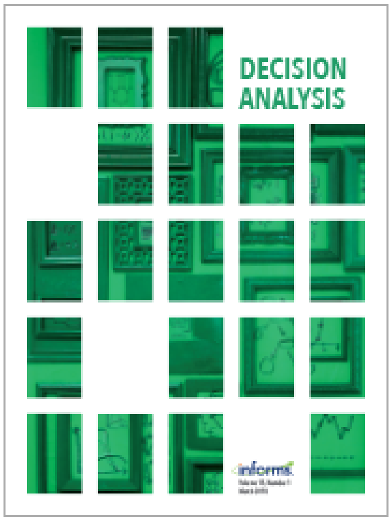

Decision Analysis

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Microfoundations of Discounting

Alexander Adamou, Yonatan Berman, Diomides Mavroyiannis, Ole Peters

http://pubsonline.informs.org/journal/deca

DECISION ANALYSIS
Vol. 18, No. 4, December 2021, pp. 257–272
ISSN 1545-8490 (print), ISSN 1545-8504 (online)

Microfoundations of Discounting

Alexander Adamou,[a] Yonatan Berman,[a] Diomides Mavroyiannis,[b] Ole Peters[a,c]

[a] London Mathematical Laboratory, London W6 8RH, United Kingdom; [b] Milestone Institute, Budapest 1067, Hungary; [c] Santa Fe Institute, Santa Fe, New Mexico 87501

Contact: a.adamou@lml.org.uk, https://orcid.org/0000-0002-4188-7045 (AA); y.berman@lml.org.uk (YB); diomavro@hotmail.com, https://orcid.org/0000-0002-5452-3910 (DM); o.peters@lml.org.uk, https://orcid.org/0000-0002-1246-8570 (OP)

Received: January 28, 2021
Revised: May 15, 2021
Accepted: June 7, 2021
Published Online in Articles in Advance: August 6, 2021

https://doi.org/10.1287/deca.2021.0436

Abstract. An important question in economics is how people choose between different payments in the future. The classical normative model predicts that a decision maker discounts a later payment relative to an earlier one by an exponential function of the time between them. Descriptive models use nonexponential functions to fit observed behavioral phenomena, such as preference reversal. Here we propose a model of discounting, consistent with standard axioms of choice, in which decision makers maximize the growth rate of their wealth. Four specifications of the model produce four forms of discounting—no discounting, exponential discounting, hyperbolic discounting, and a hybrid of exponential and hyperbolic discounting—two of which predict preference reversal. Our model requires no assumption of behavioral bias or payment risk.

Keywords: discounting • decision criteria

1. Introduction

1.1. Background

In economics and psychology, temporal discounting—or, simply, discounting—is a paradigm of how decision makers choose between rewards available at different times in the future. We write here of people and money payments, noting that discounting is also studied in other animals and for other reward types. The basic observation to be explained is that, for two payments of equal size, people prefer typically the earlier payment to the later one. In the discounting paradigm, the later payment is discounted relative to the earlier payment by multiplying it by a factor (the *discount factor* or *discount function*) that depends on the time period between the payments (the *delay*). This operation expresses the later payment as an equivalent payment at the earlier time, to be compared with the earlier payment actually on offer.

Why do people assign lower values to payments further in the future? One plausible answer is that a later payment is less likely to be received, because there is more time for something to go wrong with it. In other words, delay increases risk. Another is that, for equal payments, the later one corresponds to a lower growth rate, which, if sustained over time, would result in being poorer. Modern treatments of discounting in economics tend to follow risk-based reasoning, whereas there is a more even split between risk and rate interpretations in psychology.

This paper studies the microfoundations of discounting using the rate interpretation in a riskless setting. In our model, a decision maker chooses the payment that maximizes the growth rate of her wealth. Our model assumes no behavioral bias and does not violate standard axioms of choice. It predicts a range of situation-dependent discount functions, including those documented in the discounting literature.

This literature abounds with models (Cohen et al. 2020). In some, theoretical considerations are used to construct the decision maker's maximand, from which the discount function is derived. Such models are labeled "normative"—they say what decision makers should do if they want to act in some sense optimally. In other models, the discount function is chosen to fit empirical data, with theoretical justification sought post hoc or not at all. These models are labeled "descriptive"—they predict what decision makers actually do, regardless of whether it is in any sense optimal.

The main normative model in economics is exponential discounting, in which the discount function decays exponentially with the delay (Samuelson 1937). For money payments, this is derived straightforwardly: either by a no-arbitrage argument, assuming payments are guaranteed and the earlier payment can be invested during the delay to earn compound interest at a riskless rate, or by assuming the later payment has a constant hazard rate during the delay and the decision maker maximizes the expected payment (Kacelnik 1997). Exponential discounting is applied to utility-of-consumption flows as well as to money payments (Cohen et al. 2020).[1]

Exponential discounting is not descriptive. Experiments on human and nonhuman animals suggest that payments can be discounted more for shorter delays and less for longer delays than is well described by fitting the rate parameter of an exponential function. Furthermore, subjects exhibit a behavior known as preference reversal, which is incompatible with exponential discounting (Green and Myerson 1996, figure 2). Specifically, subjects switch their preference from the later to the earlier payment as the time to the earlier payment—which we call the *horizon*—gets shorter, while the payments and the delay between them stay fixed (Keren and Roelofsma 1995, experiment 1). The primary evidence against the main normative model is summarized by Myerson and Green (1995) and references therein. In response to its falsification, descriptive models are proposed with discount functions better able to fit observations.

The most widely used descriptive model is hyperbolic discounting, where the discount function is a hyperbola in the delay. The function has one free parameter, known as the degree of discounting, which determines its steepness. Fitting this parameter to experimental data are more efficient than fitting an exponential function, both at group level and for individuals (Myerson and Green 1995). Furthermore, preference reversal is compatible with hyperbolic discounting (Green and Myerson 1996, figure 2). Other descriptive models of discounting are proposed in the literature, such as the quasi-hyperbolic and constant-sensitivity models discussed in Yoon (2020).

Kacelnik (1997) remarks on this divergence between normative and descriptive models, noting that the hyperbola

> *is not strongly explanatory because it did not emerge from an a priori analysis but purely from its power to describe data efficiently. In contrast, because of the strong appeal of the a priori argument favouring exponential discounting, several re-elaborations have been made to rescue the rationale that led to it.*

Most such attempts to adapt the normative model introduce payment risk, which we discuss in Section 1.3. Another approach, favored in behavioral economics, is to present nonexponential discounting as a cognitive bias—a deviation from optimal behavior—to be documented and quantified in mathematical functions that encode human psychology (Loewenstein and Prelec 1992, Laibson 1997).

1.2. Our Model—Growth Rate Maximization

Here we propose a model of temporal discounting, compatible with both exponential and hyperbolic discounting, in which neither payment risk nor behavioral bias is assumed. Instead, we make the single behavioral assumption that decision makers act to maximize the growth rate of their wealth. We study what we call the *riskless intertemporal payment problem* (*RIPP*), in which a decision maker chooses between two known, certain payments at known, certain, future times. When confronted with a RIPP, a decision maker in our model compares the growth rates of wealth arising from each payment and chooses the payment with the higher growth rate.

Our model is normative in that it specifies a maximand, the growth rate of wealth, and states what decision makers should do to act optimally with respect to it. In other words, we are exploring a model world in which decision makers are assumed to be growth rate maximizers, and any statements we make about the behavior of decision makers in our model are conditioned on that assumption. We do not claim here that real people should always or do always act to maximize the growth rate of wealth. The extent to which they do will, in part, determine whether our model makes good predictions about real behavior, that is, whether it is descriptive.

The RIPP is underspecified, in that additional information about the decision maker and her circumstances is required to define and compute the growth rates of her wealth. In our model, in which growth rates are maximized, the RIPP becomes fully specified by knowing

- the decision maker's existing wealth;
- whether and how payment amounts depend on wealth—the *wealth dynamics*;
- how long the decision maker must wait until the next decision—the *time frame*.

In other models, such as expected discounted utility, different information is needed to resolve the underspecification of the RIPP. We discuss examples from the literature in Section 1.3.

The additional information is important. Choosing between wealth-independent payments, such as payments for labor, is palpably different from choosing between payments proportional to wealth, like investment income. In our terminology, this is the difference between additive and multiplicative dynamics, which are the two cases we treat in this paper. Likewise, choosing an earlier payment after which one is immediately free to pursue another opportunity is different from choosing an earlier payment after which one must wait until the entire problem has played out. We analyze these two time frames in our model.

