working paper
department
of economics

massachusetts
institute of
technology

50 memorial drive
cambridge, mass. 02139

WORKING PAPER
DEPARTMENT
OF ECONOMICS

TIME CONSISTENT OPTIMAL STOPPING
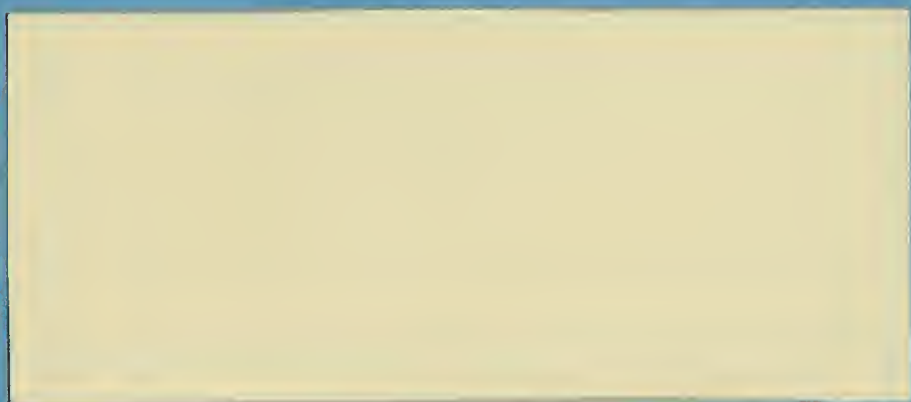
Lones Smith

No. 97-5   March, 1997

MASSACHUSETTS
INSTITUTE OF
TECHNOLOGY

50 MEMORIAL DRIVE
CAMBRIDGE, MASS. 021329

Time-Consistent Optimal Stopping

Lones Smith[*]

Department of Economics

Massachusetts Institute of Technology

Cambridge, MA 02139

March 17, 1997

Abstract

This paper investigates time consistent preferences where the strategy space consists of a stopping time, and the decision-maker also acts under uncertainty. Despite the requirement of identical vNM preferences, it is shown that in addition to the exponential discounting class identified by Strotz (1956), negative affine discounting also yields time consistent behavior.

1. INTRODUCTION

There is a rich literature starting with Strotz (1956) examining time-consistent planning.[1] The insight here is that even in a world of perfect certainty, time preference in a dynamic environment must assume a very specific functional form if individuals are to have an incentive to carry out their previously-conceived optimal plans. In a nutshell, current and future 'selves' might well disagree on how the wisest course of action for the future self.

The cited literature is interested in behavior described by a very rich action space — often infinite-horizon consumption paths. Strotz' main finding is that if the discounting at time t of future consumption at time $\tau > t$ only depends on the horizon length $\tau - t$, then preferences are time-consistent exactly when there is exponential discounting: The consumption weight is $e^{-\beta(\tau-t)}$, for some $\beta > 0$. This paper asks this same question for a simpler action space and payoff structure that arises in a microeconomics context — optimal stopping under uncertainty. Here, identical preferences alone, or sooner is preferred to later, has no cutting power. But because of the stochastic environment, time-consistent preferences must meet a sharper test: Individuals must enjoy the identical von Neumann-Morgenstern (vNM) preferences as their future selves. In other words, a harsher standard must be met by a smaller class of payoff functions than in Strotz context. On balance, I show that discounting may be either exponential or negatively affine (i.e. time-τ consumption is weighted by $\gamma - \beta(\tau - t)$ as of time t).

There are two reasons to care about this result. First, it is useful for economic theorists to know the full gamut of possible time-consistent dynamic objective functions under uncertainty. Second, the precise functional equations required by time-consistency can be easily understood in the simple optimal stopping context, and qualitatively differs from the deterministic infinite horizon consumption framework.

2. THE ANALYSIS

A. The Model.   A decision maker has a choice set $A \cup \{\text{stop}\}$ in every period. After each non-stopping action $a \in A$, a stochastic signal $y \in Y$ is observed. For instance, A may be a list of places to look each period, and y the simple fact that the process hasn't found the sought-for prize. As a function of the history h^t to time t of actions, and observed signals, strategy σ_t is a choice of a new (possibly stopping) action, which will stop the

---

[1]For a recent and quite encompassing bibliography, see Donoghue and Rabin (1996).

process with chance $\pi(h^t)$. The goal is to stop the process as soon as possible. Also, stopping at time $t > s$ is weighted by a factor $w(t - s)$ that is decreasing in t. This embeds two assumptions: first *strict time preference* (stopping sooner is preferred to later), and second *time symmetry* (since more generally one could envision weighting by $w(s, t)$).

