

Uncertainty and Hyperbolic Discounting[†]

by

Partha Dasgupta* and Eric Maskin**

Revised
November 2004

*Faculty of Economics, University of Cambridge
**Institute for Advanced Study, Princeton
and
Department of Economics, Princeton University

We are grateful to the Beijer International Institute of Ecological Economics, Stockholm, and the National Science Foundation (SES-0318103) for research support. For comments on earlier versions of this paper, we thank Roland Bénabou, Douglas Bernheim, David Colander, Nick Davies, Eddie Dekel, Avinash Dixit, Ernst Fehr, Christopher Harris, Alex Kacelnik, John Kagel, Adi Livnat, François Maniquet, Ariel Rubinstein, Larry Samuelson, Peter Sozou, and three referees.

[†] This paper was previously entitled “Uncertainty, Waiting Costs, and Hyperbolic Discounting.”

Abstract

We propose an evolutionary explanation for the pattern of intertemporal preference reversals often ascribed to “hyperbolic discounting.” We take the view that preferences—manifested, for example, in urges, cravings, and inclinations— are the outcome of evolutionary forces, and so will induce animals or humans to make survival-maximizing choices in “typical” decision problems. We show that if the typical problem involves payoffs whose realization times are uncertain, then optimal preferences give rise to relatively patient behavior when the time horizon is long but induce a switch to impatience when the horizon grows short.

Such reversals do not entail dynamic inconsistency in typical decision problems; behavior there is optimal. However, if a decision-maker is confronted with a choice for which the realization-time uncertainty falls outside the evolutionary norm, her preferences may well prompt her to behave inconsistently. We argue that, if such a choice problem recurs, her evolutionarily endowed ability to learn will lead her to make self-commitments against these urges.

JEL Classification: D11, D81, D91.

1. Introduction

Empirical studies in economics and behavioral ecology suggest that *ceteris paribus* animals and humans appear to place less weight on the future than on the present—i.e., they act as though they *discount* future payoffs. Furthermore (and more interestingly) they do so with discount rates that increase as the time before those payoffs are realized grows shorter.¹ In other words, subjects act as though they become less patient when payoffs are more imminent.

An anecdotal (human) example is offered by O'Donoghue and Rabin (1999): when offered the choice in February between a painful seven-hour task (e.g., preparing a tax return) on April 1 and a painful eight-hour task on April 15, most of us, they suggest, will opt for the earlier date. But as April 1 approaches, we are apt to change our minds, if we can, and postpone the pain to the 15th, even though it will then be greater. Thus, we behave as though we discount the later pain more as time grows short. Strotz (1956) discusses a similar phenomenon involving positive payoffs. Early in the calendar year, he notes, many people attempt to lay money aside for Christmas. However, as time goes by, they may find themselves spending the money on summer vacations or back-to-school clothes. It is as though they have become more impatient than they were back in January.

Both the O'Donoghue-Rabin and Strotz examples accord with “hyperbolic discounting,”² which has attracted considerable interest among economists because it appears to shed light on important economic phenomena such as household saving behavior (see Laibson 1997, Bernheim, Skinner, and Weinberg 2001, and Harris and Laibson 2001).

¹ The empirical literature on birds (particularly pigeons and starlings) is summarized in Mazur (1987) and Green and Myerson (1996); that on humans in Ainslie (1992).

² Strictly speaking, hyperbolic discounting requires that the discount rate should vary inversely with the time to payoffs. But the term has come to be applied more generally to any manifestation of increasing impatience as time horizons shrink. See Rubinstein (2001), who cautions that many observed behaviors are consistent with a variety of preferences in addition to those involving hyperbolic discounting.

One might ask whether there is some reason for such behavior. The point of view that we take here is that, to a considerable extent, animal and human behavior is shaped by *preferences* that are the outcome of evolutionary forces. That is, cravings, urges, or instincts—the operational manifestations of preferences—induce the animal or human to make the right choice in the “average” situation that it, he, or she is likely to face, where the “right” choice means the one that maximizes survival (and, therefore, the opportunity to reproduce).³ We show that if the “average” situation entails some uncertainty about when payoffs are realized, the corresponding preferences may well entail hyperbolic discounting, giving rise to preference reversals.

To get a rough understanding of this result, imagine that a decision-maker (DM) is offered the choice between a small (positive) payoff relatively soon (prospect P) or a big payoff relatively late (prospect P'). Furthermore, suppose that, for either P or P' , there is a small but positive probability that at any time t before the anticipated payoff dates, the payoff will be realized early at t rather than when expected. Assume that the DM initially opts for prospect P' ; the reason for making this choice, of course, is P' 's bigger payoff, enhanced by the chance of early realization.⁴ But, as time passes, if the payoff from P' does not materialize, the likelihood of early realization declines. Of course, for the same reason, the likelihood of early realization wanes for P too - - but this does not matter so much because P 's payoff is anticipated sooner anyway. Hence, with time, the more immediate prospect P , becomes increasingly attractive relative to P' and eventually the DM may switch to P .

