions are relaxed later. See Sections 5 and 6.

[4]We allow any adapted stochastic termination process, including one with jumps, so as to reflect sudden changes in the environment, either planned or unexpected.

[5]When there exist multiple maximizers, the monotonicity statements hold in the strong set order, or Veinott strong set order (Veinott (1989)), as exposed by Milgrom and Shannon (1994), on the relevant lattices.

[6]When the discount functions have absolutely continuous logarithms, so that $\alpha(t) =$

stopping time. The result holds for any cash flow (or Bernoulli utility) process, including any non–Markovian one. It holds for any agent maximizing expected discounted utility, independently of his (possibly, time–dependent) risk aversion. We also provide a simple counter–example showing why adding a negative termination value can invalidate the result.

When agents can control the cash flows before interrupting them, the problem is more difficult, because agents with varying degrees of patience may apply different controls and, hence, drive an initially common state to widely different levels. Nonetheless, we show the following: first, the value function is *always* increasing with patience, the previous caveat notwithstanding. Again, the result holds for any arbitrary stochastic payoff and nonnegative termination value. Second, when the underlying state has Markov dynamics, the *continuation domain* (i.e., the set of states for which stopping is strictly suboptimal) is increasing with patience. These results must not be confused with the claim that the optimal stopping time increases, because a more patient agent may drive the state faster to the stopping boundary. Indeed, we provide a simple example where more patience precipitates stopping.

The actual duration of a project under control can be analyzed when cash flows (or utilities) depend on a one-dimensional diffusion state. When the payoff (or utility) flow is increasing in the current state, the value function is submodular in the pair (state-discount rate), and the state path and project duration are both 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 hold for a control variable of arbitrary dimension and when stopping entails any positive termination value. We also provide conditions under which the control itself is monotonic in patience, at any given state.

When the agent's discount function is stochastic, more patience (in a suitably extended sense) may result in earlier project interruption. One may build a counter–example exploiting negative correlation between cash–flows and the agent's discount factor. Strikingly, one may even build a counter-example where cash flows are deterministic. To restore a positive result, one needs to assume positive correlation, at each time t, between i) the stochastic discount factor ratio $(\beta(t)/\alpha(t))$, and ii) the discounted cash flow received at time t by the less patient

$\alpha(0)\exp\left(-\int_0^t r_\alpha(s)ds\right)$ for some discount rate function r_α, with a similar expression for β, more patience is equivalent to the requirement that $r_\beta(t) \le r_\alpha(t)$ for all t.

agent (zero, if that agent has already stopped). This assumption is clearly satisfied if both discount functions are deterministic or, more generally, if their ratio is. We provide another set of sufficient conditions on the primitives of the problem when both utility and discount functions depend on some one dimensional state. When the utility satisfies some single-crossing condition and discount functions can be ranked according to some "reference-based" patience ordering, the optimal stopping time is also increasing with patience.

Finally, we also consider the case of a time-inconsistent agent, in the sense of Strotz (1955). Such agent is characterized by a family of discount functions, one for each instant, which determine the time-preference for the individual at each instant.[7] The patience ordering has a natural extension to that context. We first show that, when agents are naive, the more patient agent always stops later, thus extending the first theorem to that case. We then show by example that when agents are sophisticated, the result can be reversed. Intuitively, this will be the case if a patient agent anticipates that his future selves will be even more patient than he currently is, and prefers to stop immediately rather than let these future selves engage in projects with very far-off rewards. However, we show that for the most common models of time-inconsistency, the game played between the different selves of an agent results in a stopping time that arises later for a patient agent than for a less patient one. Finally, we also show that a naive agent always stops later than a sophisticated agent with the same discount function, introduce an intermediate notion of sophistication, and show, more generally, that stopping and the perceived value function decrease with sophistication.

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 positive termination value.[8] (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

[7]See, e.g., O'Donoghue and Rabin (1999) for an exposition of the β-δ model, a well–known discrete-time model of time–inconsistency.

[8]The latter extension was also studied by Flemming and Wright (1971) and Sen (1975).

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) 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. (4) We show that in any multi-armed bandit problem with at least one "safe" arm (i.e., an arm whose payoff distribution is perfectly known) with positive expected cash flow (or utility), a more patient agent experiments longer with risky alternatives.

2 Pure Optimal Stopping Problem

2.1 Nonnegative Termination Value

Let $\{u_t\}$ and $\{G_t\}$ denote real-valued stochastic processes adapted to the filtration $\mathcal{F}$ of some filtered probability space $(\Omega, P, \mathcal{F} = \{\mathcal{F}_t\})$.[9] An agent solves the following optimal stopping problem: choose $\hat{\tau} \in \mathcal{T}$ to

$$\text{maximize}\, U(\hat{\tau};\alpha) = E\left[\int_0^{\hat{\tau}} \alpha(s)u_s ds + \alpha(\hat{\tau})G_{\hat{\tau}}\right] \tag{1}$$

where $\mathcal{T}$ denotes the set of stopping times adapted to $\mathcal{F}$ and taking values in the (possibly infinite) time interval $T = [0, \bar{t}]$,[10] and E denotes the expectation operator. We assume that the *discount function* $\alpha : T \to (0, \infty)$ is deterministic and strictly positive.[11]

While we choose a continuous–time formulation, all the results of the paper, except for those of Section 4, which explicitly require some path continuity, are easily adapted to a discrete-time formulation.

We assume throughout that $E[\int_0^{\bar{t}} \alpha(s)|u_s|ds]$ is finite, that $\alpha(t)$ and $|G_t(\omega)|$ are uniformly bounded (over $t \in T$ and $\omega \in \Omega$), and that the discount function α and the paths of the

[9] For example, utility could take the form $u_t = u(x_t, t)$ where x_t is some underlying Markov process. Such restriction is not required here.

[10] If $\bar{t} = \infty$, $T = [0, \infty)$.

[11] The case of stochastic discount factors is examined in Section 5.

termination value G have bounded variation, and that all processes are right-continuous and have left limits. These assumptions guarantee that all expectations (including conditional expectations) are well-defined, that Fubini's Theorem can be used, and that some (Lebesgue-Stieltjes) integrals with respect to α and G are well defined.[12]

We denote the *value* of problem (1) by $V(\alpha)$, to emphasize its dependence on the discounting function. Formally, $V(\alpha) = U(\tau; \alpha)$ where τ solves (1).

We endow the set of stopping times with the following partial order:

$$\tau \prec \tau' \quad \Leftrightarrow \quad P\left(\{\omega \in \Omega : \tau(\omega) \leq \tau'(\omega)\}\right) = 1.$$

The set of stopping times is a lattice.[13] We also endow the set of discount functions with the following partial order

$$\alpha \prec \beta \quad \Leftrightarrow \quad s \mapsto \tfrac{\beta(s)}{\alpha(s)} \text{ is (weakly) increasing in } s.$$

For instance, suppose that the discount functions are exponential, i.e., $\alpha(s) = \exp(-rs)$ and $\beta(s) = \exp(-\bar{r}s)$, where r and $\bar{r}$ are positive scalars. Then, the ratio $\beta(s)/\alpha(s)$ is increasing in s if $\bar{r} < r$. When the discount functions have absolutely continuous logarithms, so that $\alpha(t) = \alpha(0)\exp\left(\int_0^t -r_\alpha(s)ds\right)$ for some discount rate function r_α, with a similar expression for β, more patience is equivalent to the requirement that $r_\beta(t) \leq r_\alpha(t)$ for all t.

For any lattice L, a correspondence $\alpha \mapsto L(\alpha) \subset L$ is increasing in the *strong set order*[14] if, whenever $\alpha \prec \beta$, $l \in L(\alpha)$ and $l' \in L(\beta)$, we have $l \wedge l' \in L(\alpha)$ and $l \vee l' \in L(\beta)$.

Let $\mathcal{T}(\alpha)$ denote the set of all optimal stopping times when the discount function is α.

THEOREM 1 *Suppose that G is a nonnegative process. Then, $\alpha \mapsto \mathcal{T}(\alpha)$ is increasing in the strong set order. Moreover, $\alpha \prec \beta$ implies that*

$$V(\beta) \geq \frac{\beta(0)}{\alpha(0)} V(\alpha).$$

When the optimal stopping time is unique, Theorem 1 simply says that, path by path, a more patient agent always stops later than a less patient one. The statement in terms of the

[12] These integrals are used to analyze the case of nontrivial termination values. See the Appendix.

[13] Indeed, it is well known and easy to show that the minimum and maximum of two stopping times are also stopping times. See, e.g., Karatzas and Shreve (1991).

[14] See, e.g., Milgrom and Shannon (1994).

strong set order deals with the more complicated case where there are multiple maximizers. In that more general case, Theorem 1 says that if the less patient does not stop, at some given time, there is always an optimal strategy for the more patient agent that entails not stopping at that time. Reciprocally, if it is optimal for the more patient to stop immediately, it is also optimal for the less patient agent to do so.

We now prove Theorem 1 when the termination value is zero. The general case is shown in the Appendix. The proof is based on several lemmas.

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

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

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

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

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

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

Suppose that τ is an optimal stopping time when the discounting function is α. We define the function $v : R_+ \to R$ by

$$v(s) = E[u_s 1_{s<\tau}]; \tag{2}$$

$v(s)$ is the expected payoff rate at time s, where the payoff is zero in the event that one has stopped before s (i.e., $s \geq \tau$). By the optimality of τ, $V(\alpha) = E\left[\int_0^{\tau} \alpha(s) u_s ds\right]$ and Fubini's theorem implies that

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

The next 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,

$$\begin{aligned}
\int_t^{\bar{t}} \alpha(s)v(s)ds &= E\left[\int_t^{\bar{t}} \alpha(s)u_s 1_{s<\tau} ds\right] \\
&= E\left[E\left[\int_t^{\bar{t}} \alpha(s)u_s 1_{s<\tau} 1_{t<\tau} ds|\mathcal{F}_t\right]\right] \\
&= E\left[E\left[\int_t^{\tau} \alpha(s)u_s ds|\mathcal{F}_t\right] 1_{t<\tau}\right].
\end{aligned}$$

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

The next lemma is proved in Quah and Strulovici (2009).