An optimal strategy at time s is one that maximizes

$$U_s(\sigma) = \mathcal{E}\left(\sum_{t=s}^{\infty} w(t - s)\pi(h^t)|\sigma\right) \tag{1}$$

where the expectation is over future signals $\langle y_t \rangle$ given the strategy σ, conditional on not having yet stopped. I assume that a strategy exists that stops in finite expected time.

B. The Main Result. In every period, one's actions affect the evolution of the future lotteries over public signals. Thus, it is necessary and sufficient that an individual at any two times $s_1 < s_2$ have identical vNM preferences over actions at all times $t \geq s_2 > s_1$.

Lemma *One's vNM preferences over future stopping times are constant iff there exist functions $a > 0$ and b such that*

$$w(s + t) = a(s)w(t) + b(s) \tag{2}$$

Proof: Since separate stopping times are mutually exclusive states of the world, with a pyoff realized just once, the weight w behaves like a vNM utility function. By the uniqueness part of the vNM Expected Utility Theorem, vNM preferences are constant iff these utilities are positively affinely related. Hence, it is necessary and sufficient that there exist functions $\hat{a} > 0$ and $\hat{b}$ of (s_1, s_2) such that

$$w(t - s_1) \equiv \hat{a}(s_1, s_2)w(t - s_2) + \hat{b}(s_1, s_2) \tag{3}$$

Since this must hold if we replace (s_1, s_2, s) by $(s_1 + \tau, s_2 + \tau, s + \tau)$, and because w is nonconstant in t, it follows that $\hat{a}(s_1, s_2) \equiv \hat{a}(s_1 + \tau, s_2 + \tau)$ and $\hat{b}(s_1, s_2) \equiv \hat{b}(s_1 + \tau, s_2 + \tau)$. Thus, $\hat{a}(s_1, s_2) \equiv a(s_2 - s_1)$ and $\hat{b}(s_1, s_2) \equiv b(s_2 - s_1)$. Equivalently, putting $s = s_2 - s_1$, the characterizing equation (2) for w follows. $\qquad \square$

Proposition (Time-Consistent Characterization) *The only monotonic decreasing weights w solving (2) and yielding a well-defined objective function (1) are exponential $w(t) = \alpha e^{-\beta t} + \gamma$, or negatively affine $w(t) = -\beta t + \gamma$ in time, where $\alpha, \beta > 0$.*

Proof: The functional equation (2) is generally solved in Corollary 1 in section 3.1.3 of Aczel (1966). That result also admits the possibility of a negatively exploding exponential

2

function ($\alpha, \beta < 0$) for which the objective (1) will not be defined. Conversely, assuming a finite expected stopping time, any vanishing exponential or negative affine function will yield a well-defined objective. $\square$

C. Discussion of an Example. Suppose that one wishes to find a moving target 'as soon as possible' given partial information of its location.[2] What time-consistent preferences describe this objective function? Since there is just single payoff, we may WLOG set $\gamma = 0$ in the Proposition — for it clearly will not affect behavior. In this case, an optimal search strategy is time-consistent iff the payoff from discovery in period t is $e^{-\beta t}$ or is $-t$. The latter negative affine weighting, where the player seeks to minimize the *expected discovery time*, is the most natural interpretation of the stated goal. It can also be seen as is the limit of the exponential weighting as the interest rate β vanishes: $-t = \lim_{\beta \to 0} (e^{-\beta t} - 1)/\beta$.

It is tempting to draw inferences about more general stopping stopping games such as price search or alternating offer bargaining. In those cases, where there is either a more general terminal payoff, or intermediate payoffs, the Lemma fails, as the negatively affine solution vanishes. This brings us back to the Strotz solution (but with uncertainty).

References

Aczel, J., *Lectures on Functional Equations and their Applications*, New York: Academic Press, 1966.

Donoghue, Ted and Matthew Rabin, "Doing it Now or Later," 1996. Berkeley Mimeo.

Smith, Lones, "Equilibrium and Efficiency in the Rendezvous Problem," 1994. MIT Mimeo.

Strotz, R., "Myopia and Inconsistency in Dynamic Utility Maximization," *Review of Economic Studies*, 1956, *23*, 165–180.

---

[2]This problem is quite relevant: See Smith (1994) for a compelling coordination game with this flavor.

Date Due