Trichotomic discounted utility

Craig S. Webb[1]

Abstract

Recent evidence on intertemporal choice suggests that decision-makers may exhibit both increasing and decreasing impatience simultaneously, called inverse-S discounting. This paper introduces *trichotomic discounted utility* to model inverse-S discounting. Under trichotomic discounted utility the decision-maker distinguishes between short-term delays, medium-term delays, and long-term delays. Exponential discounting holds within each category, but not necessarily across each category. We provide preference foundations for trichotomic discounted utility in the timed outcomes framework and in the consumption streams framework. The key axiom is a weakening of the time consistency axiom that allows for both increasing and decreasing impatience. We also consider a weak, local version of time consistency, and show that this plausible condition is satisfied only by discount functions with exponential "pieces".

1 Introduction

The present bias example of Thaler (1981) has been widely accepted as convincing evidence that decision-makers frequently violate exponential discounting, and that decreasing impatience prevails. The quasi-hyperbolic discounting model elegantly captures present bias. By retaining exponential discounting for all periods after the present, the quasi-hyperbolic discounting model is immediately a familiar tool for economists. As a testament to its tractability, quasi-hyperbolic discounting is by far

✉ Craig S. Webb
craig.webb@manchester.ac.uk

[1] Economics, The University of Manchester, Manchester, England

the most widely applied intertemporal choice model in behavioural economics (Dhami 2016, p. 645–691).

The present bias example considers a specific type of comparison—short-term delays with long-term delays. More recently, a literature has emerged that has provided a more complete picture of time discounting. As well as confirming the importance of decreasing impatience, several studies have highlighted that *increasing* impatience is also common (Abdellaoui et al. 2010, 2013; Attema et al. 2010, 2016; Chesson and Viscusi 2003; Gigliotti and Sopher 2004; Onay and Onculer 2007; Read et al. 2005; Rohde 2018). However, one should not be tempted to conclude that individuals exhibit only decreasing impatience or only increasing impatience. Sayman and Onculer (2009) and Takeuchi (2011) found that many exhibit *both* increasing impatience and decreasing impatience. More specifically, increasing impatience followed by decreasing impatience. Takeuchi (2011) called this pattern *inverse-S* discounting.

Inverse-S discounting arises, according to Sayman and Onculer 2009 and Takeuchi (2011), because of an extended notion of the "present". Such a decision-maker exhibits present bias, in the sense that he is more impatient in this period than for longer delays. During the extended present, however, his impatience may increase. This paper studies inverse-S discounting behaviour. We provide a model, *trichotomic discounted utility*, that accommodates inverse-S discounting whilst retaining the elegant simplicity and tractability of quasi-hyperbolic discounting. Under trichotomic discounted utility the decision-maker distinguishes between short-term delays, medium-term delays, and long-term delays. Exponential discounting holds within each category, but not necessarily across each category. Having a tractable model that accommodates inverse-S discounting is important for both theoretical applications and empirical applications. In theoretical work, it seems desirable that applications of behavioural models capture the relevant behaviour found in experiments. In applied work, the empiricist's choice of parametric model often influences the interpretation of the data. By fitting models that a priori exclude increasing impatience, or exclude inverse-S discounting, the empirical evidence could be misrepresented (Attema et al. 2016; Bleichrodt et al. 2009).

Section 2 provides preference conditions for inverse-S discounting. Section 3 introduces trichotomic discounting to accommodate inverse-S discounting. A preference foundation is provided in Sect. 4 using the timed outcome framework. An axiomatic foundation for an integral representation, useful for macroeconomic and other applications, is provided in the "Appendix". Section 5 presents related discussions. To readers familiar with recent developments in choice under risk, it will not have gone unnoticed that "inverse-S" sounds familiar. In experimental work on choice under risk, inverse-S probability weighting is the prevailing finding. In Sect. 5.1, a formal connection between inverse-S discounting and inverse-S probability weighting is established. In Sect. 5.2, a local version of time consistency is characterised. We show that this condition is satisfied only by discount functions with exponential "pieces" and discuss the implications for existing discount functions from the literature.