Although, as this example illustrates, our model predicts preference reversals that accord with hyperbolic discounting, these reversals will be entirely dynamically *consistent* as

³ In this respect, our approach is similar to that of Samuelson and Swinkels (2002), who—although their approach is otherwise quite different from ours—share our view that an animal's urges are evolutionarily determined substitutes for full information about the choices it faces.

⁴ There is a connection here with the theory of sequential search among risky prospects (Weitzman 1979), which shows that *ceteris paribus* choosing prospects with higher dispersions before those with lower dispersions is the optimal search strategy.

long as the DM confronts decision problems (i.e., the DM will actually profit from reversing herself). However, the DM's urges may lead her astray - - resulting in dynamic *inconsistency* - - if she faces decisions for which the payoff-realization times depart from the norm. Thus, evolution also enables her to *learn* how to overcome or neutralize those urges when an atypical situation recurs. For example, in experimental settings, pigeons discover how to commit themselves not to switch from "patient" to "impatient" choices (Rachlin and Green 1972; for more on this see section 5). And people susceptible to the impulse to spend their holiday savings prematurely find that they can thwart that inclination by putting their money in illiquid Christmas accounts (Strotz 1956).

In section 2, we examine the standard rationale for discounting (viz., to take account of the risk that future payoffs may disappear or depreciate) and discuss why some previous explanations for hyperbolic discounting turn out not to be consistent with the O'Donoghue-Rabin and Strotz phenomena. In section 3, we introduce uncertainty about when payoffs are realized and show that this leads to a version of hyperbolic discounting that *does* explain these reversals (Propositions 1 and 1*). In section 4, we show that our analysis extends to forms of uncertainty considerably more general than allowed for in section 3 (Proposition 2). Finally, in section 5, we turn to atypical decision problems and dynamic inconsistency.

2. Discounting and Hazard Rates

Before turning to hyperbolic discounting, we must first ask the question, "Why should a DM discount at all?" A conventional answer, provided by both the economics and zoology literatures, is that, in a typical situation that an animal or human may face, future payoffs run some risk of disappearing or depreciating (Yaari 1965). Suppose, for example, that a blackbird tends to hang around a particular raspberry bush waiting for the fruit to ripen.⁵ Before this happens, however, a flock of crows (which don't care about

⁵ We are indebted to the ornithologist Nick Davies for our bird/fruit examples.

ripeness) may descend on the fruit, devouring it all. Hence, the blackbird should discount the payoff of getting the fruit, where the discount rate is the *hazard rate* of the “crow arrival”

process (the hazard rate at time t is $\frac{1}{\Delta t}$ times the probability that the crows will arrive

between times t and $t + \Delta t$, conditional on their not having arrived already). In particular, if this process is Poisson—i.e., the hazard rate is independent of t —the implied discount rate is constant. Thus, if a prospect (e.g., eating raspberries) has a payoff V (the calories in the berries) at time T (the time at which the berries ripen) and the hazard rate is a constant r , the prospect should be evaluated as though the payoff is $e^{-rT}V$.⁶

This model suggests an immediate potential explanation for hyperbolic discounting. Suppose that the hazard rate depends on time, i.e., we can express it as $r(t)$. Assume furthermore that r is decreasing in t . Then, more distant prospects, as will be discounted at a lower rate than more proximate prospects, as hyperbolic discounting demands. One problem with this simple story, however, is that in most settings there is no particular reason why we should suppose that the hazard rate is declining (there is another, more serious difficulty with the story, which we will come to below).

An illuminating paper, Sozou (1998), offers an alternative theory. Sozou supposes that the hazard rate is constant but *unknown* to the decision-maker (where the uncertainty corresponds to a gamma distribution). He considers *two* choice problems: one in which the DM chooses between a payoff $V (> 0)$ at time T and a larger payoff V' at time $T' (T < T')$;

⁶ Another explanation for discounting—treated in an earlier version of this paper (Dasgupta and Maskin 2002)—turns on the idea that waiting for a payoff to materialize is often costly, either because the blackbird may have to use up energy while waiting (a physiological cost), or because by hanging around the raspberry bush it may lose other opportunities for food (an opportunity cost). This alternative explanation for discounting has the advantage of immediately explaining the empirical finding (see Green and Myerson 1996) that larger payoffs are discounted less than smaller ones. Specifically, suppose that a bird is offered the choice between a prospect with reward αV at time T and one with reward $\alpha V'$ at time T' , where α is a scalar and $V' > V > 0$ and $T' > T$. Experiments suggest that as α rises, the bird is more likely to favor the latter prospect. This effect is predicted by the waiting-cost model of discounting simply because, for larger α , waiting costs matter less relative to rewards.