We stress that a real decision maker cannot choose her wealth, wealth dynamics, and time frame. They are imposed on her, and on us as modelers, by her personal circumstances and the nature of the payments offered. The model components encoding these external facts should not, therefore, be viewed as arbitrary degrees of freedom. Rather, they are chosen to match real and relevant circumstances, which are unspecified in the RIPP. The resulting model is predictive: for given wealth, wealth dynamics, and time frame, it predicts what a decision maker will do if our assumption that she maximizes her growth rate is realistic.

Depending on the specification, our model predicts four different forms of discounting: no discounting, exponential, hyperbolic, and a hybrid of exponential and hyperbolic. This is not an exhaustive list—other dynamics would produce other forms of discounting. Two of the discount functions nested in our model, hyperbolic and hybrid, are compatible with preference reversal.

The main contribution of this paper is the unification of hitherto disparate discount functions within a single normative model, growth rate maximization, that does not violate standard axioms of choice. Notably, both exponential and hyperbolic discounting, the main normative and descriptive models in the discounting literature, appear as special cases of our model. We assume neither bias nor dynamic inconsistency (Cohen et al. 2020, p. 3) in the decision maker's behavior, because at all times she prefers the option with the highest growth rate. In some specifications, this translates into a reversal of preference between *payments* (not growth rates). In our perspective, this reflects a change of circumstances and not of mind. The fundamental preference for faster growth never reverses.[2]

Moreover, this paper marks a shift from psychological to circumstantial explanations of discounting. We predict that changes in the discount function arise from changes in wealth dynamics and time frame, which are properties not of the decision maker but of her circumstances. When these circumstances are included in the formalism, a single behavioral model is capable of predicting a range of observed behaviors. Because psychological risk preferences, encoded in idiosyncratic utility functions, do not appear in our model, we sidestep recent controversy in the literature about the suitability of experiments involving money payments to test models of utility-of-consumption flows (Cohen et al. 2020). Such experiments are able to falsify our model and are planned.

This paper is organized as follows. Section 2 sets out the RIPP and our decision model. In Section 3, we present different specifications of the problem and describe how a decision maker discounts future payments in each of them. We conclude in Section 4.

1.3. Related Literature

This paper follows the tradition of adapting the normative discounting model to make it consistent with observations, that is, to make it descriptive

(Kacelnik 1997). Our strategy is to postulate the growth rate of wealth as the decision maker's maximand. Computing this requires information about wealth, wealth dynamics, and time frame, about which the RIPP is silent.

Another strategy is to leave the maximand—usually an expected payment or utility gain—unchanged, and to introduce uncertainty in the amount or timing of the payment. The uncertainty is chosen so that the effective discount function fits the data. Adding risk can be viewed as another way of resolving the underspecification of the RIPP. This approach is questionable for two reasons. First, although the RIPP is silent about dynamics and time frame, it is explicit about uncertainty: the choice presented is between certain payments at certain times. Second, dynamics and time frame are always relevant in real decisions, but real risks can be negligible. Payment risk is a sound microfoundation only when the uncertainty absent from the problem formulation is important in reality. Such situations are plausible. A predator declining a small but easy prey to search for something more filling risks catching nothing at all. However, adding payment uncertainty to generate, say, hyperbolic discounting is not a general prescription. It fails when the real payment risk is negligible, as envisaged in the RIPP and presented in experiments, for example, Myerson and Green (1995).

Furthermore, highly specific forms of uncertainty or other adaptations are required to recover the hyperbola. Green and Myerson (1996) point out that an expected payment model with a hyperbolic hazard rate predicts hyperbolic discounting. They note also that the exponential function can be made consistent with observed behavior, including preference reversal, by allowing its rate parameter to vary with payment size. Sozou (1998) treats an expected payment model with an uncertain hazard rate, about which the decision maker learns through Bayesian updating. If the prior distribution of the hazard rate is exponential, then hyperbolic discounting is again obtained. Dasgupta and Maskin (2005) also assume risk neutrality but keep the hazard rate constant. To produce hyperbolic discounting, they introduce the possibility of payment occurring before the anticipated time.

Such adaptations lead to a loss of generality. They make statements of the following type: "if there is uncertainty in a future payment, and if it takes this particular form, then the discount function is a hyperbola." Such ad hoc assumptions reduce the generality of risk-based models further, because they are useful only when the real risk is of a particular nature.

The strategy we follow leaves payment certainty alone and changes the maximand. It is long established in biology and psychology that hyperbolic discounting is consistent with maximizing the rate of change of resources in a model of additive increments. This insight, traced back to the predation model of Holling (1959), is recognized in studies of human discounting (Myerson and Green 1995, Kacelnik 1997, Sozou 1998). Kacelnik (1997, figure 2) offers a pictorial representation, similar to ours in Section 3, of how rate maximization predicts preference reversal. In the rate interpretation, the degree of discounting is no longer a free psychological parameter. Rather, it is constrained to be the reciprocal of the horizon (Myerson and Green 1995). This is a testable prediction.

We extend this strand of the literature by setting the decision maker's maximand as the growth rate of resources under realistic dynamics, specifically additive and multiplicative. Furthermore, we allow the time frame of the decision—over which growth rates are computed—to be the period from the decision to either the chosen or the later payment. This captures circumstances akin to opportunity costs, specifically whether receiving the earlier payment frees the decision maker to pursue other payments.

Our work contributes to the growing field of ergodicity economics (Peters 2019) in which decision makers maximize the long-time growth rate of resources, rather than expectation values of psychologically transformed resource flows (possibly under psychologically transformed probability measures). This study joins recent evidence of strong dependence on wealth dynamics of human decisions under uncertainty (Meder et al. 2019) and may be used to design similar experiments without uncertainty.

2. Theoretical Framework

2.1. Problem Definition

We begin by formalizing a basic riskless temporal choice problem, where discounting is used to express a later payment as an equivalent payment at an earlier time. We define a riskless intertemporal payment problem as follows.

Definition 1 (Riskless Intertemporal Payment Problem).
A riskless intertemporal payment problem is a comparison between two tuples, $a \equiv (t_a, \Delta x_a), b \equiv (t_b, \Delta x_b)$. A decision maker at time t_0 with wealth $x(t_0)$ must choose between two future cash payments, whose amounts and payment times are known with certainty. The two options are

1. an earlier payment of Δx_a at time $t_a > t_0$;
2. a later payment of Δx_b at time $t_b > t_a$.

A criterion for choosing a or b is required. Here we explore what happens if that criterion is maximization of the growth rate of wealth.

2.1.1. Growth Rates and Dynamics. A growth rate is defined as the scale parameter of time in the growth function of wealth subject to dynamics. Wealth dynamics are, in essence, the rules by which wealth evolves from one point in time to another. They can be defined in several ways, most commonly using ordinary differential equations where uncertainty is negligible and stochastic differential equations where it is not.

Here we treat a riskless decision problem, the RIPP, in order not to confound temporal choice with choice under uncertainty. Therefore, we would write down an ordinary differential equation for wealth or, equivalently and more simply, the deterministic growth function,

$$x(t) = f(gt),$$

obtained by integrating the ordinary differential equation (assuming a closed-form solution exists). Here f denotes the growth function and g the growth rate, which is the scale parameter of time. Different dynamics correspond to different growth functions and, therefore, to different forms of growth rate.

It is the modeler's task to choose dynamics that resemble the real circumstances of the decision maker. Here we treat multiplicative and additive dynamics (Peters and Gell-Mann 2016), whose wealth flows resemble the common scenarios of investment and labor income, respectively. In the presence of uncertainty, these dynamics would be modeled as geometric Brownian motion and Brownian motion, used widely in applications in finance and operations research. Indeed, we are aware of recent work using growth-optimal preferences under these dynamics, that is, the approach presented here, to solve the planning problem of investing in energy efficiency (Britto et al. 2021). Where realism requires more complex wealth dynamics, these can be treated using the general framework in Peters and Adamou (2018); compare Carr and Cherubini (2020) in finance.

2.1.2. Growth Under Multiplicative Dynamics. Ignoring, for the moment, the possible payments Δx_a and Δx_b, a common assumption is that wealth grows exponentially in time. We label this dynamic as multiplicative. It corresponds to investing wealth in income-generating assets, where the income is proportional to the amount invested. In the riskless case, which approximates an investment in safe bonds or a deposit in a safe savings account, wealth grows as

$$x(t) = x(t_0)e^{g(t-t_0)},$$

and the scale parameter of time in the exponential function is g. This growth rate, g, resembles an interest rate or a rate of return on investment. Its dimension is the reciprocal of time, for example, 5% per year.