2 Inverse-S discounting behaviour

This section provides the preference conditions of inverse-S discounting behaviour. For clarity of presentation, we employ the timed outcome framework (Fishburn and Rubinstein [1982](#)). A timed outcome is an ordered pair $(t, x) \in \mathbb{R}_+^2$, interpreted as receiving outcome x and time t and zero consumption otherwise. We assume a dynamic preference $\{\succeq_d\}_{d \in \mathbb{R}_+}$ over timed outcomes that is represented by discounted utility. That is, for all $d \in \mathbb{R}_+$, $\succeq_d$ is represented by $(t, x) \mapsto D(t-d)u(x)$, with $u : \mathbb{R}_+ \to \mathbb{R}$ a utility function for outcomes with $u(0) = 0$, and $D : \mathbb{R}_+ \to \mathbb{R}$ a strictly decreasing and continuous discount function with $D(0) = 1$ and $\lim_{t \to \infty} D(t) = 0$. If $\{\succeq_d\}_{d \in \mathbb{R}_+}$ is represented by discounted utility (D, u), then it can be equivalently represented by $(\tilde{D}, \tilde{u})$ if and only if there are $a, b > 0$ such that $\tilde{D} = D^a$ and $\tilde{u} = bu^a$. *Exponential discounting* holds if preferences are represented by discounted utility and, further, there is a constant $\delta \in (0, 1)$ such that $D(t) = \delta^t$ for all $t \geqslant 0$.

A key property of exponential discounting is stationarity. Preferences satisfy *stationarity* if, for all $(x, s), (y, t) \in \mathbb{R}_+^2$ and $\Delta \geqslant \max\{-s, -t\}$ we have $(s, x) \succeq_0 (t, y)$ if and only if $(s + \Delta, x) \succeq_0 (t + \Delta, y)$. Stationarity characterises the behaviour of a decision-maker who exhibits a constant level of impatience. The first preference suggests that improving outcome x to y is insufficient to compensate for the additional waiting time $(t - s)$. If a common delay increment Δ is added to both outcomes, a decision-maker with constant impatience will continue to report that the outcome improvement is insufficient to compensate for the additional waiting time $(t - s)$.

It is well-known that violations of the stationarity axiom will occur if a decision-maker is *present-biased*. For example, a decision-maker may be unwilling to wait when comparing (now, £100) and (2 weeks, £110), but willing to wait when considering (52 weeks, £100) and (54 weeks, £110). Present bias is a special case of decreasing impatience. We are interested in characterising the preferences of a decision-maker who exhibits *both* increasing *and* decreasing impatience. The following pattern is a potential (double) violation of stationarity:

Definition 1 Preferences exhibit the *static inverse-S discounting* pattern if there exists $x < y$, $s < t$, and $0 < \Delta_1 < \Delta_2$, such that

$$(s, x) \sim_0 (t, y), \quad (s + \Delta_1, x) \succeq_0 (t + \Delta_1, y), \quad \text{and} \quad (s + \Delta_2, x) \preceq_0 (t + \Delta_2, y).$$

Inverse-S discounting describes the behaviour of a decision-maker exhibiting increasing impatience, followed by decreasing impatience. To capture this, we begin with a benchmark indifference $(s, x) \sim_0 (t, y)$. For this impatient decision-maker, the pain of delaying consumption from time s to time t is precisely offset by the outcome being improved from x to y. Under reasonable monotonicity and continuity assumptions, such an indifference can be found. Now consider a short delay of Δ_1, applied to both outcomes. If the decision-maker's impatience is increasing, then an improvement $(y - x)$ is no longer sufficient to compensate for the delay $(t - s)$, and $(s + \Delta_1, x) \succeq_0 (t + \Delta_1, y)$ holds. However, if a longer delay of Δ_2 is applied to both outcomes, to a degree that the decision-maker's impatience has decreased, then

an improvement $(y - x)$ is more than sufficient to compensate for the delay of $(t - s)$, and $(s + \Delta_2, x) \preccurlyeq_0 (t + \Delta_2, y)$ holds.

Although this study is motivated by empirical evidence, it is worth discussing the plausibility of inverse-S behaviour. Present bias is very intuitive. How, then, should we understand the inverse-S pattern? Takeuchi (2011) suggested that inverse-S discounting results from an "extended present". Rather than taking the "present" as the immediate now, it seems natural to allow for (subjective) non-degenerate periods. The next month, for example. Inverse-S discounting behaviour is a natural manifestation of present bias, once one considers what happens during this extended present. Consider the following example of inverse-S discounting behaviour:

$$(\text{now}, £100) \sim_0 (2 \text{ weeks}, £110), \quad (2 \text{ weeks}, £100) \succcurlyeq_0 (4 \text{ weeks}, £110),$$
$$\text{and} \quad (52 \text{ weeks}, £100) \preccurlyeq_0 (54 \text{ weeks}, £110).$$

Comparing each of the first two preferences with the third preference, the decision-maker in both cases reveals decreasing impatience. If this decision-maker's "present" period is the next month, then we can say that both cases reveal present-biased behaviour. The first two preferences also reveal impatience increasing during the present. Initially, with plenty of "present" time left, the decision-maker is relaxed and willing to wait. Towards the end of the present, such waiting becomes intolerable. In this sense, this present biased decision-maker naturally reveals increasing impatience, so reveals inverse-S discounting behaviour.