LEMMA 3 *Suppose that γ and h are integrable real-valued functions defined on some compact interval $[x', x'']$ of $\mathbb{R}$, with γ increasing. If $\int_x^{x''} h(s)ds \geq 0$ for all x in $[x', x'']$, then*

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

We now conclude the proof of Theorem 1 for the case of zero termination value. Consider two discount functions $\alpha \prec \beta$, and optimal stopping times $\tau \in \mathcal{T}(\alpha)$ and $\tau' \in \mathcal{T}(\beta)$. We need to show that i) $\tau \vee \tau'$ is optimal for β, and that ii) $\tau \wedge \tau'$ is optimal for α. We first show i). Consider any outcome $\omega \in \Omega$, and $t < \tau(\omega)$. It is enough to show that

$$E\left[\int_t^{\tau} \beta(s)u_s|\mathcal{F}_t\right] \geq 0$$

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

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

We now show ii). Consider $\omega \in \Omega$ such that $\tau'(\omega) < \tau(\omega)$. At time $t = \tau'(\omega)$, the optimal continuation value under β is zero, therefore, it must also equal zero under α, from (5). This implies that $\tau \wedge \tau'$ is optimal for α.

Theorem 1 tell us that *as future utility is discounted less, the optimal horizon gets longer*, independently of the particular stochastic process and payoff function under consideration.

2.2 Negative Termination Value

Here is a simple example showing that a more patient agent can stop earlier than a less patient one when the termination value can take negative values. Consider a deterministic setting where time lies in some finite interval $[0, \bar{t}]$. Suppose that interrupting cash flows at time t results in a termination value of

$$G(t) = -g \exp(\gamma t)$$

for some $g > 0$. For simplicity, we assume that the utility flow is constant, equal to 0.

An agent with constant discount rate $\tilde{r}$ must choose $\tau \in [0, \bar{t}]$. One can think of this problem as procrastinating a chore that worsens with time.

If $r' < \gamma < r$, then the more patient (r') stops immediately, while the less patient stops at the latest time $t = \bar{t}$. Indeed, the present value for stopping at t is

$$-g \exp{(\gamma - \tilde{r})t},$$

which is increasing or decreasing, depending on whether $\tilde{r}$ is less or greater than γ.

3 Optimal Stopping Combined with Optimal Control

There are many situations in which an agent, in addition to the stopping time, 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. In this section we examine the extent to which our earlier results could be extended to this setting.

Consider the following optimization problem: choose $\hat{\tau}$ and $\hat{\lambda}$ to

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

where the processes $\{x_t\}$ and $\{\hat{\lambda}_t\}$ and the stopping time $\hat{\tau}$ are adapted to some filtered probability space $(\Omega, P, \mathcal{F} = \{\mathcal{F}_t\})$ and such that the law of the process $\{x_t\}_{t\geq 0}$ is controlled by $\hat{\lambda}$ 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 $\{\hat{\lambda}_s\}_{s\leq t}$.[15] Later, we will focus on the case where x is a Markov process, but this assumption is not needed for now. We assume that $\{x_t\}$ takes values in some topological space $\mathcal{X}$ and that the process $\{\hat{\lambda}\}_t$ is an *admissible* control, i.e., an adapted process taking values in some set[16] Λ such that the process x is uniquely defined[17] and such that $E\left[\int_0^{\bar{t}} |\alpha(s) u(x(s), \lambda(s), s)| ds\right]$ is finite. We maintain our earlier assumption that the function α and the process G are uniformly bounded and have bounded variation.

Let $V(\alpha)$ denote the value of (6) when α is the discount function.

Theorem 2 (General Value Monotonicity) *Suppose that G is a nonnegative process and that α and β are discount functions for problem (6) such that $\alpha \prec \beta$. Then,*

$$V(\beta) \geq \frac{\beta(0)}{\alpha(0)} V(\alpha).$$

Proof. Consider any pair (τ, λ) that is optimal for the discount function α. Now suppose that the same pair (τ, λ) is used under the discount function β. In that case, we are brought back to the analysis in Section 2, with the flow utility $u_t = u(x_t, \lambda_t, t)$. Equation (5) (or (24), for the case of a positive termination value) shows the result. ∎

It would be tempting to conclude that the optimal stopping time is also increasing with patience. However, that result is generally false. Since optimal controls may differ under

[15] In principle, we could allow the control to be based on partial observation only. In that case, $\hat{\lambda}$ would depend on a subfiltration of $\mathcal{F}$, for example based on some injective function of $y_t = v(x_t)$ (i.e., the underlying state x_t is not fully and directly observable). While such extension would not affect the qualitative arguments of this section, it would be technically more involved, if only to guarantee that the optimal control problem is well defined in such environment. See for example Bensoussan (1992).

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

[17] See Fleming and Soner (1993).

α and β, the state processes may follow different paths, which in turn can lead to non-comparable stopping times. An example where more patience results in an earlier stopping time is provided in Section 4.3.

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 provide a simple condition under which the optimal stopping time *does* increase with patience.

3.1 Domain-Based Comparative Statics for Markov Processes

Suppose now that $\{x_t\}$ is a Markov process: for any t' and control $\{\lambda_s\}_{s\geq t'}$, the distribution of $\{x_s\}_{s\geq t'}$ depends on past history only through $x_{t'}$. Also suppose that the termination value process is of the form $G_t = G(x_t, t)$ for some function G of state and time. In this case, it is well known that the value function depends, at any time t, only on the current state x_t and time t: $V_t = V(t, x_t)$ for some function V.

For any discount function α, define the *continuation domain* at time t by

$$C(\alpha, t) = \{x : V(\alpha, t, x) > G(t, x)\}.$$

$C(\alpha, t)$ is the set of states at which stopping at time t (and, thus, getting termination value $G(t, x)$) is strictly suboptimal. In general, $C(\alpha, t)$ varies with time, but if x has time-homogeneous dynamics and $G(x, t)$ only depends on x, $C(\alpha, t)$ is a constant subset of $\mathcal{X}$. For example, if x is a one-dimensional controlled diffusion, then $C(\alpha, t)$ is typically (though not necessarily) an interval $[a, b]$ such that it is optimal to stop exactly when x_t hits one boundary of that interval.[18] The next result says that *the continuation domain expands with patience.*

THEOREM 3 *Suppose that x is a Markov process, that $G_t = G(x_t, t)$ is nonnegative, and that $\alpha \prec \beta$. Then, for all t,*

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

[18] In Sections 4, and 7.3, the interval is of the form $[a, \infty)$.

Proof. Suppose that $x \in C(\alpha, t)$. This means that $V(\alpha, t, x) > G(t, x)$. Theorem 2, applied at time t, implies that $V(\beta, t, x) \geq V(\alpha, t, x)$, so $x \in C(\beta, t)$. ■

Theorem 3 applies, for example, to experimentation problems where the state x is the belief of the agent about an arm. Theorem 3 tells us that the *experimentation domain*, i.e., the range of beliefs for which experimentation takes place, increases with the decision maker's patience, as detailed in Section 7.4.

4 Time-Based Comparative Statics for One-Dimensional Diffusions

4.1 Monotonicity of State Path and Stopping Time

Consider the following control problem:[19] choose $\hat{\tau}$ and $\hat{\lambda}$ to

$$\text{maximize}\, U(\hat{\tau}, \hat{\lambda}, ; \alpha) = E\left[\int_0^{\hat{\tau}} e^{-rs} u(x(s), \hat{\lambda}(s), s) ds + \alpha(\hat{\tau}) G(x_{\hat{\tau}}, \hat{\tau})\right], \tag{7}$$

subject to

$$dx_t = \mu(x_t, \hat{\lambda}_t, t)dt + \sigma(x_t, t)dB_t \qquad x_0 = x,$$

where $\{x_t\}$ is one dimensional and B is the standard Brownian motion. We assume that σ is bounded below away from zero and that the value function is well defined and smooth[20] to emphasize its dependence on the initial state x and on the discount rate r, we denote the value by $V(x, r)$.

For expositional simplicity, we focus on comparing exponential discount functions with fixed discount rates. The results of this section hold more generally, as detailed in Appendix 10.3.

We make the following assumption.

[19] 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.

[20] The assumption on σ guarantees that the value function is smooth enough to solve almost everywhere the HJB equation used in Lemma 6. Smoothness of V guarantees that it solves (uniquely) the HJB equation used in Lemma 6. See, e.g., Fleming and Soner (1993).

Assumption 1 (Monotonicity) *u and G are increasing in x for all λ, t.*

This assumption on u and G guarantees that the value function $V(x,r)$ is increasing in x. Indeed, starting from some state y, with $y > x$, one can always replicate the control applied when starting from x and get a higher payoff flow at all times before and upon stopping. More generally, letting $V(x,t,r)$ denote the value function at time t when the current state is x, $V(x,t,r)$ is nondecreasing in x. Given the Markov assumption, this implies that there exists a process $a(t)$ such that the optimal stopping time is given by $\tau = \inf\{x_t \notin (a_t, \infty)\}$.[21] The assumed smoothness of the value function implies that it must satisfy the following Hamilton-Jacobi-Bellman (HJB) equation:[22]

$$0 = \sup_{\lambda} \left\{ u(x,\lambda,t) + \mu(x,\lambda,t)V_x(x,t,r) + V_t(x,t,r) + \frac{1}{2}\sigma^2(x)V_{xx}(x,t,r) - rV(x,t,r) \right\}. \tag{8}$$

We first make a change of variable, replacing the initial control variable by a direct determination of the drift. Let $\upsilon(x,m,t) = \sup\{u(x,\lambda,t) : \mu(x,\lambda,t) = m\}$, with $\upsilon(x,m,t) = -\infty$ if the defining set is empty. Thus, $\upsilon(x,m,t)$ is the maximal utility one can achieve at (x,t) while providing a drift m to the state. We assume that this maximum is achieved whenever υ is finite. The optimal control problem can be reexpressed in terms of m:

$$0 = \sup_{m} \left\{ \upsilon(x,m,t) + mV_x(x,t,r) + V_t(x,t,r) + \frac{1}{2}\sigma^2(x)V_{xx}(x,t,r) - rV(x,t,r) \right\}. \tag{9}$$

Any drift-control process $\{m_t\}$ that achieves at all times the supremum in (9) is an optimal control. We assume that (9) has at least one maximizer for all (x,t). Reciprocally, a control process that fails to achieve the supremum in (9) on a set of positive measure is suboptimal.[23] Let $M(x,t,r)$ denote the set of maximizers of (9) for state x, time t, and discount r. Fixing the initial condition x, let $X(x,r) \subset \mathcal{C}([0,\bar{t}],\mathbb{R})$ denote[24] the set of paths obtained starting from x when using an admissible selection from M: at each time t, the drift satisfies $m_t \in M(x_t,t,r)$. Thus, for each realization ω of the underlying Brownian motion, $X(x,r)$ contains all trajectories of the state obtained from using an optimal control at each

[21] Indeed, recall that the continuation domain at time t consists of the states x such that $V(x,t,r) > G(x,t)$.