2.1.3. Growth Under Additive Dynamics. Another possibility is additive dynamics, where wealth grows linearly in time. This corresponds to situations where investment income is negligible and wealth changes by net flows that do not depend on wealth itself. In the riskless case—resembling, for example, net income from employment at a known salary with stable consumption costs—wealth grows as

$$x(t) = x(t_0) + g(t - t_0),$$

and the scale parameter of time in the linear function is g. This growth rate, g, resembles a net savings rate from additive income. It is measured in units of currency per unit of time, for example, $5,000 per year.

The functional form of the growth rate differs between the dynamics. The growth rate between time t and $t + \Delta t$ can be extracted from the expression for the evolution of wealth over that period. Under multiplicative dynamics it is

$$g = \frac{\ln x(t+\Delta t) - \ln x(t)}{\Delta t},$$

and under additive dynamics it is

$$g = \frac{x(t+\Delta t) - x(t)}{\Delta t}.$$

The matching of growth rate with dynamics is crucial. An additive growth rate computed for wealth following a multiplicative process would vary with time, as would a multiplicative growth rate computed for additively growing wealth. The correct growth rate extracts a stable parameter from the dynamics, allowing processes with the same type of dynamics to be compared.

Given the wealth dynamics, a RIPP implies two growth rates: g_a, associated with option a, and g_b, associated with option b. This permits defining growth-optimal behavior:

Definition 2 (Growth-Optimal Preferences).

The preference relation $\succcurlyeq$ is growth optimal if, given the wealth dynamics, a decision time t_0, an initial wealth $x(t_0)$, and payments $a \equiv (t_a, \Delta x_a)$ and $b \equiv (t_b, \Delta x_b)$,

1. $a \succ b$ [a is preferred to b] if and only if $g_a > g_b$;
2. $a \sim b$ [indifference between a and b] if and only if $g_a = g_b$;
3. $a \prec b$ [b is preferred to a] if and only if $g_a < g_b$.

In words, Definition 2 states that a growth-optimal decision maker prefers option a if her wealth grows faster under this option than under option b, and vice versa. She is indifferent if the growth rates are equal. Definition 2 satisfies the weak axiom of revealed preference: completeness is satisfied by design, and it satisfies transitivity, as preference is fully determined by a comparison of scalar growth rates.[3]

2.2. Model Setup

Figure 1 illustrates a RIPP, corresponding to the basic question that arises in temporal discounting, for example, "would you prefer to receive \$100 tomorrow or \$200 in a month's time?" Growth rates depend on time increments, not times themselves, so it is useful to define the two fundamental time increments in the problem: the period from the decision to the earlier payment, called the *horizon*,

$$H \equiv t_a - t_0,$$

and the period between the payments, called the *delay*,

$$D \equiv t_b - t_a.$$

We adopt the standard definition of the discount factor, δ, as the multiplicative factor that, when applied to a later payment, produces an equivalent "discounted" payment at an earlier time. By equivalent, we mean that a decision maker is indifferent between a payment of Δx at the later time and a payment of $\delta \Delta x$ at the earlier time.

In general, the discount factor depends on the parameters of the RIPP. Of primary interest in discounting is its dependence on the delay, D, because this expresses how the perceived value of a later payment diminishes as the time to its receipt lengthens. When conceived as a function of the delay, $\delta(D)$, the discount factor is commonly known as the *discount function*. If the discount factor depends additionally on other parameters in the RIPP, then a family of discount functions exists, one for each set of values of the additional parameters. In such cases, we write, for example, $\delta(D;H)$ to indicate additional dependence on the horizon.

Clearly, the discount function must satisfy $\delta(0) = 1$. Typically, discounting models find that $0 < \delta(D) < 1$ for $D > 0$, expressing that earlier payments are always preferred to later payments, ceteris paribus.

Multiplication by the discount function allows two payments separated by a delay to be compared as if there were no delay between them. This simplifies the decision to choosing the larger discounted payment. In our RIPP, therefore, $\delta(D)$ has the properties that $\Delta x_a > \delta(D)\Delta x_b$ when $a \succ b$, $\Delta x_a < \delta(D)\Delta x_b$ when $a \prec b$, and $\Delta x_a = \delta(D)\Delta x_b$ when $a \sim b$.

We determine the discount function in our model as follows. For each model specification, we find a RIPP in which the decision maker with growth-optimal preferences is indifferent, $a \sim b$. Indifference implies that $\Delta x_a = \delta(D)\Delta x_b$, from which we obtain the discount function as

$$\delta(D) \equiv \left.\frac{\Delta x_a}{\Delta x_b}\right|_{a\sim b}, \tag{1}$$

where the notation $|_{a\sim b}$ means "evaluated under the condition $a \sim b$." In some model specifications, the discount function will depend on other parameters in the problem, such as the horizon, payment sizes, initial wealth, and background growth rate.

Solving the temporal choice problem requires additional assumptions. To compute growth rates two assumptions are needed. The first concerns the *dynamics* under which the decision maker's wealth grows, as discussed in Section 2.1. This determines the appropriate form of the growth rate. The second assumption concerns what we call the *time frame* of the decision, specifically whether a decision maker receiving the earlier payment at t_a is free immediately to make her next

Figure 1. The Basic Setup of the Model

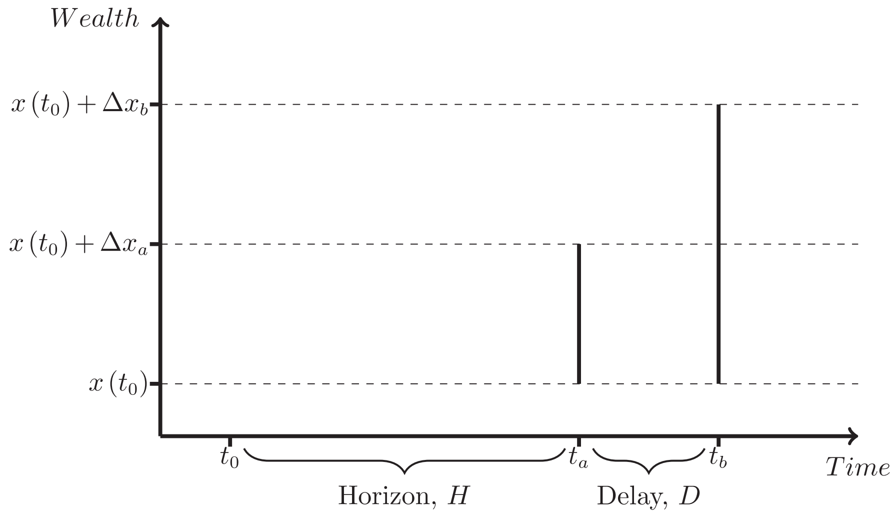

Notes. A decision maker faces a choice at time t_0 between option a, which guarantees a payment of Δx_a at time t_a, and option b, which guarantees a payment of $\Delta x_b > \Delta x_a$ at time $t_b > t_a$. We define the time between the decision and the earlier payment as the *horizon*, $H \equiv t_a - t_0$, and the time between the two payments as the *delay*, $D \equiv t_b - t_a$.

decision or whether she must wait until the later time t_b. This determines the appropriate time period for computing the growth rate under each option. Full specification allows the decision maker's maximand—the growth rate of her wealth—to be evaluated and her options compared. For concreteness, we confine our attention to $\Delta x_b > \Delta x_a$, which is the most commonly considered dilemma.

We describe four different specifications of this basic setup. In each we calculate the growth rates of wealth, g_a and g_b, associated with options a and b. Indifference, $a \sim b$, corresponds to equality of the growth rates, $g_a = g_b$. From this analysis we infer the discount function using Equation (1). The four specifications predict four forms of discounting—no discounting, exponential, hyperbolic, and a hybrid of exponential and hyperbolic—and, in some cases, preference reversal.

We stress that model specification in practice is not done arbitrarily. Neither decision maker nor modeler can choose the decision maker's initial wealth, wealth dynamics, and time frame. They are imposed on her by the reality of her personal circumstances and of the payments being offered.