The static inverse-S discounting pattern is one way in which a decision-maker may reveal increasing and decreasing impatience. Such a tendency may also be revealed through changes in the decision-maker's plans over time. A central property of exponential discounting is time consistency. Preferences satisfy *time consistency* if, for all $(s, x), (t, y) \in \mathbb{R}_+^2$ and $0 \leqslant d_1, d_2 \leqslant \min\{s, t\}$ we have $(s, x) \succcurlyeq_{d_1} (t, y)$ if and only if $(s, x) \succcurlyeq_{d_2} (t, y)$. Here, the decision-maker is comparing outcomes arriving at fixed calendar times. Time consistency requires that the decision-maker does not reverse his preferences as the decision time changes. If this decision-maker today (in 2018) prefers £1100 on 1st January 2022 over £1000 on 1st January 2020, then he will not change his mind at time before 1st January 2020. It is well known that present bias, as implied by decreasing impatience, can lead to violations of time consistency. One who today prefers £1100 on 1st January 2022 over £1000 on 1st January 2020, may well prefer the £1000 when asked on 1st January 2020. That is, when immediate consumption is available, one's capacity to wait is diminished. The following implication of increasing and decreasing impatience is a potential (double) violation of time consistency:

Definition 2 Preferences exhibit the *dynamic inverse-S discounting* pattern if there exists $x < y$, and $d_1 < d_2 < d_3 < s < t$, such that:

$$(s, x) \sim_{d_3} (t, y), \quad (s, x) \succcurlyeq_{d_2} (t, y), \quad \text{and} \quad (s, x) \preccurlyeq_{d_1} (t, y).$$

In this case, the decision-maker is indifferent, at decision time d_3, between (s, x) and (t, y). Previously, however, he has expressed two different preferences. Much

earlier, at decision time d_1, the decision-maker planned to wait for outcome y. That is, when the outcomes were further away from decision time, his desire was to be more patient. At decision time d_2, closer to the timed outcomes, the decision-maker reveals his most impatient self and prefers the sooner, but smaller outcome. This sense of urgency calms down by decision time d_3, where indifference is revealed.

The static and dynamic versions of inverse-S discounting behaviour are logically independent. In general, a decision-maker may reveal one type of behaviour without revealing the other. It would be useful for empirical research to examine these conditions separately. Under discounted utility, as assumed in this paper, the conditions are equivalent.

Theorem 1 *Under discounted utility, the following statements are equivalent:*

1. *Preferences exhibit the static inverse-S pattern.*
2. *Preferences exhibit the dynamic inverse-S pattern.*

3 Trichotomic discounting

This section introduces a discounting function that captures inverse-S discounting behaviour. Besides rationalising inverse-S discounting, our main consideration is simplicity. That is, we seek the minimal departure from exponential discounting compatible with inverse-S discounting. There have been many non-exponential functional forms of discounting proposed in the literature (see Sect. 6). Of these, *quasi-hyperbolic* discounting is the most widely studied (Hayashi 2003; Attema et al. 2010; Olea and Strzalecki 2014; Anchugina 2017) and applied (Asheim 1997; Laibson 1997; Barro 1999; Diamond and Koszegi 2003; Luttmer and Mariotti 2003). Quasi-hyperbolic discounting captures present bias, but retains exponential discounting in all other respects. Exponential discounting assumes discount factors $1, \delta, \delta^2, \ldots$ at times $0, 1, 2, \ldots$ Under quasi-hyperbolic discounting, times later than the immediate are penalised by an additional factor $\beta \in (0, 1)$, resulting in discount factors $1, \beta\delta, \beta\delta^2, \ldots$ A continuous time version, called *continuous quasi-hyperbolic* discounting, was developed by Pan et al. (2015) and Webb (2016). Continuous quasi-hyperbolic discounting assumes that there is a point in time S such that $D(t) = \delta_S^t$ for all $0 \leqslant t < S$, and $D(t) = \beta\delta_L^t$ for all $S \leqslant t < \infty$, where $\delta_S, \delta_L \in (0, 1)$ an $\beta = \left(\frac{\delta_S}{\delta_L}\right)^S$. If $0 < S \leqslant 1$, then continuous quasi-hyperbolic discounting generates the same discount factors as discrete time quasi-hyperbolic discounting. Under continuous quasi-hyperbolic discounting, the time period $[0, S]$ could be called the "present". Takeuchi's (2011) suggestion was that inverse-S discounting results from an "extended present", as discussed in the previous section. Here, we formalise this idea. Consider the following discount function:

$$
D(t) = \begin{cases} \delta_S^t & \text{if } 0 \leqslant t < S, \\ \beta\delta_M^t & \text{if } S \leqslant t < M, \\ \beta\gamma\delta_L^t & \text{if } M \leqslant t < \infty, \end{cases}
$$

Fig. 1 Trichotomic discounting and the inverse-S pattern ($\delta_M < \delta_S < \delta_L$)

where $\delta_S, \delta_M, \delta_L \in (0, 1)$ and $\beta = \left(\frac{\delta_S}{\delta_M}\right)^S$ and $\gamma = \left(\frac{\delta_M}{\delta_L}\right)^M$. Under these specifications, D is strictly decreasing and continuous.