[22] The HJB equation is based on a model where the discount function, from any time s onwards, is given by $\alpha(s,t) = \exp(-r(t-s))$, for $t \geq s$. The factor $\exp(rs)$ does not affect the agent's decisions, compared to a model where the discount function equals $\alpha(s,t) = \exp(-rt)$. See Section 6 for the case of time inconsistency.

[23] Precisely, one may show that any HJB-maximizing control strictly dominates any control that does not solve (9) for a time set of positive measure.

[24] $\mathcal{C}([0,\bar{t}],\mathbb{R})$ denotes the set of real-valued continuous functions defined on $[0,\bar{t}]$ (or $[0,\infty)$, if $\bar{t} = +\infty$).

time. When (9) has everywhere a unique maximizer, $X(x,r)$ simply contains the optimal trajectory of the state. We endow the set of real-valued processes with time domain $[0,\bar{t}]$ with the partial order

$$x \prec y \Leftrightarrow Pr\left(\{\omega : x_t(\omega) \leq y_t(\omega) \text{ for all } t \leq \bar{t}\}\right) = 1.$$

Finally, let $\mathcal{T}(x,r)$ denote the set of optimal stopping times at discount rate r. The set of stopping times is endowed with the same partial order as in Section 2.

Theorem 4 *Suppose that $G(x,t)$ is nonnegative for all x and t. Then, $M(x,t,r)$, $X(x,r)$ and $\mathcal{T}(x,r)$ are decreasing in r in the strong set order.*

When the optimizers are unique, Theorem 3 means that i) the optimal drift, ii) the state path, and iii) the optimal stopping time are all increasing in patience, in their respective domain. As for Theorem 1, the strong set order statement deals with the more general case where there exist multiple optimizers. For instance, for any trajectory of the state optimally generated by a less patient agent, there exists an optimal trajectory for the more patient agent that is everywhere above the first trajectory.

We show the result for the case of a zero termination value. The general case is shown in the Appendix (Section 10.2). Suppose that $(x,t) \mapsto \lambda(x,t)$ is a maximizer of (8), and, hence defines an optimal control. Let

$$h_s = E\left[u(y_s, \lambda(y_s,s), s)1_{\tau(y)\geq s} - u(x_s, \lambda(x_s,s), s)1_{\tau(x)\geq s}\right], \tag{10}$$

where the expectation is taken with respect to the Wiener measure on the Brownian noise.

Lemma 4 *For all t, $\int_t^\infty e^{-rs} h_s ds \geq 0$.*

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

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

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 $\gamma(s) = s$, we conclude from Lemma 4 that

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

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

$$E\left[\int_0^{\tau(y)} se^{-rs}u(y_s, \lambda(y_s, s), s)ds\right] \geq E\left[\int_0^{\tau(x)} se^{-rs}u(x_s, \lambda(x_s, s), s)ds\right] \tag{11}$$

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),[25]

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

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 (11). ■

LEMMA 6 *$M(x, r)$ is decreasing in r in the strong set order.*

Proof. The drift-control HJB equation for the problem (simplified for the time–homogeneous case) is

$$0 = \sup_m \left\{v(x, m) + mV_x(x, r) + \frac{1}{2}\sigma^2(x)V_{xx}(x, r) - rV(x, r)\right\}.$$

[25]The parameter r does not affect the distribution of the underlying process, which justifies the application of the theorem.

The objective is submodular in (m, r). Indeed, its cross-derivative with respect to m and r equals V_{xr}, which is less than zero, since V is submodular in (x, r) from Lemma 5. Therefore, the set of maximizers, $M(x, r)$, is decreasing in r in the strong set order, for each x (see Topkis (1978) and Milgrom and Shannon (1994)). ■

We can now conclude the proof of Theorem 4 for the case of zero termination value.

Proof. Given some initial condition x, let x'_t denote any optimal trajectory for rate r', and let μ'_t denote the corresponding drift control. For $r < r'$, Lemma 6 implies that there exists a selection $\bar{\mu}(x)$ of $M(x, r)$ for all x such that $\bar{\mu}(x'_t) \geq \mu'_t$ for all t. Now consider any optimal control μ_t for discount rate r that coincides with $\bar{\mu}(x'_t)$ whenever the resulting trajectory x is such that $x_t = x'_t$. By construction, the paths of x_t and x'_t cannot cross because, given that they are continuous, they must first touch, and the moment they touch (i.e. have the same value of the state variable at the same time), the assumption of the drifts guarantees that the path of the lower discount rate receives a higher drift.[26] We 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, so $\tau(x, r) \geq \tau(x, r')$. ■

4.2 Control Monotonicity

We now show that, when the drift is increasing in the control, then the control is increasing in patience for any state. Suppose that λ is one-dimensional and that the drift function μ is increasing in λ. Then, proceeding as in Lemma 6 but using the original HJB equation 8 instead of 9, one can show that the set $\Lambda(x, r)$ of maximizers of the equation is decreasing in r. Indeed, the objective jointly depends on λ and r through $u(x, \lambda) + V_x(x, r)\mu(x, \lambda)$, which is submodular in r and λ. In fact, a similar argument holds if λ is multidimensional and the functions u are μ is supermodular in λ.

THEOREM 5 *Suppose that μ is increasing in λ and that u and μ are supermodular in λ (which always holds if λ is one dimensional). Then, $\Lambda(x, r)$ is decreasing in r, in the strong set order.*

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

4.3 Deterministic Counterexample for Time-Based Comparative Statics

We provide an example of a problem where the more patient agent stops earlier.

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 agent will never use the control $\lambda(x) = 2$ for x sufficiently close to 0, because that control yields negative instantaneous utility, while the reward (M) is 'far away', arriving only when $x_t = 1$. By contrast, a patient agent puts more value on future cash flows; from his perspective, the small negative cash flow 0.01 is only incurred for a short time, and leads to high cash flows. Thus it seems possible that for sufficiently high values of M, the more patient agent will choose $\lambda(x) = 2$ for *all* $x \in [0, 1)$ and therefore stop sooner than a less patient agent.

We now confirm 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

$$V'(x) \geq 1.01. \tag{12}$$

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

$$V(x) = \frac{-0.01}{r} + c \exp\left(\frac{rx}{2}\right). \tag{13}$$

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 $\lambda = 1$ (so $T = 9$). This implies that

$$c = \frac{e^{-r/2}}{r} \left[M\left((1 - e^{-rT}\right) + 0.01\right], \tag{14}$$

and hence that

$$V'(x) = \frac{\exp(r(x-1)/2)}{2} \left[M\left(1 - e^{-rT}\right) + 0.01\right] \geq \frac{\exp(-r/2)}{2} \left[M\left(1 - e^{-rT}\right) + 0.01\right], \tag{15}$$

which is uniformly, arbitrarily large for M large and r fixed. In particular, at a given discount rate $\tilde{r}$, there is $\tilde{M}$ such that (12) is satisfied. This shows that the function V defined by (13), with c defined in (14), and the control $\lambda = 2$ solve the HJB equation, and hence setting $\lambda(x) = 2$ for all $x \in [0, 1)$ is optimal for this agent.[27]

On the other hand, for a fixed $\hat{x} \in (0, 1)$ and with $M = \tilde{M}$, it is clear from (15) that (12) will be violated if r is sufficiently high above $\tilde{r}$. Such an agent's optimal control will involve choosing $\lambda(x) = 1$ for at least some $x \in [0, 1)$; consequently, she will stop *later* than the agent with discount rate r.

5 Stochastic Discount Factor

Dynamic stochastic models in macroeconomics and finance typically involve stochastic discounting, where the stochastic discount is a "state price deflator," representing some (random) intertemporal rate of substitution between (in discrete time) consecutive periods.[28]

Another, related, reason for considering stochastic discounting is to think of discounting as the borrowing rate of a firm. The firm does not know today what borrowing rate will be available in a year from now. If there is procyclicality between the interest rate and the

[27] From straightforward inspection of the righthand side of (15), 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 all $x \leq 1$.

[28] See, e.g., Duffie (2001).

firm profit (or if on the contrary, its profits are countercyclical with respect to the general borrowing rate, or lending rate), the discounting rate and cash flows are correlated.

Discounting may also be stochastic because the borrowing rate of the firm depends on its credit rating. In this case, there is an obvious connection between cash flows and discounting. Performance pricing loans and step up bonds are examples of debt contracts where the interest rate depends on some performance measure of the borrower.

This section contains two examples showing that even in simple settings, stochastic discounting can result in the more patient agent stopping earlier, and two theorems providing general conditions under which this is ruled out. The first theorem is based on some positive correlation condition. The second theorem assumes a single crossing condition on the utility function, combined with some reference-based notion of patience. A natural extension of the patience order on stochastic discount function is that, path by path, the ratio process $\beta(s)/\alpha(s)$ is increasing in s. The first counter example exploits correlation between that ratio and the cash-flows. Striking, however, the second counter example is based on *deterministic* cash flows. Intuitively, stochastic discounting causes discounted future cash flows to behave as if cash flows were also stochastic. From the agent's viewpoint, the two sources of uncertainty play a similar role on his present value. For simplicity, the counter examples are built in discrete time.

5.1 Counter Example I: Discount-Cash Flow Correlation

Consider a risky investment over three periods $\{0, 1, 2\}$, with two states of the world, "high" and "low," which are equally likely. The cash-flows in Period 0 and 1 are -1 and -2, respectively. The cash-flow in Period 2 is 21 if the state is high, and -4 if the state is low.

Suppose that agent A has the deterministic discount function $\alpha(t) = \left(\frac{1}{2}\right)^t$. Agent A strictly prefers to make the investment, since he gets a present value of $-1 + \frac{1}{2}(-2) + \frac{1}{4}\left(\frac{1}{2}(21) + \frac{1}{2}(-4)\right) > 0$.