2.3. Elicitation Procedures

In this study, we propose a behavioral model—growth optimality—and make predictions about how payments are discounted (encoded in discount functions) by analyzing a decision maker's preference in a binary choice between payments. These predictions can be compared with observations of discounting elicited from real decision makers by confronting them with similar choices.

However, a choice between payments is not the only problem involving discounting with which a decision maker can be faced. Frederick (2003) provides an overview of elicitation procedures used in the context of life-saving policy decisions, extended recently by Guyse et al. (2020). Guyse and Simon (2011) discuss similar procedures for decisions involving monetary payments, as in the present study. As well as *choice*, analyzed explicitly here, the following elicitation procedures are relevant to our work:

- *matching*, where a decision maker must name the amount of a payment which would make them indifferent between it and another specified payment;
- *rating*, where a decision maker must quantify how much better or worse one payment is than another; and
- *sequence*, where a decision maker must choose between sequences of payments.

Under our approach, a payment has its desirability quantified by its implied growth rate. This quantification allows predictions to be made in the choice paradigm, simply by determining which payment has the

Figure 2. The Four Model Specifications

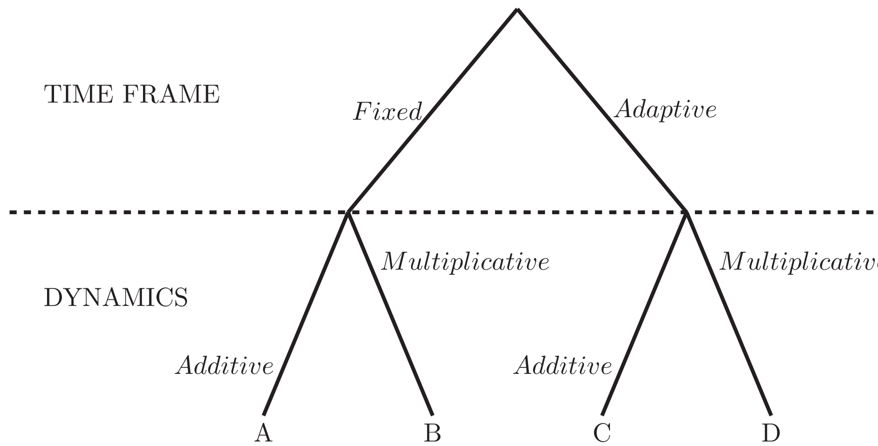

Notes. Each specification is determined by the time frame (fixed or adaptive) and wealth dynamics (additive or multiplicative). For ease of reference, each case is given a label: A, B, C, or D.

higher growth rate. Extension to the matching and rating paradigms is straightforward: one payment is matched with another by making their growth rates equal, and one payment can be rated relative to another by the ratio of the growth rates.

The sequence paradigm is particularly interesting from our perspective, because it confronts the decision maker with a choice between streams of outcomes; compare Lichtendahl and Bodily (2012) for health and consumption decisions. In a choice between two single payments, we must make assumptions about how each will be repeated, in order to infer future trajectories and compute growth rates. This is unnecessary if the trajectory is already specified, and converting the sequence of payments into a growth rate is simple. Because it obviates the need for additional assumptions in our approach, elicitation by sequence may provide the most direct validation of our model.[4]

3. Results

3.1. Specification

We begin by describing the four different model specifications. Each specifies two aspects necessary to quantify the growth rate of wealth: the time frame of the decision and the dynamics under which wealth evolves.

3.1.1. Time Frame. A key aspect, often left unspecified or implicit in the literature, is whether receiving the earlier payment frees the decision maker to pursue the next payment. We treat this by specifying the time frame of the decision, which we illustrate with the following scenarios:

1. Denise is a day laborer. Every evening she looks at a job market forum and chooses a job for the next day. Jobs pay different wages and take different amounts of time, although always less than a day. Denise is paid as soon as she completes the job and goes home. She cannot do more than one job each day.

2. Fiona is a freelancer. She works on projects ranging from a few days to many months and she can work on only one project at a time. As soon as she finishes a project, she gets paid and can move on to the next project.

In the first scenario, the important element to note is that no matter which choice is made, it does not affect the timing of future choices. Denise's next decision is always made the next evening. The time frame is independent of the choice, so we say it is *fixed*.

In the second scenario, the time frame depends on the choice made. The timing of Fiona's next choice is determined by her current decision, for example, choosing a shorter project frees her sooner for the next opportunity. We call this the *adaptive* time frame.

In our model, we must compute the growth rate of wealth using the time period over which the growth rate is effective. This is the time period between successive decisions. With a fixed time frame, it is the period from decision to later payment, that is, $t_b - t_0 = H + D$. In Denise's case, this is one day. With an adaptive time frame, it is the period from decision to chosen payment, that is, $t_a - t_0 = H$ for option a and $t_b - t_0 = H + D$ for option b. This specification is appropriate to Fiona's situation.

3.1.2. Dynamics. As described in Section 2, the wealth dynamics can also take different forms. We address two common cases: additive and multiplicative wealth dynamics. We note that under the multiplicative dynamics, it is assumed that the payment itself is reinvested at a background exponential growth rate, r. For additive dynamics, there is essentially no reinvestment of the payment. Income in this dynamic is independent of wealth.

We discuss the four specifications, as illustrated in Figure 2. In each case, we compute the growth rates g_a and g_b associated with each option, compare them to determine the conditions under which each option is preferred, determine whether preference reversal is predicted, and, finally, elicit the form of temporal discounting equivalent to our decision model.

3.2. Case A—Fixed Time Frame with Additive Dynamics

In this case, the period for computing the growth rate is that between the decision and the later payment, $t_b - t_0 = H + D$, and wealth dynamics are additive. Here, irrespective of the initial wealth and in addition to the chosen payment, wealth grows at a background additive growth rate, k. In other words, the decision maker always receives an amount $k(t_b - t_0)$ over the period, on top of what else she chooses.

In the economic context, we might interpret this sinecure-like income as coming from a job or welfare scheme that does not interact with the payments under consideration. In biological contexts, we might interpret negative k as a metabolic rate of energy expenditure. Provided it takes the same value under the two choices, k appears in neither preference criterion nor discount function. We include it only for completeness.

We begin by writing down the final wealth under each of the two options, evaluated at t_b:

$$x_a(t_b) = x(t_0) + k(t_b - t_0) + \Delta x_a;$$
$$x_b(t_b) = x(t_0) + k(t_b - t_0) + \Delta x_b.$$

The growth rates are

$$g_a = \frac{x_a(t_b) - x(t_0)}{t_b - t_0} = \frac{\Delta x_a}{H + D} + k;$$
$$g_b = \frac{x_b(t_b) - x(t_0)}{t_b - t_0} = \frac{\Delta x_b}{H + D} + k.$$

Because $\Delta x_b > \Delta x_a$, option b is always preferred to option a. This is a trivial case: under additive wealth dynamics and comparing growth rates over the same time period, only payment size matters. There is no discounting and the discount function δ is undefined, because the indifference condition is never satisfied.

3.3. Case B—Fixed Time Frame with Multiplicative Dynamics

In this case, the period for computing the growth rate is that between the decision and the later payment, $t_b - t_0 = H + D$, and wealth dynamics are multiplicative. This specification corresponds to the classical temporal discounting, where wealth compounds continuously at the background multiplicative growth rate, r, and payments are reinvested at this rate.

The earlier payment, Δx_a, if chosen, is treated as growing exponentially from its receipt at t_a to t_b. Therefore, the wealths evolve from t_0 to t_b as follows:

$$x_a(t_b) = x(t_0)e^{r(t_b - t_0)} + \Delta x_a e^{r(t_b - t_a)};$$
$$x_b(t_b) = x(t_0)e^{r(t_b - t_0)} + \Delta x_b.$$

The corresponding growth rates are[5]

$$g_a = \frac{\ln x_a(t_b) - \ln x(t_0)}{t_b - t_0} = \frac{1}{H + D} \ln\left(1 + \frac{\Delta x_a e^{rD}}{x(t_0)e^{r(H+D)}}\right) + r;$$
$$g_b = \frac{\ln x_b(t_b) - \ln x(t_0)}{t_b - t_0} = \frac{1}{H + D} \ln\left(1 + \frac{\Delta x_b}{x(t_0)e^{r(H+D)}}\right) + r.$$

The discount function is obtained by setting the growth rates to be equal, which happens when $\Delta x_a e^{rD} = \Delta x_b$. This yields

$$\delta(D; r) = \frac{\Delta x_a}{\Delta x_b}\bigg|_{g_a = g_b} = e^{-rD},$$

which is the classical exponential discounting result (Samuelson 1937). We see that the discount function depends on a single time period, the delay D, and on the background growth rate of wealth, r.[6] The interpretation is straightforward: if it is possible to reinvest the earlier payment such that it will exceed the later payment amount at the later payment time, then option a is preferred to option b.