the other in which she chooses between the same payoffs V and V' , but at times $T + t$ and $T' + t$, respectively. Sozou shows that if the DM has the same uncertainty about hazard rates in the two problems, she will exhibit greater patience in the latter problem, i.e., she is more apt to choose the payoff at $T' + t$ in that problem (this accords with hyperbolic discounting, which predicts that a decision-maker will use a lower discount rate for a choice problem involving a longer time horizon).⁷ To understand this result, notice that the choice between (V, T) and (V', T') matters only if neither payoff disappears before time T . Hence, given the uncertainty about the hazard rate, the DM should evaluate the payoffs using the expected discount rate *conditional* on reaching T without the payoffs disappearing. Similarly, she should use the expected discount rate conditional on reaching $T + t$ to make the choice between $(V, T + t)$ and $(V', T' + t)$. But the latter discount rate is lower than the former (the longer the payoffs do not disappear, the lower is the conditional hazard rate), and so the result follows.

However, this model cannot explain the sort of preference reversals in the O'Donoghue-Rabin tax story or the Strotz Christmas story. Indeed, it predicts just the *opposite*. Imagine that a DM decides early on that she will do her taxes on April 1. As that date approaches and the need to perform the onerous task has failed to disappear (e.g., Congress has declined to repeal the tax laws, or the DM's CPA brother-in-law is not available to do the taxes for her), the conditional hazard rate is now lower than it was *ex ante*, and so the DM should discount the future *less*. That is, she should be *even more* resolved than before to refrain from postponing the pain to the 15th.⁸

⁷ The same result has also been obtained by Weitzman (2001), but for the purpose of arguing that discount rates used in public investment projects ought to be declining functions of time.

⁸ There is at least one existing explanation for hyperbolic discounting, based on the foraging model of Holling (1959), which *can* account for preference reversal in the tax example. In a setting where animals make repeated choices over time, this theory postulates that animals act to maximize payoff per unit of time (see Green and Myerson, 1996, and Kacelnik, 1997). It is not clear, however, how to reconcile this postulate with the more conventional maximand of *total* payoff.

The story we started with, in which it is simply assumed that the hazard rate declines with time, also fails to work. It implies that preferences between doing taxes on April 1 and April 15 should not change at all from February to the end of March (the hazard rate, although varying with time, is the same for both options in that time interval).

The failures of this simple story and the Sozou model to explain preference reversal in the tax and Christmas savings scenarios illustrate that there are two different meanings in the literature for the term “hyperbolic discounting.” One meaning (used, for example, by Sozou 1998) pertains to comparisons *across* decision problems, e.g., to the comparison between (V, T) and (V', T') on the one hand and $(V, T + t)$ and $(V', T' + t)$ on the other. The other—invoked by those interested in preference reversals (e.g., O’Donoghue and Rabin 1999, Strotz 1956, Rachlin and Green 1972)—applies to a comparison within the *same* decision problem, but at different times. Our focus is on the latter concept.

3. Hyperbolic Discounting

Let us turn to our own explanation of hyperbolic discounting. Our model, like Sozou’s, emphasizes uncertainty, but now the uncertainty pertains to *when* payoffs are realized (rather than to whether they will be realized at all). Recall the blackbird waiting for the raspberries to ripen. The bird may be pretty sure that the berries will be ripe by tomorrow morning, but let us imagine, as realism dictates, that there is some chance that they will ripen earlier (or later). As before, let the prospect of eating the raspberries be $P = (V, T)$, where $V (> 0)$ is the payoff from consuming the berries and T is the time when the payoff is most likely to be realized (i.e., when the berries will be ripe). To capture the uncertainty about the realization time, suppose that at any time $t < T$ there is a probability $q \Delta t$ that V will be realized in the interval between t and $t + \Delta t$, i.e., q is the probability density of early

realization.⁹ Thus the total probability that V will be realized before T is qT and the probability that it will be realized at T is $1 - qT$.¹⁰ If the hazard rate (the crow-arrival process) is r , the blackbird will derive an expected payoff of

$$(1) \quad \left(\int_0^T qe^{-rt} dt + (1 - qT)e^{-rT} \right) V$$

from the prospect.