Now consider another agent, B, who has the stochastic discount function $\beta(0) = 1$ and $\beta(1) = \frac{1}{2}$, and $\beta(2) = \frac{1}{4}$ if the state is high, and $\beta(2) = \frac{1}{2}$ if the state is low.

At each time $t \in \{0, 1\}$, B knows only his discount rate between the current period and the next one. In particular, B does not know $\beta(2)$ when he decides, at time 0, whether to invest in the project.

Agent B is more patient than agent A: the ratio β/α is increasing over time for each realization of uncertainty, and strictly so if the state is low. Moreover, B strictly prefers not to invest, because the present value of the project, using his stochastic discount function, is $-1 + \frac{1}{2}(-2) + \frac{1}{4}\frac{1}{2}(21) + \frac{1}{2}\frac{1}{2}(-4)$. This shows that the more patient agent stops earlier.

5.2 Counter Example II: Deterministic Cash Flows

There are four periods $\{0, 1, 2, 3\}$, with corresponding cash flows -1, -1, -4, and 2.1. There are two states of the world, which are equally likely. For computational simplicity, we choose the following discount functions:[29]

In State I, the discount functions are given by $\alpha(0) = \beta(0) = \alpha(1) = \beta(1) = \alpha(2) = \beta(2) = 1$ and $\alpha(3) = \beta(3) = 5$.

In State II, the discount functions are given by $\alpha(0) = \beta(0) = \alpha(1) = \beta(1) = \alpha(2) = \alpha(3) = 1$, and $\beta(2) = \beta(3) = 3$.

Thus, agents differ in patience only in State II, and in that case, β/α strictly increases between Periods 1 and 2. For each state, β/α is weakly increasing in time, so that agent B is more patient that A, and strictly so in State II.

At each time $t \in \{0, 1, 2\}$, agents only know their discount rate between the current period and the next period. In particular, they do not know which state (I or II) has realized until period 1.

Agent A strictly prefers to go all the way to Period 3, whereas Agent B strictly prefers to stop immediately. Indeed,

$$PV^I = -1 \times 1 \quad - 1 \times 1 \quad - 4 \times 1 \quad + 2.1 \times 5 \quad = 4.5,$$

conditional on State I. In State II, both agents have a negative present value for the cash

[29]It is easy to modify the example so that discount functions be decreasing in time and/or less than one, but that makes the computations less obvious.

flows:

$$PV^{II,A} = -1 \times 1 \quad - 1 \times 1 \quad - 4 \times 1 \quad + 2.1 \times 1 \quad = -3.9,$$

and

$$PV^{II,B} = -1 \times 1 \quad - 1 \times 1 \quad - 4 \times 3 \quad + 2.1 \times 3 \quad = -7.7,$$

conditional State II. Aggregating across states, Agent A has a slight but strict preference in favor of the project. However, Agent B values future cash flows more precisely when the present value of these future cash flows is negative, because the discount function is higher in that case. Overall, thus, the value of undertaking the project is negative for him.

5.3 Positive Results with Stochastic Discounting

A stochastic discount function is a positive process defined on $[0, \bar{t}]$. Given two discount functions α and β, say that β exhibits more patience than α, still denoted $\alpha \prec \beta$,

$$Pr\left(\left\{\omega : \frac{\beta(s)}{\alpha(s)} \text{ is increasing on } [0, \bar{t}]\right\}\right) = 1.$$

An agent solves the following optimal stopping problem (using the notation of Section 2):

$$V(\alpha) = \sup_{\hat{\tau} \in \mathcal{T}} E\left[\int_0^{\hat{\tau}} \alpha(s) u_s.\right] \tag{16}$$

The only difference with Section 2 is that α is now stochastic. For simplicity we also focus the extension on the case without termination value. We impose otherwise the same integrability and other technical conditions as in Section 2.

The positive result requires the following assumption.

Assumption 2 (Positive Correlation) *Let τ denote any optimal stopping time for α. For all $s \in [0, \bar{t}]$,*

$$Corr\left(\frac{\beta(s)}{\alpha(s)}, \alpha(s) u_s 1_{s<\tau}\right) \geq 0. \tag{17}$$

Theorem 6 *Suppose that $\alpha \prec \beta$ and that Assumption 2 holds. Then, $\mathcal{T}(\alpha) \prec \mathcal{T}(\beta)$ in the strong set order. Moreover,*

$$V(\beta) \geq \frac{\beta(0)}{\alpha(0)} V(\alpha).$$

Intuitively, Assumption 2 means that the Agent B (β discount) is more likely to value future cash flows more relative to Agent A (α discount) when i) these cash-flows are higher, ii) A values these cash flows more, and iii) A is less likely to have stopped already.

Proof. Without loss of generality, it suffices to show that if it is strictly optimal for A to continue beyond $t = 0$, then it is also strictly optimal for B to do so. Optimality of τ implies that $E[V_t 1_{t<\tau}] \geq 0$ and, hence,[30] that

$$\int_t^{\bar{t}} E[\alpha(s) u_s 1_{s<\tau}] ds \geq 0$$

for all $t \leq \bar{t}$. Since Agent B can always choose the stopping time τ, we have (using, again, the same manipulation as in the proof of Theorem 1),

$$V(\beta) \geq \int_0^\infty E\left[\frac{\beta(s)}{\alpha(s)} \alpha(s) u_s 1_{s<\tau}\right] ds \geq \int_0^\infty E\left[\frac{\beta(s)}{\alpha(s)}\right] \times E[\alpha(s) u_s 1_{s<\tau}] ds,$$

where the inequality comes from Assumption 2. The function $E[\beta(s)/\alpha(s)]$ is increasing in s, since $\beta(s)/\alpha(s)$ is increasing path by path. Applying Lemma 3, we conclude that

$$V(\beta) \geq \frac{\beta(0)}{\alpha(0)} V(\alpha).$$

This shows that if $V(\alpha)$ is strictly positive, then so is $V(\beta)$ and, hence, that it is also strictly suboptimal for B to stop immediately. The rest of the proof follows that of Theorem 1. ■

The proof only requires that $E[\beta/\alpha]$ be increasing, which is a weaker requirement than pathwise monotonicity of the ratio. Assumption 2 involves the optimal stopping time τ used for discount α. The next result relies on assumptions directly expressed on primitives. Patience is compared, across discount functions, with respect to some reference point of the utility function.

Suppose that $u_s = u(x_s, s)$ where x is some one-dimensional state, and that, for each s, $u(\cdot, s)$ has the single crossing property: there exists $\bar{x}_s$ such that $u(x, s) \geq (\leq) 0$ for $x \geq (\leq) \bar{x}_s$. Also suppose that the discount factors are stochastic only through the state x. That is,

$$\alpha_s = \alpha(x_s, s)$$

for some function α, with a similar expression for β. Such specification where payoffs and the discount factor depend on some common underlying process arises, for instance, in affine term structure models.[31]

[30] See Lemma 2.

[31] See Duffie (2001, Chapter 7) for a discussion of such models.

We modify the definition of patience as follows:

Assumption 3 (Reference-Based Patience) *Say that β exhibits more reference-based patience than α, denoted $\alpha \prec^{ref} \beta$ if the two following conditions hold:*

1. $\frac{\beta(x,s)}{\alpha(x,s)}$ *is increasing in x, for each $s \in T$.*
2. $\frac{\beta(\bar{x}_s,s)}{\alpha(\bar{x}_s,s)}$ *is increasing in s.*

When the reference point $\bar{x}_s$ is independent of s, the second condition is identical to the patience order introduced in Section 2. In general, that condition applies the patience ordering by evaluating players' discounting at time s at the pivotal (or reference) state level $\bar{x}_s$. The first condition is akin to a monotone likelihood ratio property and comes in pair with the single crossing property to obtain comparative statics.

Theorem 7 *Suppose that $\alpha \prec^{ref} \beta$. Then, $\mathcal{T}(\alpha) \prec \mathcal{T}(\beta)$ in the strong set order. Moreover,*

$$V(\beta) \geq \frac{\beta(0)}{\alpha(0)} V(\alpha).$$

Proof. The proof is identical to that of Theorem 6, except for the following. Single-crossing of the utility function u and monotonicity of the ratio β/α with respect to x imply that

$$E\left[\frac{\beta(x_s,s)}{\alpha(x_s,s)}\alpha(x_s,s)u(x_s,s)1_{s<\tau}\right] \geq \frac{\beta(\bar{x}_s,s)}{\alpha(\bar{x}_s,s)} \times E[\alpha(x_s,s)u(x_s,s)1_{s<\tau}]$$

The second part of Assumption 3 then implies that we can apply Lemma 3 as in the proof of Theorem 6, and reach the same conclusion.

6 Time Inconsistency

How are the results affected if the agent is time-inconsistent? Following the literature on time-inconsistency, starting with Strotz (1955), we consider both "naive" and "sophisticated" agents. A time-inconsistent agent is characterized by a discount function $\alpha(s,t)$ for $s \leq t$ such that, at time s, his time-preference for future cash–flows are determined by the discount function $\alpha(s,\cdot)$. An agent is time consistent if $\alpha(s,t)$ is independent of s.[32] Time–inconsistency gives rise to a game between the selves of an agent at different times, who have misaligned preferences over cash flow streams. A naive agent completely ignores the fact that he is time inconsistent, while a sophisticated agent fully takes this problem into account.

Definition *B is more patient than A, denoted $\alpha \prec^{inc} \beta$, if at each time t, the ratio*

$$\frac{\beta(s,t)}{\alpha(s,t)}$$

is increasing in t, for $t \geq s$.

In this section, we first show that the results of Section 2 extend to the case of naive agents. We then provide an example showing that when agents are sophisticated, the more patient agent can stop earlier than the less patient one. We show that, when time inconsistency takes the form of *decreasing patience*, a condition satisfied by hyperbolic and quasi-hyperbolic discounting, the equilibrium stopping time of a sophisticated agent increases with patience. Furthermore, we show that a naive agents always stop later than a sophisticated one with the same discount function. We also introduce a notion of partial sophistication, and show that the equilibrium stopping time is decreasing in the agent's degree sophistication, provided that the discount function exhibits decreasing patience. This result can also be interpreted in the context of an agency with successive agents who take control in turn over the stopping decision and are decreasingly patient: more changes of control result in earlier stopping.