3.4. Case C—Adaptive Time Frame with Additive Dynamics

In this case, the period for computing the growth rate is that between the decision and the chosen payment, either $t_a - t_0 = H$ or $t_b - t_0 = H + D$, and wealth dynamics are additive. Like in Case A, regardless of the initial wealth and in addition to the chosen payment, wealth grows at the background additive rate, k. Unlike in Case A, options a and b are evaluated at t_a and t_b, respectively:

$$x_a(t_a) = x(t_0) + k(t_a - t_0) + \Delta x_a;$$
$$x_b(t_b) = x(t_0) + k(t_b - t_0) + \Delta x_b.$$

The growth rates are

$$g_a = \frac{x_a(t_a) - x(t_0)}{t_a - t_0} = \frac{\Delta x_a}{H} + k;$$
$$g_b = \frac{x_b(t_b) - x(t_0)}{t_b - t_0} = \frac{\Delta x_b}{H + D} + k.$$

The discount function is obtained by setting the growth rates to be equal. We have

$$\delta(D; H) = \left.\frac{\Delta x_a}{\Delta x_b}\right|_{g_a = g_b} = \frac{H}{H + D} = \frac{1}{1 + D/H}. \tag{2}$$

Thus, we recover the widely used descriptive model of discounting in which the discount function is a hyperbola of the delay. We note that in this specification, the discount function is parametrized additionally by the horizon, which we indicate by writing it as $\delta(D; H)$ in Equation (2). Indeed, the degree of discounting parameter—usually treated as a psychological parameter—appears in our model as $1/H$, the reciprocal of the horizon (Myerson and Green 1995). As the horizon gets shorter, $1/H$ becomes larger, δ gets smaller, and the later payment becomes less favorable. No knowledge of the decision maker's psychology is required in this setup, other than that she prefers her wealth to grow faster rather than slower.[7]

3.4.1. Preference Reversal.

Preference reversal is the observed change in preference between fixed payments separated by a fixed delay, as some other parameter of the decision problem is varied. The preference reversal reported most commonly in the discounting literature is the switch from preferring the later to the earlier payment as the horizon gets shorter. Typically, this phenomenon is elicited in experiments by presenting subjects with RIPPs with different horizons but otherwise identical parameters, as in for example, Keren and Roelofsma (1995, experiment 1) and references therein.[8] Preference reversal is prima facie puzzling because it is not predicted by exponential discounting, the main normative model of discounting in economics. We test whether it is predicted in our model of growth-optimal preferences by varying H while holding other variables constant.

In the present specification, indifference occurs at a critical horizon, $H = H^{\text{PR}}$, for which $g_a = g_b$, that is,

$$\frac{\Delta x_a}{H^{\text{PR}}} = \frac{\Delta x_b}{H^{\text{PR}} + D} \quad \Rightarrow \quad H^{\text{PR}} = \frac{D \Delta x_a}{\Delta x_b - \Delta x_a}. \tag{3}$$

For $H < H^{\text{PR}}$, the growth rate under option a exceeds that under option b and the earlier payment is preferred. The converse is true for $H > H^{\text{PR}}$. Figure 3 illustrates how the dependence of growth rate on horizon leads to preference reversal under additive dynamics with an adaptive time frame; compare Kacelnik (1997, figure 2).

3.5. Case D—Adaptive Time Frame with Multiplicative Dynamics

In this case, the period for computing the growth rate is that between the decision and the chosen payment, either $t_a - t_0 = H$ or $t_b - t_0 = H + D$, and wealth dynamics are multiplicative.

We follow the same steps as in the previous cases. Wealth evolves to

$$x_a(t_a) = x(t_0)e^{r(t_a - t_0)} + \Delta x_a;$$
$$x_b(t_b) = x(t_0)e^{r(t_b - t_0)} + \Delta x_b.$$

The corresponding growth rates are

$$g_a = \frac{\ln x_a(t_a) - \ln x(t_0)}{t_a - t_0} = \frac{1}{H}\ln\left(1 + \frac{\Delta x_a}{x(t_0)e^{rH}}\right) + r; \tag{4}$$

$$g_b = \frac{\ln x_b(t_b) - \ln x(t_0)}{t_b - t_0} = \frac{1}{H + D}\ln\left(1 + \frac{\Delta x_b}{x(t_0)e^{r(H+D)}}\right) + r. \tag{5}$$

The discount function is obtained by setting the growth rates to be equal. This yields

$$\delta(D; H, r, x(t_0), \Delta x_b) = \frac{x(t_0)e^{rH}}{\Delta x_b}\left[\left(1 + \frac{\Delta x_b}{x(t_0)e^{r(H+D)}}\right)^{\frac{H}{H+D}} - 1\right],$$

which depends on the delay and several other parameters of the problem.

Figure 3. Preference Reversal in Case C

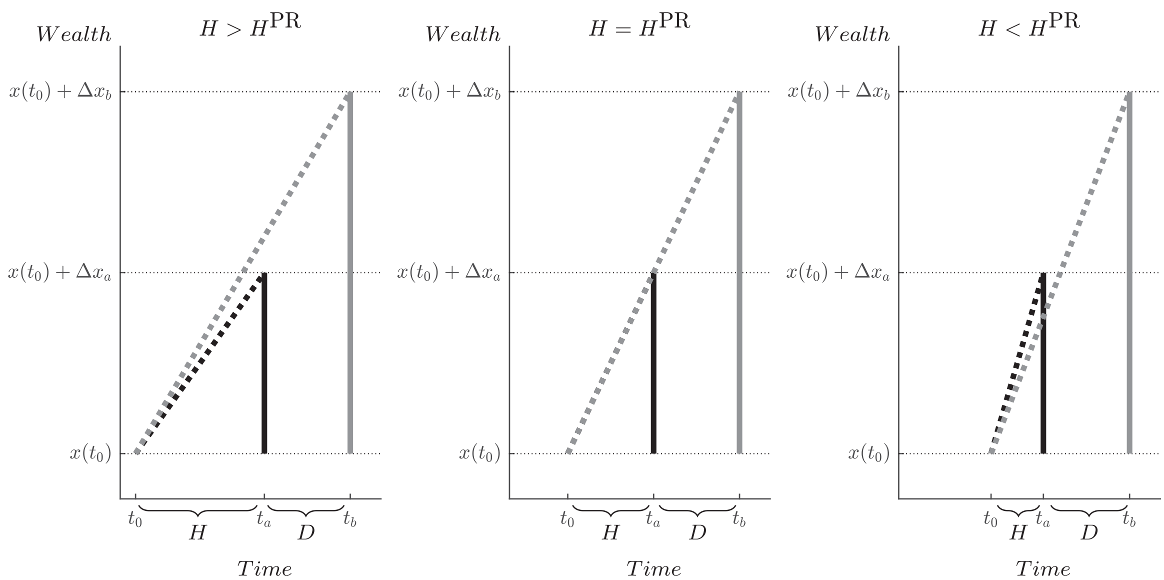

Notes. Horizon H decreases from the left to the right panel, whereas all other parameters are held constant. In the left panel, for long horizons, option b is preferred, having the higher growth rate (slope of dashed line). In the middle panel, at the critical horizon, H^{PR} in Equation (3), both options imply the same growth rate. In the right panel, for shorter horizons, option a has the higher growth rate and is preferred.

3.5.1. Preference Reversal. When the later payment is above some threshold, $\Delta x_b > \Delta x_a e^{rD}$, preference reversal is predicted.[9] The critical horizon is not expressible in closed form for general parameter values, although it becomes tractable in the limit of small payments, which we present below. If the later payment is below the threshold, $\Delta x_b < \Delta x_a e^{rD}$, then the earlier payment is always preferred.

3.5.2. Wealth Effect. Our model predicts another type of preference reversal here, elicited by varying the initial wealth, $x(t_0)$, rather than the horizon. As $x(t_0) \to 0$, the earlier payment is always preferred, regardless of the size of the later payment. Provided the later payment is above some threshold,

$$\Delta x_b > \Delta x_a e^{rD}\left(\frac{H+D}{H}\right),$$

then it becomes preferable to the earlier payment as $x(t_0)$ increases. Thus, the decision maker switches from preferring earlier to later payments as her wealth increases. We call this the *wealth effect*. It is illustrated pictorially in Figure 4, which shows the variation of growth rates, g_a and g_b, for each option as initial wealth, $x(t_0)$, increases from left to right. Other parameters are held constant.