The extended present notion here is modelled using two initial periods, 0 to S and S to M, which can be called the "short term" and "medium term". "Long term" delays are those longer than M. When applied to discounted utility, these periods should be understood in terms of delay from the decision time. We allow these intervals to be subjective, so possibly different for different individuals.[1] The notion that a decision-maker associates delays with separate categories is quite plausible; *classical categorisation* is a foundational concept of cognitive psychology. To reflect the decision-maker's trichotomous categorisation of short-term delays, medium-term delays, and long-term delays, we will call the above *trichotomic discounting*.

Consider any d_1, d_2, d_3, s, t such that $S+d_2 < M+d_1 < s < t < M+d_2 < S+d_3$, as in Fig. 2. Utility is continuous, hence there exists $x < y$ such that $(s, x) \sim_{d_3} (t, y)$. Equivalently, $\delta_S^{s-d_3} u(x) = \delta_S^{t-d_3} u(y)$. Then, we obtain:

$$
\begin{aligned}
(s, x) \succsim_{d_2} (t, y) &\Leftrightarrow \delta_M^{s-d_2} u(x) > \delta_M^{t-d_2} u(y) &\Leftrightarrow \delta_M < \delta_S, \\
(s, x) \precsim_{d_1} (t, y) &\Leftrightarrow \delta_L^{s-d_1} u(x) < \delta_L^{t-d_1} u(y) &\Leftrightarrow \delta_L > \delta_S.
\end{aligned}
$$

Such preferences, therefore, exhibit the dynamic inverse-S pattern if and only if $\delta_M < \delta_S < \delta_L$, as in Fig. 1. Trichotomic discounting has the flexibility to allow both increasing and decreasing impatience simultaneously. In spite of this flexibility, the simplicity of quasi-hyperbolic discounting is retained. For example, if $0 < S \leqslant 1 < M \leqslant 2$, the discrete time discount factors at times $0, 1, 2, 3, \ldots$ are $1, \beta\delta, \beta\gamma\delta^2, \beta\gamma\delta^3 \ldots$ In the next section, we provide a preference foundation for trichotomic discounting.

[1] Takeuchi (2011, p.472) found an average period of 22 days for this combined period.

Fig. 2 Explaining the inverse-S pattern

4 Axiomatisation

Under discounted utility, exponential discounting is characterised by *time consistency*. Trichotomic discounting, presented above, must therefore allow violations time consistency. In this section, we identify the precise nature of these violations. We provide an axiom, called *trichotomic time consistency*, that characterises trichotomic discounting. This axiomatisation serves at least three purposes. First, it provides the critical test of the model - the distinguishing, falsifiable condition for empirical tests. Second, because our preference foundation is expressed in terms of time consistency, it clarifies the normative content of the model. Third, the axiom is suggestive of the psychology underlying the proposed model.

To formulate our axiom, we will consider a restricted version of time consistency. Consider the following:

Definition (*Time consistency in* $[\sigma, \tau)$) For all $(t, x), (s, y) \in \mathbb{R}^2_+$ with $0 \leqslant \sigma \leqslant t - d \leqslant s - d < \tau$ and $0 \leqslant \sigma \leqslant t - d' \leqslant s - d' < \tau$, we have $(t, x) \succcurlyeq_d (s, y)$ if and only if $(t, x) \succcurlyeq_{d'} (s, y)$.

Time consistency in $[\sigma, \tau)$ restricts the time consistency condition to hold only for delays contained in the interval $[\sigma, \tau)$. For example, time consistency in [1 month, 12 months) is the condition that the decision-maker never reverses preferences for timed outcomes occurring within one month to a year from decision time. If this decision-maker, on 1st January, prefers (1st July, £120) to (1st June, £100), then he will express the same preference at all times until at least 1st May. After 1st May the delay between the decision time and 1st June is less than one month and this decision-maker may reverse his preference.

The full time consistency axiom corresponds to time consistency in $[0, \infty)$. Time consistency in $[\sigma, \tau)$ is a falsifiable condition provided, that is, that the interval $[\sigma, \tau)$ is specified. To formulate an axiom, a problem remains—in which intervals should we expect time consistency to hold? As a starting point, one might consider a notion of similarity. Delays of 13 months and 14 months, for example, could perhaps be judged "similar". Then, violations of time consistency in that interval, where one compares "like with like", would seem less frequent. Here, we make a simplifying assumption and assume at most three intervals. Time consistency will hold in each interval, but

Fig. 3 The necessity of trichotomic time consistency

not necessarily across intervals. These intervals could be called the short term, the medium term and the long term. Of course, this categorisation is subjective, and we seek a condition that allows these intervals to be different for different decision-makers. Our aim is to provide an axiom that delivers time consistency in the short term, time consistency in the medium term, and time consistency in the long term, all without assuming knowledge of these particular intervals. The following axiom achieves this.