To set the stage for analyzing hyperbolic discounting, let us suppose that, besides the raspberries, there is another prospect $P' = (V', T')$ corresponding to a blackberry bush. The blackberries are far enough away from the raspberries so that the bird can monitor only one bush at a time (if it fails to monitor a bush, the fruit will be all gone before it gets there).¹¹ Blackberries have a higher payoff than raspberries, i.e., $V' > V$ (they have more calories per unit volume). But are likely to ripen later ($T' > T$). As with raspberries, q is the probability density of early ripening.

We can now establish our main result:

Proposition 1: Assume that there exists $t^* (< T)$ at which the DM is indifferent between $P = (V, T)$ and $P' = (V', T')$, with $0 < V < V', T < T'$, and a probability density q of early realization (i.e., realization before T or T'). Then, the DM prefers P' to P at all $t < t^*$, but prefers P to P' at all t such that $t^* < t < T$.

⁹ We should also allow for a probability of realization *after* T . But since this would not affect our analysis in any significant way, we defer this complication to section 4, which also relaxes some of our other assumptions about the probability of early realization.

¹⁰ As we have formulated it, there is a probability atom at $t = T$. This gets at the idea that T is the “likely” realization time. It will be clear from the argument, however, that our results do not literally require atoms; the conclusions continue to hold for continuous densities that are sufficiently “peaked” at T , an assumption that seems to apply quite well to the sorts of problems that blackbirds and certain other kinds of birds typically face.

¹¹ Furthermore, if it monitors, say, the raspberry bush, it can eat its fill of raspberries, once they ripen. And so there is no point in its moving on to the blackberry bush, afterward (and visa versa).

Note: Proposition 1 implies that the DM discounts hyperbolically. When there is sufficient time remaining ($t < t^*$), she is willing to wait for the later reward, but when the horizon has grown short ($t > t^*$) she becomes impatient.

Proof: From (1), the expected net benefit from prospect P , conditional on time t ($< T$) having elapsed (and on neither prospect having yet been realized), is

$$(2) \quad \int_t^T e^{-r\tau} \hat{q}(t) d\tau V + (1 - \hat{q}(t)(T - t)) e^{-rT} V,$$

where

$$(3) \quad \hat{q}(t) = \frac{q}{1 - qt}.$$

Similarly, the conditional expected net benefit from P' is

$$(4) \quad \int_t^{T'} e^{-r\tau} \hat{q}(t) d\tau V' + (1 - \hat{q}(t)(T' - t)) e^{-rT'} V'.$$

The time derivatives of (2) and (4) are respectively

$$(5) \quad -e^{-rt} \hat{q}(t) V + \hat{q}(t) e^{-rT} V + \left(\int_t^T e^{-r\tau} d\tau - (T - t) e^{-rT} \right) V \frac{d\hat{q}}{dt}(t),$$

and

$$(6) \quad -e^{-rt} \hat{q}(t) V' + \hat{q}(t) e^{-rT'} V' + \left(\int_t^{T'} e^{-r\tau} d\tau - (T' - t) e^{-rT'} \right) V' \frac{d\hat{q}}{dt}(t).$$

Because $\frac{d\hat{q}}{dt}(t) / \hat{q}(t) = \hat{q}(t)$, we can rewrite (5) and (6) as

$$\hat{q}(t) \left[\int_t^T e^{-r\tau} d\tau \hat{q}(t) V + (1 - \hat{q}(t)(T - t)) e^{-rT} V - e^{-rt} V \right] \quad (7)$$

and

$$\hat{q}(t) \left[\int_t^{T'} e^{-r\tau} d\tau \hat{q}(t) V' + (1 - \hat{q}(t)(T' - t)) e^{-rT'} V' - e^{-rt} V' \right]. \quad (8)$$

From hypothesis, the sum of first two terms in the square-bracketed expression of (7) equals that of (8) when $t = t^*$. Hence, since $V' > V$, we conclude that (7) exceeds (8) at $t = t^*$.

Furthermore, the same argument implies that (7) exceeds (8) at any other point where the

payoffs from P and P' are equal, and so (2) can cross (4) just once. Hence, because the DM is indifferent between P and P' at $t = t^*$ she prefers P' to P (i.e., (4) exceeds (2)) at $t < t^*$, and prefers P to P' ((2) exceeds (4)) at $t > t^*$.

Q.E.D.

The idea in the proof of Proposition 1 is readily explained informally. The passage of time has two marginal effects on the expected payoff from either prospect P or P' : (i) it brings nearer the time (T or T') at which the payoff is likely to be realized (an effect that is proportional to the current expected payoff from the prospect); and (ii) it reduces the probability of early realization (an effect that for P is proportional to V and for P' to V'). Effect (i) is the *same* for prospects P and P' at a time t^* for which their expected payoffs are equal, but effect (ii), which is negative, is bigger for P' than for P (because $V' > V$). Hence, a *single-crossing* property holds: the expected payoff for P' declines faster with time than that for P whenever the prospects' expected payoffs are equal.