For simplicity, we focus on the pure-stopping setting of Section 2.

6.1 Naive Agents

As in Section 2, let $\mathcal{T}(\alpha)$ denote denote the set of stopping times such that

$$\tau \in \mathcal{T}(\alpha) \Leftrightarrow V_\tau(\alpha(\tau,\cdot)) = 0,$$

where, for any $s \in T$ and discount function $\beta : T \to \mathbb{R}_{++}$, $V_s(\beta)$ is the value function, at time t, of a time-consistent agent with discount function β. A naive agent, at time t, believes that

[32]More precisely, if the ratio $\alpha(s,t)/\alpha(s',t)$ is independent of t, for any s, s'.

his future selves all have the discount function $\alpha(t, \cdot)$. Therefore, he believes that stopping is optimal if and only if $V_s(\alpha(s, \cdot) = 0$. $\mathcal{T}(\alpha)$ is thus the set of all "optimal" stopping times, seen from a naive agent's viewpoint.

THEOREM 8 *Suppose that A and B are naive agents and that B is more patient than A. Then, $\alpha \mapsto \mathcal{T}(\alpha)$ is increasing in the strong set order. Moreover, $\alpha \prec^{inc} \beta$ implies that*

$$V(\beta) \geq \frac{\beta(0)}{\alpha(0)} V(\alpha).$$

Proof. Let $\tau \in \mathcal{T}(\alpha)$. For any $s < \tau$, $V_s(\alpha(s, \cdot)) > 0$. That value function is exactly the same as the one computed in Section 2. Since $\alpha(s, \cdot) \prec \beta(s, \cdot)$, Theorem 1 applied to time s and discount functions $\beta(s, \cdot)$ and $\alpha(s, \cdot)$ implies that

$$V_s(\beta(s, \cdot)) \geq \frac{\beta(s, s)}{\alpha(s, s)} V_s(\alpha(s, \cdot)).$$

Therefore, B finds it strictly suboptimal to stop at time t whenever A does. The rest of the proof is identical to that of Theorem 1. ∎

6.2 Sophisticated Agents: A Counter Example

We first provide an example, with sophisticated agents, where the more patient agent stops earlier. In that case, the multiple selves (one for each time) of an agent play an equilibrium where, at each time, exactly one of the selves makes a single move. For a formal definition of such equilibria, see Section 6.3. For the example, an informal description suffices. For simplicity, we consider a discrete-time counter example.

Consider a deterministic stream of cash flows given by the sequence

$$1, -M, \frac{M}{n}, \frac{M}{n}, \ldots, \frac{M}{n}, 0, 0, \ldots,$$

where $M = 10$, say, n is a (large) integer, and there are exactly $n+1$ periods where the cash flow equals $\frac{M}{n}$. An agent must decide when to interrupt this cash flow.

Suppose that the impatient agent, A, is time–consistent with a discount rate of 100% (i.e., the discount factor between two periods is $1/(1+1) = 0.5$, and the discount function satisfies

$\alpha(t) = 0.5^t$). Then, A takes the first cash flow (1), and stops before getting $-M$, because the present value of $-M$ followed by the (arbitrarily) small amounts M/n is negative.

Suppose that the patient agent, B has the same discount function as A at period 0 (i.e., a constant discount rate of 100%), but becomes extremely patient after period 1. For simplicity, suppose that the discount function of the agent after period 1 is flat.[33] Formally, B has the discount function $\beta(0,t) = 0.5^t$ ($t \geq 0$) at time 0 and the discount function $\beta(s,t) = 1$ ($t \geq s$) at any time $s > 0$.

If B does not stop immediately, his Period-1 self will want to incur the loss $-M$ in order to get future total benefits close to $(n+1)M/n > M$. For the Period 1 self, who is perfectly patient, it is indeed optimal to go through the entire stream of cash flows.

However, seen from the Period-0 B's perspective, the stream of cash flow $-M, M/n, ...$ has a large negative present value, which more than offsets the benefit of the positive Period 0 cash flow. Therefore, B prefers to forgo that immediate positive cash flow by stopping immediately, rather than have his future selves go through the entire stream of cash flows.

Therefore, the more patient agent, B, stops at Period 0, while the less patient, A, stops at Period 1.

6.3 Sophisticated Agents with Decreasing Patience

Time is discrete, with final period $\bar{t}$. Cash flows $\{u(t)\}_{t \leq \bar{t}}$ follow any arbitrary (possibly non Markovian) process with finite unconditional and conditional expectations. We focus on pure-strategy equilibria: at each time s, the s self of the agent optimally decides whether or not to interrupt the cash flows, if no previous self has already done so, given the strategy profile of his future selves. Formally, let $T = \{0, 1, \ldots, \bar{t}\}$ denote the set of periods, and $\mathcal{H}^t$ denote the set of cash-flow histories from time 0 to time t excluded.[34] The set of outcomes is then given by $\Omega = \mathcal{H}^{\bar{t}+1}$. A pure-strategy equilibrium is a binary map $e : \{(t,h) : t \in T, h \in \mathcal{H}^t\} \mapsto \{C, S\}$ that determines, for each t and $h \in \mathcal{H}^t$, whether the t-self of the

[33] Since cash flows are equal to zero after finitely many periods, such assumption is not problematic.

[34] By convention, we assume that at time t, the agent must decide whether to continue before observing the time t cash flow. The opposite convention, where current-period uncertainty is resolved before the agent's decision to continue, does not affect the analysis.

agent continues (C) or stops (S), given that previous selves have continued. Optimality requires that $e(t,h) = C$ only if the value function conditional on continuing is positive. (For simplicity, we rule out here any termination value, so that stopping yields zero.) We also assume, to pin down the equilibrium, that any indifferent agent stops. With these assumptions, the equilibrium is unique and can be computed by backward induction.

For $s \leq t$, let $\alpha(s,t)$ denote the discount factor of the s-self at time t. Let $\mathcal{S}(\alpha)$ denote the set of pairs (t,h), with $h \in \mathcal{H}^t$, at which the agent stops, in equilibrium. This set is completely determined by backward induction, and is independent of play history up to time t (excluded): whether the t-self decides to continue or stop only depends on the behavior of his future selves, not of his past selves. If $\omega \in \Omega$ denotes the realized sequence of cash flows, let $\omega^t \in \mathcal{H}^t$ denotes the truncation of ω up to time t. The equilibrium stopping time τ^α of an agent with discount function α is given by

$$\tau^\alpha(\omega) = \inf\{t : (t, \omega^t) \in \mathcal{S}(\alpha)\}.$$

Let also τ_t^α denote the stopping time of the agent restricted to $\tau_t^\alpha \geq t$, i.e., when all selves before time t are forced to continue. From the previous observation, we have

$$\tau_t^\alpha(\omega) = \inf\{t' : (t', \omega^{t'}) \in \mathcal{S}(\alpha) \text{ and } t' \geq t\}.$$

The stopping times are stochastic, because the cash-flow histories (ω^t) are stochastic.

For any $s \leq t$ and history $h \in \mathcal{H}^t$, let $\bar{V}_s^t(h)$ denote the continuation value of the s-self of the agent, from time t onwards, if all selves before time t are forced to continue, and cash-flow history h has been realized. In particular,

$$\bar{V}_s^t(h) = E\left[\sum_{t'=t}^{\tau_t^\alpha - 1} \alpha(s,t') u(t') | h\right],$$

where by convention the sum equals zero if the index set is empty, and where $\tau = \bar{t} + 1$ if the agent never stops.

Decreasing patience is formalized as follows.

Assumption 4 (Decreasing Patience) *α exhibits* decreasing patience *if, for any $s < s'$, $\alpha(s', \cdot) \prec \alpha(s, \cdot)$ on the time interval $T \cap [s', \bar{t}]$:*

$$\frac{\alpha(s,t)}{\alpha(s',t)} \text{ is increasing in } t \text{ for } t \geq s'.$$

Decreasing patience is satisfied by the two most popular ways of modeling time–inconsistency. The proof of the following examples is trivial and omitted: it suffices to check that $\alpha(s,t)/\alpha(s',t)$ is increasing in $t \geq s'$, for any $s < s'$.

EXAMPLE 1 (HYPERBOLIC DISCOUNTING) *Decreasing Patience holds for hyperbolic discounting, defined as*

$$\alpha(s,t) = 1/(1 + k(t-s))$$

for $t > s$, where k is some positive parameter.

EXAMPLE 2 (QUASI-HYPERBOLIC DISCOUNTING) *Decreasing Patience holds for quasi-hyperbolic discounting, defined as*

$$\alpha(s,t) = 1 \textit{ if } t = s$$

and

$$\alpha(s,t) = \beta \times \delta^{t-s} \textit{ for } t > s.$$

LEMMA 7 *Suppose that α exhibits decreasing patience. Then, for any $s \leq t$ and history $h \in \mathcal{H}^t$, we have*

$$\bar{V}_s^t(h) \geq 0.$$

Proof. We show by backward induction on s that for any $t \geq s$ and $h \in \mathcal{H}^t$,

$$\bar{V}_s^t(h) \geq 0.$$

For $s = \bar{t}$ (and thus, $t = s$), the final self can always avoid the last cash flow and get 0, which proves the condition, since $\tau_{\bar{t}}^{\alpha}$ is chosen by the last self. Suppose now that the induction hypothesis holds for all $s' > s$. We will show that it holds for s. For $t = s$, the inequality holds because the s self can always interrupt cash flows immediately and get 0, and would indeed do so if his present value, conditional on continuing, were negative. Now suppose that $t > s$ and $h \in \mathcal{H}^t$. We have

$$\bar{V}_s^t(h) = \sum_{t'=t}^{\bar{t}} \alpha(s,t')v(t'),$$

where $v(t') = E\left[u(t')1_{t' < \tau_t^{\alpha}} \,\middle|\, h\right]$. We now show that for all $\tilde{t} \geq t$,

$$\sum_{t'=\tilde{t}}^{\bar{t}} \alpha(t,t')v(t') \geq 0. \tag{18}$$

We can write this expression as

$$E\left[1_{\tilde{t}\le\tau_t^\alpha}\sum_{t'=\tilde{t}}^{\tau_t^\alpha-1}\alpha(t,t')u(t')\middle| h\right] = E\left[1_{\tilde{t}\le\tau_t^\alpha}E\left[\sum_{t'=\tilde{t}}^{\tau_{\tilde{t}}^\alpha-1}\alpha(t,t')u(t')\middle|\tilde{h}\right]\middle| h\right],$$

where $\tilde{h}$ is the history at time $\tilde{t}$, and we use the fact that $\tau_t^\alpha = \tau_{\tilde{t}}^\alpha$ whenever $\tilde{t}\le\tau_t^\alpha$ (see observation above). The inner expectation is equal to $\bar{V}_t^{\tilde{t}}(\tilde{h})$ and is nonnegative, by the induction hypothesis. This shows (18).