Figure 5 shows the difference in growth rates, $g_a - g_b$, as a function of initial wealth, $x(t_0)$, for the same parameters as in Figure 4. The earlier payment is preferred when this difference is positive, which happens for small enough wealth. For larger wealth, the growth rate difference is negative and the later payment is chosen.

We interpret this as follows. Assuming multiplicative dynamics and an adaptive time frame, it is growth optimal for people of lower wealth to choose an earlier, smaller payment; and growth optimal for wealthier individuals to hold out for the later, larger payment. This is consistent with the findings of Epper et al. (2020, p. 1200), that "relatively patient individuals are consistently positioned higher in the wealth distribution." It is likely consistent with Green et al. (1996), in which people with higher incomes were observed to discount less steeply, because income is positively correlated with wealth.

3.5.3. Sign Effect. Our model predicts another apparent behavioral anomaly, referred to by Frederick et al.

Figure 4. Wealth Effect in Case D

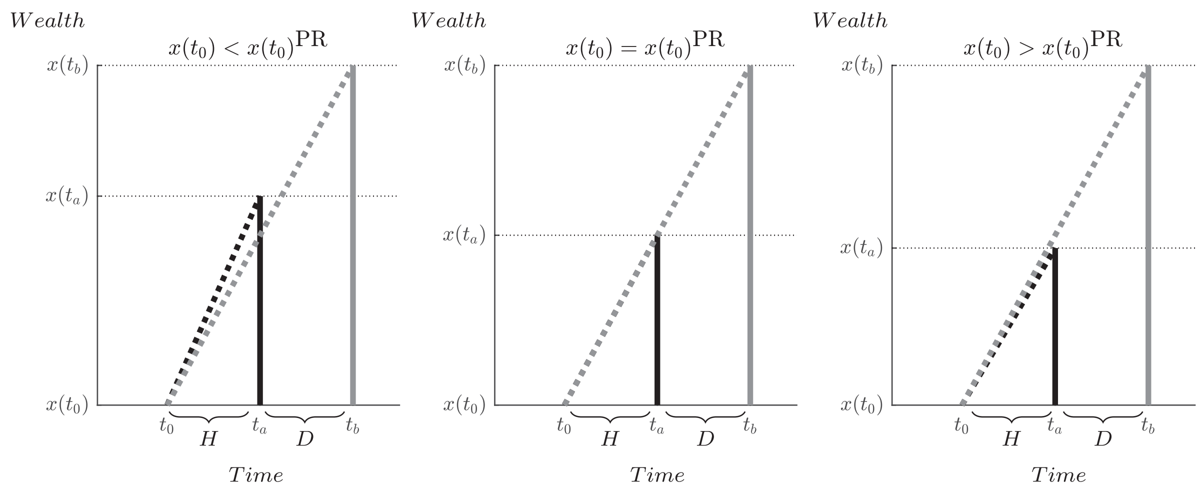

Notes. Initial wealth $x(t_0)$ increases from left panel to right panel (\$500, \$2,277, \$5,500) with all other parameters held fixed (t_0 = today, t_a = 1 year from today, t_b = 2 years from today, $\Delta x_a = \$1{,}000$, Δx_b=\$2,500, $r = 0.03$ per year). At small initial wealth, option *a* is preferred, having the higher growth rate according to Equations (4) and (5). At a critical initial wealth, $x(t_0)^{PR} \approx \$2{,}277$, both options imply equally fast growth and indifference between payments. Preference reversal occurs for larger initial wealth, where option *b* is preferred. The vertical wealth scale is logarithmic.

(2002, p. 360) as the *sign effect*, in which "gains are discounted more than losses". In our framework, this means that the discount function for a RIPP with positive payments, Δx_a and Δx_b, would be lower than for an otherwise identical RIPP with negative

Figure 5. The Difference in Growth Rates, $g_a - g_b$, as a Function of Initial Wealth, $x(t_0)$, in Case D

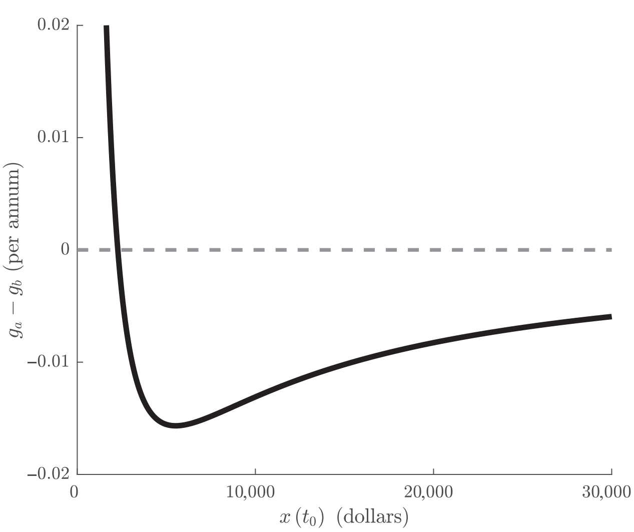

Notes. For small initial wealth the earlier, smaller payment is preferred, whereas for large initial wealth, the later, larger payment is preferred. Parameters are as used in Figure 4.

payments, $-\Delta x_a$ and $-\Delta x_b$. In Cases A–C, the discount function is independent of Δx_a and Δx_b, so the sign effect is not predicted. However, in the present specification, Case D, the discount function depends on Δx_b.

To show that our model recovers the sign effect of Frederick et al. (2002), we define

$$\delta^+ \equiv \delta(D; H, r, x(t_0), \Delta x_b),$$
$$\delta^- \equiv \delta(D; H, r, x(t_0), -\Delta x_b),$$

and ask whether $\delta^+ < \delta^-$, that is, whether gains are more heavily discounted than losses. We study the ratio,

$$\frac{\delta^+}{\delta^-} = -\left[\frac{\left(1+\frac{\Delta x_b}{x(t_0)e^{r(H+D)}}\right)^{\frac{H}{H+D}}-1}{\left(1-\frac{\Delta x_b}{x(t_0)e^{r(H+D)}}\right)^{\frac{H}{H+D}}-1}\right].$$

If $0 < \Delta x_b < x(t_0)e^{r(H+D)}$, that is, for a positive later payment below some threshold, the fraction in square braces has a positive numerator and a negative denominator, meaning that the expression is positive. Furthermore, because $0 < H/H + D < 1$, we know by concavity that

$$\left|\left(1+\frac{\Delta x_b}{x(t_0)e^{r(H+D)}}\right)^{\frac{H}{H+D}}-1\right| < \left|\left(1-\frac{\Delta x_b}{x(t_0)e^{r(H+D)}}\right)^{\frac{H}{H+D}}-1\right|,$$

from which it follows that $\delta^+/\delta^- < 1$ and the sign effect is predicted.

Let us consider a simple numerical example. We assume $H = 1$ day, $D = 1$ day, $x(t_0)$=\$10, and $r = 10\%/$ day. With these parameters, a growth-optimal decision maker will be indifferent between an early (positive) payment of \$1 and a later payment of \$2.31. However, she will also be indifferent between an early loss of \$1 and a later loss of \$2.11. In other words, the gain is discounted more than the loss.

3.5.4. Magnitude Effect. Frederick et al. (2002, p. 363) describe a *magnitude effect* where "small outcomes are discounted more than large ones". In our framework, this means that the discount function for a RIPP with payments Δx_a and Δx_b would be lower than for an otherwise identical RIPP with payments $\kappa\Delta x_a$ and $\kappa\Delta x_b$, where $\kappa > 1$. The magnitude effect is not predicted in Cases A–C, where the discount function does not depend on payment sizes. We ask whether it is predicted in the present specification, Case D.