There are two ways of interpreting the preference reversal exhibited in Proposition 1. The first is to think of the DM as making a once-and-for-all choice of P' over P at time 0 and then, unexpectedly, having the opportunity to choose again at some time t , in which case the DM will switch to P if $t > t^*$. The other is to imagine that she can switch between prospects at any time, in which case the DM will select P' at time 0 and stick with it until time $t = t^*$, at which point she will switch over to P .

Proposition 1 pertains to positive payoffs, but payoffs in the O'Donoghue-Rabin tax story we began with are negative. Nevertheless, we can straightforwardly reformulate the proposition to handle negativity, thereby accounting, in particular, for the preference reversal in the tax story:

Proposition 1*: Maintain the hypotheses of Proposition 1 except now assume that $V' < V < 0$.

Then, the DM prefers P to P' for $t < t^*$ but P' to P for $t > t^*$.

Proof: Almost identical to that of Proposition 1.

4. More General Uncertainty

In Propositions 1 and 1* it was assumed that (i) all early realization times are equally probable, (ii) the probability density of early realization is the same for both prospects, and (iii) there is no probability of late arrival. Because there is no particular reason why the average problem that an animal or human might face should satisfy these assumptions, we will relax them. We first note, however, that we cannot get by without *any* assumptions on the distribution of arrival times. To see this, consider prospects P and P' with positive payoffs, and imagine that, in some given time interval, the more distant prospect P' has a *zero* probability of early realization, but that the more immediate prospect P has a positive probability. Then, should that interval elapse with no early realization, the updated chance of early realization for P' will not have changed (because there was no chance of realization in that interval anyway), but it will have grown dimmer for P . Hence, P will *not* become relatively more attractive, contrary to the conclusion of Proposition 1.

Nevertheless, hyperbolic discounting holds for a large and, we would argue, fairly natural class of probability distributions. Consider prospects $P = (V, T)$ and $P' = (V', T')$, with $0 < V < V'$ and $T < T'$. For $t \neq T$, let $q(t)$ be the probability density for realization of P at time t . Similarly, for $t \neq T'$ let $q'(t)$ be the probability density for realization of P' at t .

For $t < \tau$, define

$$\hat{q}(\tau, t) = \begin{cases} \frac{q(\tau)}{1 - \int_0^t q(x) dx}, & \tau < T \\ \frac{q(\tau)}{\int_t^\infty q(x) dx}, & T < \tau \end{cases}$$

Thus, $\hat{q}(\tau, t)$ is the density that P is realized at τ , conditional on its not being realized before t . Similarly, for $t < \tau$,

$$\hat{q}'(\tau, t) = \begin{cases} \frac{q'(\tau)}{1 - \int_0^t q'(x) dx}, & \tau < T' \\ \frac{q'(\tau)}{\int_t^\infty q'(x) dx}, & T' < \tau \end{cases}$$

is the corresponding density for P' . Assume

$$(9) \quad \hat{q}(t, t) \geq \hat{q}'(t, t) \quad \text{for all } t < T$$

and

$$(10) \quad V \hat{q}(t, t) < V' \hat{q}'(t, t) \quad \text{for all } t < T.$$

If the density of a prospect's early realization drops off monotonically as we move away from the anticipated realization time (T or T'), then (9) seems reasonable since, because $t < T < T'$, the density should have dropped off less at t for P than for P' .

Assumption (10) serves to rule out the sort of example illustrated in the first paragraph of this section. Because $V < V'$, it accommodates the possibility that the density of early arrival for P' may be strictly less than that for P , but prevents the former from being *too* much less than the latter. For a simple example (beyond that of section 3) satisfying these hypotheses, suppose that the density of early realization increases linearly with time at a rate inversely proportional to the likely realization time, i.e., $q(t) = \frac{\alpha t}{T}$ for $t < T$ and $q'(t) = \frac{\alpha t}{T'}$ for $t < T'$.

Then, for α sufficiently small, (9) and (10) will hold.