Applying the discrete version of Lemma 3 (see Quah and Strulovici (2009)), and using that $\alpha(s,\cdot)\succ\alpha(t,\cdot)$, we conclude that

$$\sum_{t'=t}^{\bar{t}}\alpha(s,t')v(t')\ge\frac{\alpha(s,t)}{\alpha(t,t)}\sum_{t'=t}^{\bar{t}}\alpha(t,t')v(t')\ge 0,$$

which proves the induction step and the lemma. ■

To clarify the notation, we let $\bar{V}_s^t(h;\alpha)$ denote the value function $\bar{V}_s^t(h)$ for agent α.

Theorem 9 *Suppose that α and β both exhibit decreasing patience, and that $\beta\succ^{inc}\alpha$. Then,*

$$\mathcal{S}(\beta)\subset\mathcal{S}(\alpha)$$

and, for all (s,h) with $h\in\mathcal{H}^s$,

$$\bar{V}_s^s(h;\beta)\ge\frac{\beta(s,s)}{\alpha(s,s)}\bar{V}_s^s(h;\alpha).$$

Proof. Let $\mathcal{S}^s(\alpha)$ denote the set of pairs (t,h) in $\mathcal{S}(\alpha)$ such that $t\ge s$, i.e., the pairs of time-histories, after time s, at which the agent stops. We show by backward induction on s that $\mathcal{S}^s(\beta)\subset\mathcal{S}^s(\alpha)$. The inclusion holds as an equality for $s=\bar{t}$, since both agents continue at $\bar{t}$ if and only if the expected utility of that period is positive, and get that expected utility, multiplied by their discount factor. Suppose that the inclusion holds for all $s'>s$. We will show that it also holds for s. Consider any $h\in\mathcal{H}^s$. If A (the α agent) stops at (s,h), then there is nothing to show. Suppose that A continues at (s,h). We need to show that B (the β agent) also continues. Since A continues, we have, conditional on h, $\tau_s^\alpha=\tau_{s+1}^\alpha$, and

$$E\left[\sum_{t=s}^{\tau_s^\alpha-1}\alpha(s,t)u(t)\middle| h\right]>0.$$

If B continues, he gets

$$E\left[\sum_{t=s}^{\tau_{s+1}^{\beta}-1}\beta(s,t)u(t)\middle| h\right]. \tag{19}$$

By the induction hypothesis, $\tau_{s+1}^{\beta} \geq \tau_{s+1}^{\alpha} = \tau_s^{\alpha}$. Therefore, (19) may be reexpressed as

$$E\left[\sum_{t=s}^{\tau_s^{\alpha}-1}\beta(s,t)u(t) + \bar{V}_s^{\tau_s^{\alpha}}(h(\tau_s^{\alpha});\beta)\middle| h\right],$$

where $h(\tau_s^{\alpha})$ is the (random) realized history up until time τ_s^{α}. We show that the first term is strictly positive and that the second term is nonnegative. The first term equals

$$\sum_{t=s}^{\bar{t}}\beta(s,t)v(t),$$

where $v(t) = E[u(t)1_{t<\tau_s^{\alpha}}|h]$. We now show that for all $s' \geq s$,

$$\sum_{t=s'}^{\bar{t}}\alpha(s,t)v(t) \geq 0. \tag{20}$$

Proceeding as in Lemma 7, that expression is equal to

$$E\left[1_{s'\leq\tau_s^{\alpha}}E\left[\sum_{t=s'}^{\tau_{s'}^{\alpha}-1}\alpha(s,t)u(t)\middle| h'\right]\middle| h\right],$$

where h' is the history at time s'. The inner expectation is equal to $\bar{V}_s^{s'}(h';\alpha)$, and is nonnegative by Lemma 7. This shows (20). An application of the discrete version of Lemma 3, along with the fact that $\beta(s,\cdot) \succ \alpha(s,\cdot)$, then shows that

$$\sum_{t=s}^{\bar{t}}\beta(s,t)v(t) \geq \frac{\beta(s,s)}{\alpha(s,s)}\sum_{t=s}^{\bar{t}}\alpha(s,t)v(t) \tag{21}$$

and, therefore, that the first term is strictly positive. The second term is nonnegative, because $\bar{V}_s^t(h;\beta)$ is always nonnegative, from Lemma 7.[35] Therefore, it is strictly optimal for B to continue at (s,h), which shows the induction step. Finally, (21) implies that $V_s^s(h;\beta) \geq \frac{\beta(s,s)}{\alpha(s,s)}V_s^s(h,\alpha)$, which shows the second part of the theorem. ■

[35]Recall that nonnegativity of $\bar{V}_s^t$ was obtained by backward induction, and holds for all histories leading up to time t.

6.4 Naive Agent vs. Sophisticated Agent

Let $\tau^N(\alpha)$ ($\tau^S(\alpha)$) denote the optimal (equilibrium) stopping time of a naive (sophisticated) agent with discount function α.[36] Also let $V^N(\alpha)$ denote the value function that a naive agent with discount α *thinks* he is getting, based on his naive assumption that his future selves have the same discount function as he does. Finally, let $V^S(\alpha)$ denote the value function of a sophisticated agent with discount α.

COROLLARY 1 *[Naive vs. Sophisticated Agent] Suppose α exhibits decreasing patience. Then,*

$$\tau^S(\alpha) \leq \tau^N(\alpha),$$

and $V^S(\alpha) \leq V^N(\alpha)$.

Proof. Consider a time-consistent agent, with discount function $\beta(s,t) = \alpha(0,t)$ for all s, t. β trivially exhibits decreasing patience, and $\alpha \prec^{inc} \beta$, because α exhibits decreasing patience. Theorem 9 the implies that the β agent stops at time 0 only if the α agent stops at time 0. Since the naive agent thinks, at time 0, that he is the β agent, he behaves as the β agent: he stops only if the sophisticated α agent stops. Iterating the argument for times $t = 1, 2, \ldots, \bar{t}$ shows that the stopping time of the naive agent is always greater than the stopping time of the sophisticated agent. Finally, Theorem 9 also implies that the naive agent's perceived value, which is the value of the β agent, weakly exceeds that of the sophisticated agent. ∎

In fact, the comparison of value functions in Corollary 1 has another simpler proof and intuition, which does not require the assumption of decreasing patience. The naive agent thinks that he is time-consistent. Equivalently, he believes that his future selves will do exactly what he wants them to do, in contrast to the sophisticated agent, who realizes that he has to compromise with his future selves. Therefore, it is natural that the naive agent's *perceived* value function should be higher than the sophisticated agent's. For the same reason, the naive agent's *perceived* continuation value, if he does not stop immediately, is higher than the real continuation value of the sophisticated agent, and so he only wants to stop if the sophisticated agent wants to stop.[37]

[36] Recall our simplifying assumption, in this section, that if the agent is indifferent between continuing and stopping, he stops. That assumption uniquely pins down the optimal and equilibrium stopping times.

[37] This argument is reminiscent of the control sharing effects in Strulovici (2010): an individual incentives to pursue experimentation are always higher if he is a dictator than if he shares control with other members

Sharing Control

Corollary 1 can be generalized as follows. Let S be a subset of T containing 0. Given α, we define a new discount function by $\hat{\alpha}_S$ in the following way: for any s and $t \geq s$,

$$\hat{\alpha}_S(s,t) = \alpha(s',t)$$

where

$$s' = \min\{\tilde{s} \in S : \tilde{s} \leq s\}.$$

In other words, s' is largest element of S that is smaller than or equal to s. The following claims are easy to show.

PROPOSITION 1 *Suppose that α exhibits decreasing patience. Then,*

a) For any $S \subset T$ containing 0, $\hat{\alpha}_S$ has decreasing patience.

b) If $S' \subset S''$, then $\hat{\alpha}_{S'} \succ^{inc} \hat{\alpha}_{S''}$.

COROLLARY 2 (GRADUAL CONTROL) *Let S' and S'' be two subsets of T containing 0 with $S' \subset S''$. Then*

$$\mathcal{S}(\hat{\alpha}_{S'}) \subset \mathcal{S}(\hat{\alpha}_{S''}),$$

and

$$V(\hat{\alpha}_{S''}) \leq V(\hat{\alpha}_{S'}).$$

Given claims (a) and (b), this corollary follows immediately from Theorem 9.

Interpretation: When evaluating the future at date $t = 0$, the naive agent uses the discount function corresponding to $S = \{0\}$ and the (fully) sophisticated agent uses $S = T$. Between these extremes, we can imagine differing degrees of sophistication corresponding to different subsets of T. For example, suppose $S' = \{0, 3, 5\}$. Then the agent thinks that the date-0 self determines the stopping rule at dates 0, 1, and 2; the date-3 self takes over at dates 3 and 4, and finally the date-5 self determines the stopping rule from date 5 onwards. We say that the agent at $t = 0$ who uses the discount function $\hat{\alpha}_{S''}$ is *more sophisticated* than the agent who uses the discount function $\hat{\alpha}_{S'}$ if $S' \subset S''$.

---

of society (see Theorems 1 and 8 in that paper). The relevant interpretation here is that the naive agent erroneously believes that he is a dictator over his future selves.

When the different temporal selves are interpreted as successive agents, one gets another interpretation. Consider an agency with successive policymakers at each date of set S, a previous policy maker is replaced by another, *less patient* decision maker. Corollary 2 then implies that the more policymaker replacements there are, and the worse the valuation from the viewpoint of the current (date 0) policymaker, and thus the more likely he is to stop. The current policymaker dislikes not only having his tenure shortened, but also when the tenure of *his successors* is shortened as well. For example, $V\left(\hat{\alpha}_{\{0\}}\right) \geq V\left(\hat{\alpha}_{\{0,3\}}\right) \geq V\left(\hat{\alpha}_{\{0,3,5\}}\right)$.