We define

$$\delta \equiv \delta(D; H, r, x(t_0), \Delta x_b),$$
$$\delta^\kappa \equiv \delta(D; H, r, x(t_0), \kappa\Delta x_b).$$

The magnitude effect is predicted if $\delta < \delta^\kappa$. By considering the ratio,

$$\frac{\delta}{\delta^\kappa} = \kappa\left[\frac{\left(1+\frac{\Delta x_b}{x(t_0)e^{r(H+D)}}\right)^{\frac{H}{H+D}} - 1}{\left(1+\frac{\kappa\Delta x_b}{x(t_0)e^{r(H+D)}}\right)^{\frac{H}{H+D}} - 1}\right],$$

we find that $\delta > \delta^\kappa$ for $\Delta x_b > 0$ and that $\delta < \delta^\kappa$ for $\Delta x_b < 0$.[10] Thus, our model does not support the magnitude effect for RIPPs as typically presented, that is, with positive payments, but supports it for negative cash flows.

3.5.5. Discounting in the Small Payment Limit. In many applications of discounting, it is plausible to assume that the payments are small relative to wealth:

$$\Delta x_a \ll x(t_0)e^{rH};$$
$$\Delta x_b \ll x(t_0)e^{r(H+D)}.$$

We can express the critical horizon and the discount function in closed form in this limit. Setting $g_a = g_b$ and using the first-order approximation $\ln(1+\epsilon) \approx \epsilon$ for $\epsilon \ll 1$, we get

$$H^{\mathrm{PR}} \equiv \frac{D\Delta x_a e^{rD}}{\Delta x_b - \Delta x_a e^{rD}},$$

and

$$\delta(D; H, r) = \left.\frac{\Delta x_a}{\Delta x_b}\right|_{g_a=g_b} \approx \frac{He^{rH}}{(H+D)e^{r(H+D)}} = \frac{e^{-rD}}{1+D/H},$$

which is a product of hyperbolic and exponential discount functions. This hybrid case has interesting behavior in the long- and short-delay limits. As shown in Figure 6, discounting is almost hyperbolic for short delays and almost exponential for long delays. Thus, for the same dynamic and the same time frame, and assuming small payments relative to initial wealth, our choice criterion predicts both approximately hyperbolic and approximately exponential discounting, depending on the delay.

4. Discussion

This paper explores temporal discounting under the postulate that decision makers maximize the growth rate of their wealth. We consider the basic temporal choice problem between two known, certain payments at known, certain, future times. To compute

Figure 6. The Discount Function in the Small Payment Limit in Case D

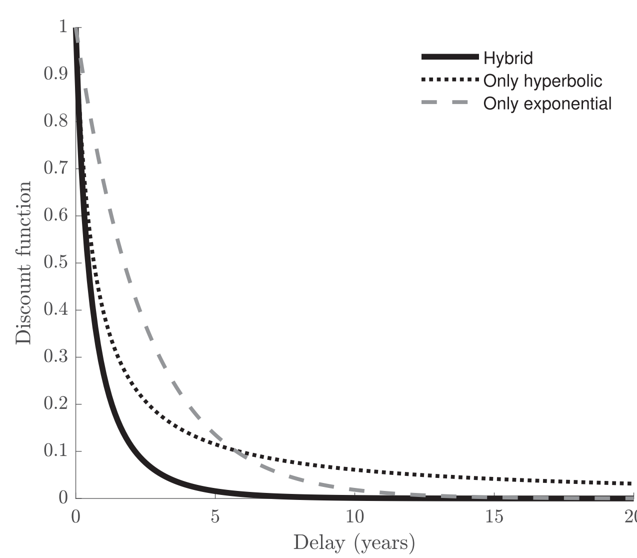

Notes. The solid black curve is the hybrid discount function $\delta = e^{-rD}/(1 + D/H)$, for $r = 0.4$ per year and $H = 0.65$ years. This is close to the hyperbolic discount function $1/(1 + D/H)$ (black dotted) for short delays and to the exponential discount function e^{-rD} (gray dashed) for long delays.

growth rates, the problem must be further specified. We incorporate in our model information about wealth dynamics, treating additive and multiplicative cases, and the time frame of the decision, meaning the period over which growth is evaluated. This information reflects circumstances that exist and are relevant in real choice problems, but about which the basic formulation is silent.

Preference reversal is an observed behavior in which decision makers switch from preferring later to earlier payments for shorter horizons. It is incompatible with the classical normative model of exponential discounting. Our model generates four different forms of discounting, depending on the decision maker's circumstances—no discounting, exponential, hyperbolic, and a hybrid of exponential and hyperbolic. The hyperbolic and hybrid forms predict preference reversal without, as is commonly needed, assumptions of behavioral bias or payment risk.

The hybrid case—corresponding to multiplicative dynamics and an adaptive time frame—suggests another type of preference reversal, called the wealth effect. Here, a decision maker switches from an earlier to a later payment as her wealth increases. In other words, richer people discount less steeply than poorer people, in line with empirical findings (Green et al. 1996, Epper et al. 2020).

Our main contribution is the prediction of nonexponential discounting and preference reversal in a normative model, growth rate maximization, that does not violate standard axioms of choice. Changes in the discount function arise only from changes in initial wealth, wealth dynamics, and time frame, which reflect in our model the real circumstances of the decisions being made. This marks a shift from psychological to circumstantial explanations of discounting. Our model assumes no dynamic inconsistency, in that the decision maker prefers at all times the option with the highest growth rate. If corroborated empirically, it would be a descriptive model. Experimental tests are feasible because the model works directly with money payments rather than utility-of-consumption flows (Cohen et al. 2020).

The temporal choice problem we study is riskless. A planned extension of this work is to explore the consequences for our model of payment uncertainty. Uncertainty decreases the long-time growth rate associated with a payment (Peters and Gell-Mann 2016). This would make a risky payment less desirable in our model, without reference to the risk preferences of the decision maker.

The dynamics discussed here do not cover the entire range of wealth dynamics. Although multiplicative and additive wealth dynamics are common and intuitive, other wealth dynamics are possible, which would lead to other forms of discounting. Our decision criterion can be adapted to general dynamics using the growth rates described in Peters and Adamou (2018). However, we stress that neither wealth dynamics nor time frame should be viewed as arbitrary free parameters in our model. Rather, they should be matched to the reality of the decisions under study.

We end with a caveat. In the small horizon limit, that is, as the earlier payment approaches an immediate payment, the growth rate corresponding to it diverges in the adaptive time frame. This indicates a loss of model realism. We link this to the breakdown of an implicit assumption: that the growth rates we compute are sustained over sufficiently long periods to be meaningful. They are, in effect, the growth rates of wealth achieved under repetition of the choice. We share the view of Kacelnik (1997, p. 60) that "the discounting process used for one-off events seems to obey a law that evolved as an adaptation to cope with repetitive events." As the horizon shrinks, this imagined repetition occurs at a frequency so high that the model decision no longer resembles a plausible real scenario.

Acknowledgments

The authors thank R. Singer for insightful discussions that led to an early outline for this manuscript. They are grateful to M. Gentry and to seminar participants at the 15th European Meeting on Game Theory, the 2019 International Conference on Public Economic Theory, the Danish Research Centre for Magnetic Resonance, the University of Cyprus, Université Paris-Dauphine, and Royal Holloway, University of London, for their invaluable feedback. They thank the associate editor and reviewers, whose careful reading and thoughtful comments improved this manuscript considerably, and Baillie Gifford, the Mercers' Company, J. Anderson, M. Elsberg, W. Miller, F. Mustafa, and C. Vermont for supporting the Ergodicity Economics program at the London Mathematical Laboratory.

Endnotes

[1] Notably, our model does not treat consumption, except inasmuch as the money payments we envisage could be net of it. In effect, we build a model of "money earlier or later" decisions (Cohen et al. 2020, p. 299). Modeling financial flows is a standard approach in theoretical and experimental studies of discounting. Cohen et al. (2020) note that inconsistencies with observed behavior have led to criticism of this approach as focusing inappropriately on money rather than consumption. One aim of this paper is to show that modeling purely monetary decisions can resolve these inconsistencies, without the need to invoke consumption.

[2] Frederick et al. (2002, p. 352) distinguish time discounting from time preference. They "use the term *time discounting* broadly to encompass any reason for caring less about a future consequence" and "use the term *time preference* to refer, more specifically, to the preference for immediate utility over delayed utility." Because growth rate maximization does not posit an innate preference for earlier over later payments, but rather seeks to predict preferences from circumstantial considerations, ours is a model of time discounting in this terminology.

[3] When comparing three payments, $(t_a, \Delta x_a)$, $(t_b, \Delta x_b)$, and $(t_c, \Delta x_c)$, the preference between them is determined by the ordering of their respective growth rates, g_a, g_b, and g_c, which are scalars, thus guaranteeing transitivity. We note that growth-optimal preferences also satisfy the independence and continuity axioms in a risky setting (see, e.g., Berman and Kirstein 2020) and, therefore, satisfy the von Neumann–Morgenstern axioms (von Neumann and Morgenstern 1944).