Proposition 2: Suppose that (9) and (10) hold. Assume that there exists $t^* < T$ at which the DM is indifferent between P and P' . Then the DM prefers P' to P at all $t < t^*$ and prefers P to P' at all $t \in (t^*, T)$. Hence, hyperbolic discounting obtains.¹²

Proof: The net payoffs from P and P' conditional on reaching time $t (< T)$ are, respectively,

$$(11) \quad \int_t^\infty \hat{q}(\tau, t) e^{-r\tau} d\tau V + \left(1 - \int_t^\infty \hat{q}(\tau, t) d\tau\right) e^{-rT} V$$

and

$$(12) \quad \int_t^\infty \hat{q}'(\tau, t) e^{-r\tau} d\tau V' + \left(1 - \int_t^\infty \hat{q}'(\tau, t) d\tau\right) e^{-rT'} V'.$$

Hence, their time derivatives are respectively

$$(13) \quad -\hat{q}(t, t) e^{-rt} V + \int_t^\infty \frac{\partial \hat{q}}{\partial t}(\tau, t) e^{-r\tau} d\tau V + \hat{q}(t, t) e^{-rT} V \\ - \int_t^\infty \frac{\partial \hat{q}}{\partial t}(\tau, t) d\tau e^{-rT} V$$

and

$$(14) \quad -\hat{q}'(t, t) e^{-rT'} V' + \int_t^\infty \frac{\partial \hat{q}'}{\partial t}(\tau, t) e^{-r\tau} d\tau V' + \hat{q}'(t, t) e^{-rT'} V' \\ - \int_t^\infty \frac{\partial \hat{q}'}{\partial t}(\tau, t) d\tau e^{-rT'} V'.$$

Because

$$\frac{\partial \hat{q}}{\partial t}(\tau, t) = \hat{q}(\tau, t) \hat{q}(t, t) \quad \text{and} \quad \frac{\partial \hat{q}'}{\partial t}(\tau, t) = \hat{q}'(\tau, t) \hat{q}'(t, t),$$

we can rewrite (13) and (14), respectively, as

$$(15) \quad \hat{q}(t, t) \left[\int_t^\infty \hat{q}(\tau, t) e^{-r\tau} d\tau V + \left(1 - \int_t^\infty \hat{q}(\tau, t) d\tau\right) e^{-rT} V \right] - \hat{q}(t, t) e^{-rt} V$$

and

¹² Although Proposition 2 allows for payoff realization times after T and T' , notice that the conclusion about hyperbolic discounting applies only to times before T .

$$(16) \quad \hat{q}'(t,t) \left[\int_t^\infty \hat{q}'(\tau,t) e^{-r\tau} d\tau V' + \left(1 - \int_t^\infty \hat{q}'(\tau,t) d\tau \right) e^{-rT} V' \right] - \hat{q}'(t,t) e^{-rt} V'.$$

The square-bracketed expression in (15) is just (11), while that in (16) is (12). Hence, from (9) and (10), (15) is greater than (16) whenever (11) equals (12), and in particular when $t = t^*$. We conclude that (11) exceeds (12) for $t \in (t^*, T)$ and (12) exceeds (11) for $t < t^*$.

Q.E.D.

5. Dynamic Inconsistency and Self-commitment

We have shown that preference reversals conforming with hyperbolic discounting follow from a model with uncertainty about payoff realization times. That is, the choices that a DM makes at $t < t^*$ reflect a discount rate lower than those at times $t > t^*$. As the model stands so far, however, there is no dynamic inconsistency in these choices, despite the reversals. In particular, notice that the DM at time $t = 0$ would not choose to “tie her hands” in order to prevent herself from switching at $t > t^*$. Thus, the model does not yet explain the self-commitment strategies observed in pigeon experiments or in people trying to save for Christmas.

To get at dynamic inconsistency and self-commitment, let us further develop our evolutionary model as follows: Suppose that, over a long period of time (“evolutionary” time), a species faces prospects of the sort we have been discussing, i.e., those of the form (V, T, q) (where the parameters V , T , and the function q may change from prospect to prospect). Assume that whenever an individual DM of this species faces a particular choice between prospects $P^* = (V^*, T^*, q^*)$ and $P^{**} = (V^{**}, T^{**}, q^{**})$ with

$$(17) \quad 0 < V^* < V^{**} \text{ and } T^* < T^{**},$$

she can observe (V^*, T^*) and (V^{**}, T^{**}) but *not* the functions q^* and q^{**} that determine uncertainty.¹³ Thus, the best that she can do is to choose between P^* and P^{**} according to

¹³ This assumption gets at the quite natural hypothesis that it will ordinarily be easier for a DM to assess the payoff V and the likely realization time T of a prospect than the entire probability distribution of realization

which one maximizes her expected payoff, where the expectations are taken with respect to all possible values of q^* and q^{**} (and the weight given to a particular value of q^* or q^{**} is determined by her species's evolutionary experience). And so evolution should endow her with correspondingly appropriate urges and inclinations. As we have seen in section 4, these instincts, provided that (9) and (10) hold, will induce her to behave as though she discounts hyperbolically (in that section, the argument proceeded as though q^* and q^{**} were drawn from point distributions, but it is not hard to see that (9) and (10) and the proof of Proposition 2 can readily be generalized to handle diffuse distributions). Let us suppose in particular, that they cause her to switch from P^{**} to P^* after some time has elapsed.