Comparing Actual Values of Naive and Sophisticated Agents

The *actual* value of a naive agent can be higher or lower than that of the sophisticated agent. Consider first the case of an agent whose only time-inconsistency is between his time-0 self, on the one hand, and all of his future selves, on the other hand. That is the agent becomes time-consistent from time 1 onwards. In that case, the time 0 agent is playing a game with only one other self, who will get his way from time 1 onwards. The time 0 agent clearly benefits from being aware of that: the sophisticated agent gets a higher value than the real value of the naive agent. However, the naive agent can sometimes get a higher value. Let the payoff stream (with five dates, beginning with date 0) be

$$0, -1, 1.5, -1, 1.5,$$

and suppose $\alpha(s,t) = 1/(1 + 0.6(t - s))$.

Clearly, the agent at date 3 will not continue. A sophisticated agent at date 1 knows that and he too will not continue. Therefore, the sophisticated agent stops immediately at date 0 and receives nothing.

The naive agent at date 1 does not know that the date 3 agent will stop. Since

$$-1 + 1.5(1.6)^{-1} + (-1)(2.2)^{-1} + 1.5(2.8)^{-1} > 0 > -1 + 1.5(1.6)^{-1},$$

the date-1 agent will continue, thinking that the project will go on until the end, although in fact the project will be stopped at date 3. Therefore, the actual payoff of the naive agent, discounted at date 0, is $(-1)(1.6)^{-1} + 1.5(2.2)^{-1}$; this is positive and thus higher than actual payoff of the sophisticated agent.

7 Applications

7.1 Internal Rate of Return

The internal rate of return of a 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 rates of return. Non-uniqueness has concerned eminent economists, including Samuelson (1937) and Arrow and Levhari (1969); amongst other things, it undermines 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) termination value. This shows the remarkable robustness of the monotonicity studied in their paper and others.

THEOREM 10 *For any decision problem of type (1) or (6), the set of discount rates r for which $V(r) = 0$ is an interval.*[38]

The proof is a direct application of Theorem 2.

7.2 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

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

[38] $V(r)$ is the value of the problem when the discount function is $\alpha(s) = e^{-rs}$.

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 u is increasing in k. Then, one can apply Theorem 4 to conclude that the trajectories of capital $\{k_t\}_{t\in T}$ and the stopping time are decreasing in r.

7.3 Bankruptcy Decisions

As an application of Theorem 1, consider the model of endogenous default 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.[39] 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

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

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 $V(x,r)$ 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

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

[39] 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,[40] 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 termination value.

7.4 Duration of Experimentation

Consider now the standard multi-armed bandit problem originally described by Bellman (1956) and first resolved by Gittins (1979).[41] Suppose that at least one of the arms has a known payoff distribution. Without loss, we may suppose that there is only one such safe arm, since the agent will only ever use the safe arm with the highest expected payoff. It is easy to show, in that case, that as soon as the agent chooses the safe arm, he plays that arm forever after, since he does not learn anything new about his problem, and thus faces the same situation over and over again. We may thus say that the agent stops experimentation as soon as he plays the safe arm. Moreover, we may without loss of generality assume that this payoff is equal to zero. Indeed, observe that the agent's optimization problem is unchanged if all cash–flows are translated upwards or downwards by the same amount. Therefore, it is always possible to translate these cash-flows to guarantee that the safe arm has a zero expected cash-flow. Finally, experimentation stops exactly when the Gittins index of each of the arms drops below 0 (the index of the safe arm), and the number of times that an arm is pulled before experimentation stops is exactly the (possibly infinite) time it takes for its Gittins index to drop below zero. That is, the number of times that a given arm is pulled is exactly the number of times that it would be pulled in a two-armed bandit problem with that arm and the safe arm with zero expected payoff (see Whittle, 1980). That latter problem is

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

[41]Fudenberg and Levine (1998) provide a short exposition.

a pure stopping problem: the only decision is when to stop playing the risky arm. Moreover, stopping experimentation results in a zero termination value: the cash flow of each period where the safe arm is played has zero expectation, and so the present value of all such cash flows is zero, independently of the discount rate used by the agent.[42] Applying Theorem 1 then yields the following result.

Corollary 3 *In any multi-armed bandit with a safe arm, a more patient agent stops experimentation later.*

Cohen and Solan (2008), who characterize the optimal experimentation policy for a two-armed bandit problem where one arm is safe and the other is a Lévy process with unknown parameters, show the result explicitly, using their characterization. This result generalizes a feature of the cooperative solutions in both the Brownian setting of Bolton and Harris (1999) and the Poisson setting of Keller, Rady, and Cripps (2005). Corollary 3 generalizes Cohen and Solan's result to an arbitrary learning process.[43]

In the above setting, the agent's speed of experimentation is fixed: he can only pull one arm at any time (and in models where the agent can split resources across arms, his resource per unit of time is constant over the entire horizon). When the agent can choose the speed of experimentation, his value function is still monotonic, but the actual *time* spent experimenting is not necessarily increasing with patience. As illustrated by Counterexample 4.3, a more patient agent may experiment at a faster rate than a less patient one, resulting in the experimentation boundary being reached faster, when multiple experimentation levels are available. Moreover, all positive results pertain to the case of a single agent. It would be interesting to find an extension for the case of strategic experimentation (see Bolton and and Harris (1999) and Keller, Rady, and Cripps (2005)), or in the context of collective experimentation with voting (see Strulovici (2010)).

[42]This argument breaks down with stochastic discounting, if the discount factor is correlated with the cash-flows.

[43]All the papers cited in this paragraph allow the agent to split resources, at any time, between the safe arm and the risky arm. For a single agent problem, however, the solution is always "bang-bang," so that the problem is equivalent to forcing the agent to putting all its effort into a single arm, at each time.

8 Discussion of Switching Models

A stopping problem is a particular instance of a switching problem. A switching problem is one where the agent alternates between different technologies, projects, or, generally, streams of cash flows. For example, Dixit (1989) studies a model where a firm switches between entry and exit of a commodity market, whose the commodity price fluctuates over time. There is typically a switching cost for moving from one stream to another. Stopping amounts to a particular switching problem with two streams to choose from: the first one is the one available before stopping (project), and the second is identically zero. The switching cost from the first stream to the second one is the termination value of the project, and the switching cost from the second stream to the first one is infinite, so that one never reverts back to the project after stopping.

It is natural to ask whether the results of earlier sections have an extension to switching. In discrete time, a switching problem is simply a kind of control problem, where the state being controlled is the technology or stream of cash flows among those available. The agent can switch to another state, at some cost. Theorem 2 directly implies that the value is increasing in patience.[44]

In continuous time however, our previous results cannot be directly applied, because switching now takes the form of an "impulse" control, resulting in a lump-sum cost (i.e., not absolutely continuous with respect to the Lebesgue measure on the time domain). We sketch an argument for why value monotonicity should also hold in continuous time.

When switching is costless, the value function of a switching problem is always increasing with patience, because switching can be seen as a control applied on the cash–flow stream, so that Theorem 2 applies as is, modeling switching as a control. When switching entails lump–sum costs, the argument is more subtle. The first switching time can be thought of as a stopping time, whose termination value is the continuation value of the switching problem. Provided that this continuation value is nonnegative and increasing in patience, the value function of the problem will be increasing in patience as well, by an obvious modification of Theorem 2. Showing that the continuation value is monotonic in patience may be proved by a backward induction argument, setting an artificial end of time with a fixed termination

[44]As pointed out in Section 2, all results of the paper except those of Section 4, hold in discrete time.

value, and proceeding backward in time.

However, while one should expect the value function of a switching problem to be increasing in patience (again, provided that the value function is nonnegative at all times), the switching times themselves have no reason to be ordered in a particular way. One easy way to see this is to think of switching as a control problem (i.e., zero switching costs), and observe that in that case, the intuition from Counterexample 4.3 applies. Another way to see this is to look at the analysis of Dixit (1989), who shows that the hysteresis interval expands monotonically with the discount rate, which is inconsistent with a more patient agent always switching earlier.

9 Conclusion

This paper shows that the expected present value of any cash–flow (or utility) stream which may be interrupted optimally by its undertaker, and which yields a nonnegative termination value, is increasing with patience, even when the cash–flows or utils are stochastic and controlled by the agent. Because of the generality of the setting, the result holds even in situations where control has a long term impact on the payoffs, as in the case, for example, of gas emissions, or of learning about a process. The result must be qualified in several ways.

First, while the capital value of a project is decreasing in the interest rate under very general conditions, the *timing* of the project interruption need not be monotonic. This means, for instance, that a lower discount rate may result in a project being managed in such a way that it is actually abandoned earlier than under a higher discount rate. This observation is new to our knowledge, and may bear important consequences. For example, one cannot exclude, a priori, that giving a firm access to a lower cost of capital may result in decisions that actually will precipitate its default. However, when the utility flow of the agent is an increasing function of some one-dimensional state, a property that holds in many economic models, we do show that the timing of the project is monotonic in patience, no matter how complicated the control problem of the agent is (in particular, the control variable can have arbitrarily many dimensions).

Second, if the agent has a stochastic discount factor, the result need not hold any more.

This is potentially important, since stochastic discounting is a realistic feature in many applications. Going back the Keynesian relation between investment and interest rate, what if agents anticipate that their future borrowing rate will be correlated with their income or project's income? We provided two examples where the expected present value of a patient agent was actually lower than that of less patient one. One example exploited some negative correlation between the stochastic discount factor and the cash flows, while the other had purely deterministic cash–flows. Fortunately, monotonicity is restored under some correlation condition relating discount factors and cash flows. In particular, what matters is the variation of the ratio of discount factors between the patient and the impatient agents, rather than the discount factors themselves. One loose way to interpret the condition is a procyclicality between how much more patient the patient agent is relative to the other agent, and the cash flows faced by the agent.

Third, the results may be reversed if the agent is time–inconsistent. Intuitively, an agent who is moderately patient today, but anticipates that his future self will be much more patient and ready to incur high immediate costs for long run benefits, may decide to prevent such investment from being available. Thus, increased overall patient may precipitate project interruption. A similar intuition applies if we think of the agent as a representative agent, whose identity at each time represents a generation of individuals. Under such intergenerational interpretation, time–inconsistency is particularly important, as observed by Phelps and Pollak (1968), as the discount function can be interpreted as a measure of altruism towards future generations. If a first, moderately altruistic generation anticipates the second generation to be very altruistic towards far away generations, at the expense, say, of the third generation (which the first generation cares about), then it may behave so as to prevent the second generation from engaging in .