[4] As an example, Frederick (2003) describes an experiment in which subjects must choose between two programs that save lives. One program saves 300 lives in the first year, 200 in the second year, and 100 in the third year. The other program reverses the order. Seventy percent of subjects preferred the reversed order. Classically, this implies a discount factor greater than one, which makes the result rather puzzling. However, in the light of our approach, the result is expected, because decision makers will interpret this as a choice between positively and negatively growing trajectories.

[5] Note that the payments, Δx_a and Δx_b, appear in the growth rates, g_a and g_b, only as their ratios to the initial wealth, $x(t_0)$. An equivalent choice problem to a decision maker with different wealth has payments scaled up or down, such that these ratios are unchanged. The same comment applies to Case D.

[6] Indeed, this result corresponds to the historical use of the term "rate of discount" to describe a riskless interest rate in the money market; see, for example, Jevons (1863, p. 16).

[7] This result is consistent with subadditive discounting (Read 2001, Read and Roelofsma 2003). Read (2001) calls the time to the earlier payment the *interval* (denoted here by i) and the time to the later payment the *delay* (denoted here by d). In these terms, the discount function is

$$\delta = \frac{H}{H+D} = \frac{d-i}{d},$$

where $0 < i < d$. It follows immediately that "the discount rate will be greater the shorter the interval", which is the main premise of subadditive discounting (Read 2001, p. 11).

[8] As an example of a typical experimental observation of preference reversal, 63% of subjects in Keren and Roelofsma (1995, experiment 1) preferred f110 (Dutch guilders) in 30 weeks to f100 in 26 weeks ($a < b$ in our notation). When the same choice was presented with a shorter horizon, 82% of the same subjects preferred f100 immediately to f110 in four weeks ($a > b$).

[9] This can be shown by comparing the $H \to 0$ and $H \to \infty$ limits of g_a and g_b in Equations (4) and (5).

[10] To show this, we write the ratio δ/δ^{κ} as

$$f(y) = \frac{\kappa\left[(1+y)^{\alpha} - 1\right]}{(1+\kappa y)^{\alpha} - 1},$$

where $y = \Delta x_b / x(t_0) e^{r(H+D)}$ and $\alpha = H/H + D$. The function $f(y)$ has the following properties: it is decreasing (except at $y = 0$); it is discontinuous at $y = 0$, diverging positively as $y \to 0^+$ and negatively as $y \to 0^-$; and $f \to 1$ as $y \to \pm\infty$. Together, these facts mean that $f > 1$ for $y > 0$ and $f < 1$ for $y < 0$.

References

Berman Y, Kirstein M (2020) Risk preferences in time lotteries. Preprint, submitted February 15, https://researchers.one/articles/20.02.00007.

Britto A, Dehler-Holland J, Fichtner W (2021) Optimal investment in energy efficiency as a problem of growth rate maximisation. Working Paper 56, Series in Production and Energy, Institute for Industrial Production, Karlsruhe Institute of Technology, Karlsruhe Germany.

Carr P, Cherubini U (2020) Generalized compounding and growth optimal portfolios: Reconciling Kelly and Samuelson. Preprint, submitted March 9, http://dx.doi.org/10.2139/ssrn.3529729.

Cohen J, Marzilli Ericson K, Laibson D, White JM (2020) Measuring time preferences. *J. Econom. Lit.* 58(2):299–347.

Dasgupta P, Maskin E (2005) Uncertainty and hyperbolic discounting. *Amer. Econom. Rev.* 95(4):1290–1299.

Epper T, Fehr E, Fehr-Duda H, Kreiner CT, Lassen DD, Leth-Petersen S, Rasmussen GN (2020) Time discounting and wealth inequality. *Amer. Econom. Rev.* 110(4):1177–1205.

Frederick S (2003) Measuring intergenerational time preference: Are future lives valued less? *J. Risk Uncertainty* 26(1):39–53.

Frederick S, Loewenstein G, O'Donoghue T (2002) Discounting and time preference: A critical review. *J. Econom. Lit.* 40(2):351–401.

Green L, Myerson J (1996) Exponential vs. hyperbolic discounting of delayed outcomes: Risk and waiting time. *Amer. Zoologist* 36(4): 496–505.

Green L, Myerson J, Lichtman D, Rosen S, Fry A (1996) Temporal discounting in choice between delayed rewards: The role of age and income. *Psych. Aging* 11(1):79–84.

Guyse JL, Simon J (2011) Consistency among elicitation techniques for intertemporal choice: A within-subjects investigation of the anomalies. *Decision Anal.* 8(3):233–246.

Guyse JL, Keller LR, Huynh CH (2020) Valuing sequences of lives lost or saved over time: Preference for uniform sequences. *Decision Anal.* 17(1):24–38.

Holling CS (1959) Some characteristics of simple types of predation and parasitism. *Canadian Entomologist* 91(7):385–398.

Jevons WS (1863) *A Serious Fall in the Value of Gold Ascertained, and Its Social Effects Set Forth* (E. Stanford, London).

Kacelnik A (1997) Normative and descriptive models of decision making: Time discounting and risk sensitivity. Bock GR, Cardew G, eds. *Characterizing Human Psychological Adaptations* (John Wiley & Sons Ltd., Chichester, UK), 51–70.

Keren G, Roelofsma P (1995) Immediacy and certainty in intertemporal choice. *Organ. Behav. Human Decision Processes.* 63(3):287–297.

Laibson D (1997) Golden eggs and hyperbolic discounting. *Quart. J. Econom.* 112(2):443–478.

Lichtendahl KC Jr, Bodily SE (2012) Multiplicative utilities for health and consumption. *Decision Anal.* 9(4):314–328.

Loewenstein G, Prelec D (1992) Anomalies in intertemporal choice: Evidence and an interpretation. *Quart. J. Econom.* 107(2):573–597.

Meder D, Rabe F, Morville T, Madsen KH, Koudahl MT, Dolan RJ, Siebner HR, Hulme OJ (2019) Ergodicity-breaking reveals time optimal economic behavior in humans. Preprint June 11, submitted, https://arxiv.org/abs/1906.04652.

Myerson J, Green L (1995) Discounting of delayed rewards: Models of individual choice. *J. Experiment. Anal. Behav.* 64(3):263–276.

Peters O (2019) The ergodicity problem in economics. *Nature Phys.* 15:1216–1221.

Peters O, Adamou A (2018) The time interpretation of expected utility theory. Preprint, submitted January 11, https://arxiv.org/abs/1801.03680.

Peters O, Gell-Mann M (2016) Evaluating gambles using dynamics. *Chaos* 26:23103.

Read D (2001) Is time-discounting hyperbolic or subadditive? *J. Risk Uncertainty* 23:5–32.

Read D, Roelofsma PHMP (2003) Subadditive vs. hyperbolic discounting: A comparison of choice and matching. *Organ. Behav. Human Decision Processes* 91(2):140–153.

Samuelson PA (1937) A note on measurement of utility. *Rev. Econom. Stud.* 4(2):155–161.

Sozou PD (1998) On hyperbolic discounting and uncertain hazard rates. *Proc. Biol. Sci.* 265(1409):2015–2020.

von Neumann J, Morgenstern O (1944) *Theory of Games and Economic Behavior* (Princeton University Press, Princeton, NJ).

Yoon H (2020) Impatience and time inconsistency in discounting models. *Management Sci.* 66(12):5850–5860.

Alexander Adamou is the director of the London Mathematical Laboratory and a fellow in its Ergodicity Economics program. He holds a PhD in applied mathematics from Imperial College London and has worked professionally in gambling, risk management, and financial technology.

Yonatan Berman is a fellow in the Ergodicity Economics program at the London Mathematical Laboratory, an associated researcher in the Stone Center on Socio-Economic Inequality at City University of New York, and a fellow of the World Inequality Database project. His main research theme is the evolution of income and wealth distributions, social mobility, and their interrelationship. He is also interested in decision theory.

Diomides Mavroyiannis teaches economics at the Milestone Institute in Budapest. He received his PhD in economics from PSL-Paris Dauphine in 2019. His research interests include decision theory, industrial organization, and innovation.

Ole Peters is a fellow at the London Mathematical Laboratory and the principal investigator of its Ergodicity Economics program. He is also an external professor at the Santa Fe Institute.