Axiom (*Trichotomic time consistency*) For all $0 \leqslant \sigma \leqslant \tau < \infty$, at least one of: time consistency in $[0, \sigma)$, time consistency in $[\sigma, \tau)$, or time consistency in $[\tau, \infty)$ holds.

To understand this axiom, it will help to show why trichotomic time consistency is necessary for trichotomic discounting. Suppose that trichotomic discounting holds with parameters S and M. Now consider any pair of delays σ and τ, with $0 \leqslant \sigma \leqslant \tau < \infty$. Although σ and τ can be chosen freely, there are only six relevant configurations: both short-term delays, a short-term and a medium-term delay, a short-term and a long-term delay, both medium-term delays, a medium- and a long-term delay, or both long-term delays. Respectively, from top to bottom in Fig. 3: $\sigma \leqslant \tau \leqslant S$, $\sigma \leqslant S < \tau \leqslant M$, $\sigma \leqslant S < M < \tau$, $S < \sigma \leqslant \tau \leqslant M$, $S < \sigma \leqslant M < \tau$, and $M < \sigma \leqslant \tau$. In the first case, where both delays are short term, preferences satisfy time consistency in $[0, \sigma)$ and in $[\sigma, \tau)$, but not necessarily in $[\tau, \infty)$. In the second case, a short-term and a medium-term delay, only time consistency in $[0, \sigma)$ holds. Continuing through each of the remaining cases, we will see that at least one of time consistency in $[0, \sigma)$, time consistency in $[\sigma, \tau)$, or time consistency in $[\tau, \infty)$ must always hold.

Under trichotomic discounting, trichotomic time consistency necessarily holds. Notice that, if σ and τ above were chosen equal to S and M, then all three of the conditions of trichotomic time consistency would hold simultaneously. Indeed, for preferences outside of the exponential discounting and continuous quasi-hyperbolic discounting classes, this is the unique occasion when this happens. The following theorem establishes that trichotomic time consistency characterises our model in the timed outcomes framework. The proof of sufficiency essentially reverses the argument—we show that a unique choice of σ and τ exists such that all three conditions hold simultaneously, so that this choice then defines the short-, medium- and long-term intervals.

Theorem 2 *Under discounted utility, the following statements are equivalent:*

1. *Trichotomic time consistency holds.*
2. *There exists $S, M \in T$ and constants $\delta_S, \delta_M, \delta_L \in (0, 1)$ such that:*

$$D(t) = \begin{cases} \delta_S^t & \text{if } 0 \leqslant t < S, \\ \beta\delta_M^t & \text{if } S \leqslant t < M, \\ \beta\gamma\delta_L^t & \text{if } M \leqslant t < \infty, \end{cases}$$

where $\beta = \left(\frac{\delta_S}{\delta_M}\right)^S$ and $\gamma = \left(\frac{\delta_M}{\delta_L}\right)^M$.

The representation (D, u) in Theorem 2 can be replaced with $(\tilde{D}, \tilde{u})$ if and only if there are $a, b > 0$ such that $\tilde{D} = D^a$ and $\tilde{u} = bu^a$. If it happens that $\delta_S = \delta_M = \delta_L$, then exponential discounting holds. In that case, S and M can be chosen arbitrarily. If $\delta_S = \delta_M$ and $\delta_M \neq \delta_L$, then M is uniquely defined, but S can be chosen arbitrarily in the interval $[0, M]$. Similarly, if $\delta_S \neq \delta_M$ and $\delta_M = \delta_L$, then S is uniquely defined, but M can be chosen arbitrarily in the interval $[S, \infty)$. If it holds that $\delta_S \neq \delta_M$ and $\delta_M \neq \delta_L$, which is the relevant case for inverse-S discounting, then S and M are uniquely defined.

5 Related discussion

5.1 Discounting and probability weighting

Many readers' previous exposure to the term "inverse-S" will be in relation to choice under *risk*, rather than time. For choice under risk, there is overwhelming evidence that many decision-makers treat probabilities in an inverse-S way. Small probabilities are overweighted, and large probabilities are underweighted. Anecdotally, such a decision-maker might simultaneously play the lottery and purchase insurance. Wakker (2010: 203–243, 290–292) contains references to many studies that found the inverse-S pattern for choice under risk and uncertainty.