Assume now that for some particular choice, q^* and q^{**} are substantially different from the average values. To be concrete, let us suppose that

$$(18) \quad q^* = q^{**} \equiv 0$$

(as is the case, for example, in the pigeon experiments we will discuss below). Notice that if the DM *knew* that (18) held, then it would not be optimal for her to reverse her preferences. The optimum instead would involve her choosing one of the prospects (say, P^{**}) and sticking with it;¹⁴ preference reversal here would be dynamically inconsistent. Yet, because she does not observe q^* and q^{**} , she will, in fact, exhibit preference reversal in a once-off choice between P^* and P^{**} , despite the dynamic inconsistency of this behavior.¹⁵

Still, in addition to urges, evolution equips many species with the capacity to *learn* from recurrent events so that animals can adjust their behavior accordingly. Following this idea, let us consider a DM who confronts the same choice between P^* and P^{**} (with (17) and

times.

¹⁴ If she knew that (18) held, the DM would choose P^{**} over P^* provided that $V^{**} e^{-rT^{**}} > V^* e^{-rT^*}$. But this inequality implies that $V^{**} e^{-r(T^{**}-t)} > V^* e^{-r(T^*-t)}$ for all t , and so she would continue to choose P^{**} at any $t > 0$.

¹⁵ This underscores the importance of the assumption that the uncertainty cannot be directly observed. If it *were* observable, then evolution ought to endow the DM with the ability to make the right choice, which would rule out dynamic inconsistency.

(18) holding) many times, and suppose that choosing P^{**} (and sticking with it) is the optimal choice. The simple learning mechanism that we have in mind is one in which, for each recurrence, the DM chooses a behavior *probabilistically*, where behaviors that have been successful in the past (i.e., have generated higher payoffs) have higher probabilities of being chosen, provided that she does not have a strong urge to behave in contradictory fashion.

Let us examine how this sort of learning grafted onto the evolutionary model of sections 3 and 4 applies to the large body of experimental work pioneered by Rachlin and Green (1972) and Ainslie (1974). In a stylized rendition of those experiments, a pigeon got to make a choice (say, by pecking either of two keys) at time t_1 between prospects $P^* = (V^*, T^*, q^*)$ and $P^{**} = (V^{**}, T^{**}, q^{**})$ satisfying (17) and (18) (the payoffs V^* and V^{**} consisted of food). It then had the opportunity to change its decision and switch to the other prospect at a later time $t_2 (< T^*)$. There was also a third key that, if pecked at time t_1 , disabled the opportunity to switch prospects at time t_2 . The pigeon repeated the experiment many times.

Not surprisingly, there was no agreement on the “right” t_1 - choice: some pigeons chose P^* and others chose P^{**} . However, roughly speaking, the general findings were: (i) the typical pigeon quickly settled down and consistently made the *same* t_1 choice in most rounds; (ii) among those pigeons usually choosing P^* at t_1 , relatively few switched to P^{**} at t_2 , whereas a good many that chose P^{**} switched to P^* at t_2 ; (iii) relatively few pigeons pecked the disabling key often in early rounds; (iv) among those that eventually did peck the disabling key regularly in later rounds, most were those that had often switched from P^{**} to P^* at t_2 in early rounds, (v) among those that often switched from P^{**} to P^* at t_2 in early rounds, a large fraction eventually pecked the disabling key and chose P^{**} in later rounds.

Findings (i)-(v) accord quite well with our evolutionary/learning model. Finding (i) indicates only that subjects seemed to have stable preferences, which, in particular, would be true if they were maximizing expected discounted calories (as the theoretical model presumes). Turning to (ii), we note that a switch from P^{**} to P^* (which corresponds to hyperbolic discounting) is consistent with the theoretical model, whereas a switch from P^* to P^{**} is not. Thus (ii) is predicted by theory.

The finding that relatively few birds invoked the disabling key early on (iii) also is to be expected theoretically. As we have seen, switching may well be optimal for “typical” uncertainty but not when (18) holds. Nevertheless, it would take repeated experience for a bird to discover that it is not facing such uncertainty, and thus the bird should not, according to the model, regularly disable the switching option until it has had such experience; that is just what appears to have happened.