Fourth, we have focused on a single agent. It would be interesting to find general conditions for multiagent environments, including R&D games or when decisions are made collectively, where experimentation is increasing with agents' patience. Robust results may also be obtained in a general equilibrium framework or games with a continuum of agents.

10 Appendix

10.1 Proof of Theorem 1 with Termination Value

Here we consider the case of an arbitrary nonnegative termination value process G_t. The objective can be reexpressed as

$$\max_{\hat{\tau}} E\left[\int_0^{\hat{\tau}} \alpha(s)\left(u_s ds + dG_s + G_s \frac{d\alpha_s}{\alpha_s}\right)\right],$$

where the last two integrals are Lebesgue-Stieltjes integral.[45] This objective is similar to the case of a zero termination value, except that the new flow utility now directly depends on α through the third term $G_s \frac{d\alpha_s}{\alpha_s}$. Repeating arguments from Section 2, one may show that if

$$E\left[\int_t^{\tau} \alpha(s)\left(u_s ds + dG_s + G_s \frac{d\alpha_s}{\alpha_s}\right)\right] \geq 0$$

for all $t < \tau$, then

$$E\left[\int_0^{\tau} \beta(s)\left(u_s ds + dG_s + G_s \frac{d\alpha_s}{\alpha_s}\right)\right] \geq \frac{\beta(0)}{\alpha(0)} E\left[\int_0^{\tau} \alpha(s)\left(u_s ds + dG_s + G_s \frac{d\alpha_s}{\alpha_s}\right)\right] \geq 0. \tag{22}$$

For this argument, we replace v in (2) by the deterministic process V of bounded variation such that

$$dV(s) = E\left[\left(u_s ds + dG_s + G_s \frac{d\alpha_s}{\alpha_s}\right) 1_{s<\tau}\right].$$

Optimality of τ for α implies that

$$\int_t^{\bar{t}} \alpha(s) dV_s \geq 0$$

for all $t < \bar{t}$. A simple generalization of Lemma 3 then implies that

$$\int_0^{\bar{t}} \beta(s) dV_s \geq \frac{\beta(0)}{\alpha(0)} \int_0^{\bar{t}} \alpha(s) dV_s,$$

which shows (22).

[45] See, for example, Royden (2007) for a definition. The integrals are well defined whenever the integrating process has bounded variation and the integrand is measurable and bounded.

Since β/α is increasing, $\frac{d\beta_s}{\beta} \geq \frac{d\alpha_s}{\alpha}$ for all s (the inequality holds even when β and/or α are discontinuous). Combining the previous inequalities and using that G is nonnegative, we conclude that if τ is optimal for α, then

$$E\left[\int_0^\tau \beta(s)u_s ds + \beta(\tau)G_\tau\right] \geq \frac{\beta(0)}{\alpha(0)} E\left[\int_0^\tau \alpha(s)u_s ds + \alpha(\tau)G_\tau\right]. \tag{23}$$

Thus

$$V(\beta) \geq \frac{\beta(0)}{\alpha(0)} V(\alpha). \tag{24}$$

The rest of the proof is identical to the case of a zero termination value (see Section 2).

10.2 Proof of Theorem 4 with Termination Value

We now assume that the termination value function $G(x,t)$ is nonnegative and increasing in x for all t.

Monotonicity of G seems reasonable when x corresponds to capital which can be sold when a project is interrupted. 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 4.

Now we modify the function h in (10) to account for the termination 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.[46] and let

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

With this new definition, it is easy[47] to check that one still has[48] $\int_t^\infty e^{-rs}h_s ds = E[V(y_t,t)1_{\tau(y)>t} - V(x_t,t)1_{\tau(x)>t}] \geq 0$, and that $V_r(y,r) - V_r(x,r) = \int_0^\infty s e^{-rs} h_s ds \geq 0$ when $y > x$, by another application of Lemma 3. The rest of the proof then follows.

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

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

[48]The argument in Lemma 4 requires nonnegativity of the value functions $V(y_t,t)$ and $V(x_t,t)$, which is guaranteed here by nonnegativity of the termination value G.

10.3 Extension of Section 4 to General Discount Functions

Suppose that $\alpha \prec \beta$, and that α and β have absolutely continuous logarithms:

$$\alpha(t) = \alpha(0) \exp\left(- \int_0^t r_\alpha(s) ds \right)$$

for some discount rate function r_α, with a similar expression for β. With that representation, $\alpha \prec \beta$ is equivalent to

$$r_\beta(t) \leq r_\alpha(t)$$

for all $t \in T$.

Let

$$\gamma_\nu(t) = \alpha(t) \left(\frac{\beta(t)}{\alpha(t)} \right)^\nu,$$

for $\nu \in [0, 1]$. It is easy to check that $\gamma_0 = \alpha$, $\gamma_1 = \beta$, and $\gamma_\nu \prec \gamma_{\nu'}$ for all $\nu \leq \nu'$.

Lemma 5 is modified as follows: let $V(x, \nu)$ denote the value function of an agent with discount function γ_ν, starting from state x. Then, $V(x, \nu)$ is submodular. Indeed, the envelope theorem implies that

$$\frac{\partial V(x, \nu)}{\partial \nu} = E\left[\int_0^{\tau_\nu} \left(- \log\left(\frac{\beta(t)}{\alpha(t)} \right) \right) \gamma_\nu(t) u(x_t, \lambda_\nu(t)) dt \right].$$

To show that $V_\nu(y, \nu) - V_\nu(x, \nu) \leq 0$ for $y \geq x$, we use the same argument as in Lemmas 4 and 5, with the difference that, instead of $\int_0^\infty s e^{-rs} h_s ds \geq 0$, we show that

$$\int_0^\infty \log\left(\frac{\beta(s)}{\alpha(s)} \right) \gamma_\nu(s) h_s ds \geq 0.$$

The result is obtained by an application of Lemma 3 to the increasing function $s \mapsto \log\left(\frac{\beta(s)}{\alpha(s)} \right)$.

The HJB equations are modified only by replacing r with the time varying rate $r_\nu(t) = r_\alpha(t) + \nu(r_\beta(t) - r_\alpha(t))$. The rest of the proof is identical to that of Section 4.

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.

Bensoussan, A. (1992) *Stochastic Control of Partially Observable Systems*, Cambridge University Press.

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

Cohen, A., and Solan, E. (2008) "Bandit Problems with Lévy Processes," Working Paper, Tel Aviv University.

Dixit, A. (1989) "Entry and Exit Decisions Under Uncertainty," *Journal of Political Economy*, Vol. 97, No. 3, pp. 620–638.

Duffie, D (2001) *Dynamic Asset Pricing Theory*, Third Edition. Princeton University Press.

Flemming, J., Wright, J. (1971) "Uniqueness of the Internal Rate of Return: A Generalization," *Economic Journal*, Vol. 81, No. 322, pp. 256–263.

Fleming, W. H., Soner, H.M. (1993) *Controlled Markov Processes and Viscosity Solutions*, Springer-Verlag, New York.

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

Fudenberg, D., Levine, D. (1998) *The Theory of Learning in Games*, MIT Press, Cambridge, MA.

Gittins, J. (1979) "Bandit Processes and Dynamic Allocation Indices," *Journal of the Royal Statistical Society*, Vol. 41, No. 2, pp. 148-177.

Hicks, J. (1973) *Capital and Time*, Clarendon Press, Oxford.

Karatzas, I., Shreve, S. (1991) *Brownian Motion and Stochastic Calculus*, Springer-Verlag, New York.

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. (2010) "Performance-Sensitive Debt," *Review of Financial Studies*, Vol. 23, No. 5, pp. 1819–1854.

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

MILGROM, P., SHANNON, C. (1994) "Monotone Comparative Statics," *Econometrica*, Vol. 62, No. 1, pp. 157–180.

NORDHAUS, W (2007) "A Review of the *Stern Review on the Economics of Climate Change*," *Journal of Economic Literature*, Vol. 45, No. 3, pp. 686–702.

O'DONOGHUE, T., RABIN, M. (1999) "Doing It Now or Later," *American Economic Review*, Vol. 89, No. 1, pp. 103–124.

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

PHELPS, E., POLLAK, R. (1968) " On Second-Best National Saving and Game-Equilibrium Growth," *Review of Economic Studies*, Vol. 35, No. 2, pp. 185–199.

QUAH, J., STRULOVICI, B. (2009) "Comparative Statics, Informativeness, and the Interval Dominance Order," *Econometrica*, Vol. 77. No. 6, pp. 1949–1992.

ROSS, S., SPATT, C., AND DYBVIG, P. (1980) "Present values and internal rates of return," *Journal of Economic Theory*, Vol. 23, pp. 66-81.

ROYDEN, H. (2007) *Real Analysis*, 4^{th} Edition, Prentice Hall.

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*, Vol. 11, pp. 340-355.

Stern, N., Peters S., Bakhshi V., Bowen A., Cameron C., Catovsky S., Crane D., Cruickshank S., Dietz S., Edmonson N., Garbett S.-L., Hamid L., Hoffman G., Ingram D., Jones B., Patmore N., Radcliffe H., Sathiyarajah R., Stock M., Taylor C., Vernon T., Wanjie H., and Zenghelis, D. (2006) *Stern Review: The Economics of Climate Change*, HM Treasury, London.

Strotz, R. (1955) "Myopia and Inconsistency in Dynamic Utility Maximization," *Review of Economic Studies*, Vol. 23, No. 3, pp. 165–180.

Strulovici, B. (2010) "Learning While Voting: Determinants of Collective Experimentation," *Econometrica*, Vol. 78, No. 3, pp. 933–971.

Topkis, D. (1978) "Minimizing a Submodular Function on a lattice," *Operations Research*, Vol. 26, pp. 305–321.

Veinott, P. (1989) *Unpublished Lecture Notes*, Johns Hopkins University.

Weitzman, M. (2007) "A Review of the *Stern Review on the Economics of Climate Change*," *Journal of Economic Literature*, Vol. 45, No. 3, pp. 703–724.

Whittle, P. (1980) "Multi-armed bandits and the Gittins index," *Journal of Royal Statistical Society*, Vol. 42, pp. 143–149.