At first sight, it is remarkable that empirical research in intertemporal choice has revealed the inverse-S pattern similar to choice under risk. However, the idea that time and risk preferences are closely related has been suggested before. The future is inherently uncertain and, conversely, uncertainty often takes time to resolve (Loewenstein and Prelec 1991). Future consumption could be subject to a variety of risks. There is always a chance that the consumption good or, indeed, the decision-maker, will not survive the required time. The relationship between violations of expected utility in risky choice and violations of exponential discounting in intertemporal choice has been studied by Loewenstein and Prelec (1991), Halevy (2008) and Saito (2015).

This section establishes a formal connection between inverse-S discounting for time and inverse-S probability weighting for risk. A simple framework is assumed. Risk preferences $\succcurlyeq^r$ are defined over binary lotteries with one non-zero outcome, $\mathscr{L} = \{(p, x) : p \in [0, 1], x \in \mathbb{R}\}$, where (p, x) denotes the lottery that gives outcome x with probability p, and zero otherwise.

Definition 3 Risk preferences $\succcurlyeq^r$ exhibit the *inverse-S* pattern if there exists $x < y$, $0 < q < p \leqslant 1$, and $0 < \beta < \alpha < 1$, such that:

$$(p, x) \sim^r (q, y), \quad (\alpha p, x) \preccurlyeq^r (\alpha q, y), \quad \text{and} \quad (\beta p, x) \succcurlyeq^r (\beta q, y).$$

The inverse-S pattern suggests the decision is pessimistic for higher probabilities and optimistic for lower probabilities. Consider the intuition when p and α close to one and β is close to zero. The first indifference indicates that the outcome reduction from y to x is precisely offset by the probability improvement q to p. For α close to one, so that αp and αq are in the pessimistic region, the decision-maker is less enthusiastic about sacrificing y for an improved chance of receiving x, hence $(\alpha p, x) \preccurlyeq^r (\alpha q, y)$. For β close to zero, so that βp and βq are in the optimistic region, the decision-maker is more enthusiastic about sacrificing y for an improved chance of receiving x, hence $(\beta p, x) \succcurlyeq^r (\beta q, y)$.

Following Halevy (2008) and Saito (2015), we suppose that time preferences can be derived from risk preferences. Consider the timed outcome (t, x) at decision time $d \leqslant t$. The probability of survival until time t, conditional on decision time d, is $p(t|d)$. Assume that the probability of termination is described by an exponential distribution with hazard rate $\lambda > 0$. The probability of survival until time t, conditional on decision time d, is $p(t|d) = e^{-\lambda(t-d)}$. Therefore, the decision-maker's dynamic preferences and risk preferences are related as follows:

$$(s, x) \succcurlyeq_d (t, y) \quad \text{if and only if} \quad (e^{-\lambda(s-d)}, x) \succcurlyeq^r (e^{-\lambda(t-d)}, y),$$

for all decision t times $d \in \mathbb{R}_+$. If this equivalence holds, we will say that time preferences are determined by risk preferences.

Theorem 3 *Let time preferences be determined by risk preferences. Then, the following statements are equivalent:*

1. *Time preferences exhibit the inverse-S pattern.*
2. *Risk preferences exhibit the inverse-S pattern.*

5.2 Local time consistency

The trichotomic discounting model, being *piecewise exponential*, implies periods of time-consistent behaviour. It is shown in this section that, under discounted utility, if time consistency holds on any time interval then the restriction of the discount function must be exponential discounting on that interval. Furthermore, for existing parametric discount functions, the same condition implies that time consistency holds globally, collapsing those models to exponential discounting. Consider the following condition.

Definition (*Local time consistency*) There exists $0 \leqslant \sigma < \tau < \infty$ such that time consistency in $[\sigma, \tau)$ holds.

Local time consistency requires that, during some interval, *however brief*, the decision-maker does not reverse preferences. Let us discuss the plausibility of this

condition. Empirically observed violations of time consistency typically occur when the subjects are comparing immediate outcomes with distant outcomes. But, suppose a decision-maker forms preferences over a set of timed outcomes that occur, for example, between 27 and 28 months from now. Is there a compelling reason why this decision-maker will change his mind, reversing any preference, during the next month? Are preferences really so unstable? Empirical tests would be useful to inform the discussion. Halevy (2015) tested time consistency, and found about half of subjects to be time consistent. Local time consistency is much weaker. Even though local time consistency in continuous time is not, strictly speaking, a falsifiable condition, a discrete time analogue could be tested. In any case, I would argue that the existence of some short period of time where a decision-maker's preferences are stable is a reasonable assumption. Consider the following theorem:

Theorem 4 *Under discounted utility, the following statements are equivalent:*

1. *Local time consistency holds.*
2. *There exists constants $\delta \in (0, 1)$ and $\alpha > 0$ such that $D(t) = \alpha \delta^t$ for all $t \in [\sigma, \tau)$.*