As for (iv) and (v), notice that it is precisely the birds that initially switch from P^{**} to P^* (in response to their urge for increasing impatience) that theoretically should ultimately disable the switching option (so as to thwart this urge to switch to P^*).¹⁶ Hence, once again, theory is in harmony with experiment.

We conclude that a model in which (a) preferences evolve to handle “typical” uncertainty and (b) a bird can learn to modify its behavior when facing a recurrent atypical problem seems to do quite a good job of explaining observed patterns of preference reversal and self-commitment in pigeons.

Of course, pigeons are only one type of bird, and even if one adds starlings (for which there is also a substantial body of experimental data) and humans, there are still only a few

¹⁶ Another conceivable evolutionary possibility would presumably have been the creation of a mechanism to “turn off” the urge toward greater impatience after the bird acquired evidence that such impatience led to dynamic inconsistency. Why we don’t seem to observe such mechanisms in reality is something of a puzzle. Perhaps turning off an urge, once it has arisen in a particular setting, is too difficult to arrange physiologically. Alternatively, the ability to learn to commit oneself not to succumb to impatient urges might be a good enough substitute in practice for a mechanism that makes such urges conditional on what the bird has learned (presumably a mechanism that is so highly contingent would be quite costly physiologically).

kinds of animals whose tendencies toward dynamic inconsistency we know very much about. Indeed, it is certainly possible that there are species that do not exhibit appreciable preference reversal at all. At least on a superficial level, this would be entirely consistent with our model: although the hypotheses of Proposition 2 are fairly weak, they need not *always* hold, and, therefore, could be violated for the typical decision problems faced by some species. But, for a proper test of the model, one would want to dig deeper and examine those decision problems explicitly. We hope that future research will do just that.

References

- Ainslie, G (1974), "Impulse Control in Pigeons," Journal of the Experimental Analysis of Behavior, 21, 485-487.
- Ainslie, G. (1992), Picoeconomics: The Strategic Interaction of Successive Motivational States with the Person (Cambridge: Cambridge University Press).
- Bernheim, B.D., J. Skinner, and S. Weinberg (2001), "What Accounts for the Variation in Retirement Wealth Among U.S. Households?", American Economic Review, 91, 832-57.
- Dasgupta, P. and E. Maskin (2002), "Uncertainty, Waiting Costs, and Hyperbolic Discounting," mimeo.
- Green, L. and J. Myerson (1996), "Exponential Versus Hyperbolic Discounting of Delayed Outcomes: Risk and Waiting Times", American Zoologist, 36, 496-505.
- Harris, C. and D. Laibson (2001), "Hyperbolic Discounting and Consumption", mimeo.
- Holling, C.S. (1959), Some Characteristics of Simple Types of Predation and Parasitism", Canadian Entomologist, 91, 385-398.
- Kacelnik, A. (1997), "Normative and Descriptive Models of Decision Making: Time Discounting and Risk Sensitivity", in G.R. Bock and G. Cardew, eds., Characterizing Human Psychological Adaptations (Chichester, UK: Wiley).
- Laibson, D. (1997), "Golden Eggs and Hyperbolic Discounting", Quarterly Journal of Economics, 112, 443-477.
- Mazur, J.E. (1987), "An Adjusting Procedure for Studying Delayed Reinforcement", in M.L. Commons, J.E. Mazur, J.A. Nevin, and H. Rachlin, eds., Quantitative Analyses of Behaviour, Vol. 5: The Effect of Delay and of Intervening Events on Reinforcement Value (Hillsdale, NJ: Lawrence Erlbaum).
- O'Donoghue, T. and M. Rabin (1999), "Doing It Now or Later", American Economic Review, 89, 103-124.
- Rachlin, H. and L. Green (1972), "Commitment, Choice, and Self-control ", Journal of the Experimental Analysis of Behavior. 17, 15-22.
- Rubinstein, A. (2001), "A Theorist's View of Experiments", European Economic Review, 45, 615-628.
- Samuelson, L. and J. Swinkels (2002), "Information, Evolution, and Utility", mimeo.
- Sozou, P.D. (1998), "On Hyperbolic Discounting and Uncertain Hazard Rates", Proceedings of the Royal Society of London (Series B), 265, 2015-2020.

Strotz, R.H. (1956), "Myopia and Inconsistency in Dynamic Utility Maximization", Review of Economic Studies, 23, 165-180.

Weitzman, M.L. (1979), "Optimal Search for the Best Alternative", Econometrica, 47, 641-654.

Weitzman, M.L. (2001), "Gamma Discounting", American Economic Review, 91, 260-271.

Yaari, M. (1965), "Uncertain Lifetime, Life Insurance, and the Theory of Consumer", Review of Economic Studies, 32, 137-158.