The above theorem guarantees that, if local time consistency holds on any interval, however small, then the discount function is exponential on that interval. To the best of my knowledge, this result has not been previously obtained in the literature. Let us comment on the proof. To characterise exponential discounting over the non-negative real numbers $\mathbb{R}_+$, one can establish that a Cauchy equation $D(s + t) = D(s)D(t)$ holds for all $s, t \geqslant 0$. The difficulty here is that when s and t belong to the interval $[\sigma, \tau)$, one cannot guarantee that $s + t$ also belongs to the interval $[\sigma, \tau)$. Hence, proving Theorem 4 amounts to an extension problem. That is, finding an extension of D from $[\sigma, \tau)$ to $\mathbb{R}_+$ that satisfies the required functional equation. In the proof, we provide such an extension from $[\sigma, \tau)$ to $[\sigma, \tau) \cup [2\sigma, 2\tau)$, and can then apply results from Aczel and Skof (2007) to obtain the extension to $\mathbb{R}_+$.

There have been many parametric forms of time discounting proposed in the literature. Time preferences that are locally time consistent are characterised by a discount function with an "exponential piece". This includes trichotomic discounting, which satisfies time consistency in $[0, S)$, in $[S, M)$, and in $[M, \infty)$. Olea and Strzalecki (2014) considered a class of discount functions, called *semi-hyperbolic*, which are "eventually exponential". That is, there is some M such that time consistency in $[M, \infty)$ holds, which clearly includes trichotomic discounting.

Table 1 contains recent examples of discount functions proposed in the literature. For those discount functions, we have the following corollary of Theorem 4.

Corollary *Let discounted utility hold, with the discount function belonging to one of the classes in Table 1. Then, the following statements are equivalent:*

1. *Local time consistency holds.*
2. *Exponential discounting holds.*

The "locally exponential" property is equivalent to "globally exponential" for the parametric discount functions in Table 1. The list is not exhaustive. The above corollary suggests a problem with the involved parametric discount functions. It is important,

Table 1 Continuous discount functions

Constant sensitivity	Generalised hyperbolic
$D(t) = \delta^{t^{\xi}}$	$D(t) = (1 + \alpha t)^{-\frac{\beta}{\alpha}}$
for $t \geqslant 0$, with $\delta \in (0, 1)$, $\xi > 0$.	for $t \geqslant 0$, with $\alpha \geqslant 0$, $\beta > 0$
	Note: $D(t) = \delta^t$ is the $\alpha = 0$ case
Constant absolute DI	Double exponential
$D(t) = \beta^{\alpha^t}$	$D(t) = \omega \delta^t + (1 - \omega)\gamma^t$
for $t \geqslant 0$, with $\alpha \geqslant 0$, $\beta \in (0, 1)$.	for $t \geqslant 0$, with $\delta, \gamma \in (0, 1)$
Note: $D(t) = \delta^t$ is the $\alpha = 0$ case	and $\omega \in (0, 1]$

Constant sensitivity, also called constant relative decreasing impatience, Weibull discounting, and unit invariance discounting, see Prelec (1989, 2004), Ebert and Prelec (2007), Bleichrodt et al. (2009). Constant absolute decreasing impatience, see Prelec (1989), Bleichrodt et al. (2009). Generalised hyperbolic, see Loewenstein and Prelec (1992), the $\alpha = 1$ case see Harvey (1986), the $\alpha = \beta$ case see Mazur (1987) and Harvey (1995). Double exponential, see McClure et al. (2007)

from a descriptive point of view, to characterise models that violate time consistency. But one may ask if these discount functions are too inconsistent.[2] If the motivation for time-inconsistent models is only to explain biases that emerge when comparing short-term with long-term delays, then dispensing with time consistency of "like for like" comparisons goes further than required. The view here is that discount functions with exponential pieces are a promising alternative. To maintain tractability, using as few pieces as possible to capture the relevant behaviour seems desirable.

6 Closing comments

This paper has introduced a simple model of inverse-S discounting called *trichotomic discounting*. Exponential discounting and continuous quasi-hyperbolic discounting remain as special cases, but we add the flexibility to allow for the type of inverse-S discounting behaviour found in experiments. By using a piecewise exponential functional form, the model retains much of the normative content of exponential discounting. It will be interesting for future research to provide applications of trichotomic discounting. Trichotomic discounting is a tractable tool for analysing the connections between inverse-S discounting behaviour and economic outcomes. If this form of discounting behaviour is indeed prevalent, one would hope that its incorporation into economic models would shed light on previously unexplained phenomena.

---

[2] To address this question, it would be useful for future research to characterise these functions in terms of restricted time consistency properties.

Appendix A: Integral representation

In this section the trichotomic discounted utility model is extended to an integral representation. In obtaining such a representation, we provide the appropriate foundations for applications with consumption streams, used in finance and macroeconomics. We will make use of the tools developed by Kopylov (2010). In particular, we exploit only the richness naturally provided by the time dimension. No richness is required of the set of outcomes. Hence, the theory can be applied to any type of outcomes, be they monetary, health related, indivisible goods, and so on.

The set of *outcomes* is X, time is $T = [0, \infty)$, and the set of half-open intervals $[a, b)$ is $\mathcal{T}$. *Consumption streams* are $\mathcal{T}$-measurable functions $\mathbf{x} : T \rightarrow X$, the set of which is $\mathcal{C}$. For a decision time $d \in T$, let $\mathcal{C}_d = \{\mathbf{x}|_{[d,\infty)} : \mathbf{x} \in \mathcal{C}\}$ denote the set of consumptions streams restricted to times not earlier than d. A *dynamic preference* is a collection of static preference relations $\mathcal{R} = \{\succcurlyeq_d\}_{d \in T}$ where each $\succcurlyeq_d$ is defined over $\mathcal{C}_d$. A *dynamic model* $\mathcal{V} = \{V_d\}_{d \in T}$ is a collection of real-valued functions $V_d : \mathcal{C}_d \rightarrow \mathbb{R}$. A dynamic preference $\mathcal{R}$ is *represented* by $\mathcal{V}$ if for each $\succcurlyeq_d \in \mathcal{R}$ there is a $V_d \in \mathcal{V}$ such that, for all $\mathbf{x}, \mathbf{y} \in \mathcal{C}_d$, $\mathbf{x} \succcurlyeq_d \mathbf{y}$ if and only if $V_d(\mathbf{x}) \geqslant V_d(\mathbf{y})$.

Some further notation regarding consumption streams is useful. For $\mathbf{x}, \mathbf{z} \in \mathcal{C}$ and $0 \leqslant a \leqslant b < \infty$, we use $\mathbf{x}[a, b)\mathbf{z}$ to denote the stream that coincides with $\mathbf{x}$ in the interval $[a, b)$, and coincides with $\mathbf{z}$ elsewhere. For an outcome $x \in X$, we use $\langle x \rangle$ to denote the *constant* stream that yields outcome x at all points in time. Given a consumption stream $x \in \mathcal{C}$ and $d \in T$, denote by $\mathbf{x}_d$ the stream $\mathbf{x}_d \in \mathcal{C}_d$ such that $\mathbf{x}_d(t) = \mathbf{x}(t - d)$ for all $t \geqslant d$. The following axioms are assumed:

Axiom 1 (*Weak ordering*) For all $\succcurlyeq_d \in \mathcal{R}$, $\succcurlyeq d$ over $\mathcal{C}_d$ is complete and transitive.

Axiom 2 (*Common outcome independence*) For all $\succcurlyeq_d \in \mathcal{R}$, $\mathbf{x}, \mathbf{y}, \mathbf{z}, \tilde{\mathbf{z}} \in \mathcal{C}_d$, and $a \leqslant b$, $\mathbf{x}[a, b)\mathbf{z} \succcurlyeq_d \mathbf{y}[a, b)\mathbf{z}$ only if $\mathbf{x}[a, b)\tilde{\mathbf{z}} \succcurlyeq_d \mathbf{y}[a, b)\tilde{\mathbf{z}}$.

Axiom 3 (*Interval monotonicity*) For all $\succcurlyeq_d \in \mathcal{R}$, $\mathbf{x}, \mathbf{y} \in \mathcal{C}_d$, if $\langle \mathbf{x}(t) \rangle \succcurlyeq_d \langle \mathbf{y}(t) \rangle$ for all $t \geqslant d$, then $\mathbf{x} \succcurlyeq_d \mathbf{y}$. If it also holds that $\langle \mathbf{x}(t) \rangle \succ_d \langle \mathbf{y}(t) \rangle$ for all $t \in [a, b)$, for some $a < b$, then $\mathbf{x} \succ_d \mathbf{y}$.

Axiom 4 (*Weak outcome separability*) For all $\succcurlyeq_d \in \mathcal{R}$, $a \leqslant b$, $c \leqslant d$, and all $x, y, \tilde{x}, \tilde{y} \in X$, with $\langle x \rangle \succ_d \langle y \rangle$ and $\langle \tilde{x} \rangle \succ_d \langle \tilde{y} \rangle$, $\langle x \rangle[a, b)\langle y \rangle \succ_d \langle x \rangle[c, d)\langle y \rangle$ only if $\langle \tilde{x} \rangle[a, b)\langle \tilde{y} \rangle \succ_d \langle \tilde{x} \rangle[c, d)\langle \tilde{y} \rangle$.

Axiom 5 (*Strong monotone continuity*) For all $\succcurlyeq_d \in \mathcal{R}$, $\langle x \rangle, \mathbf{y}, \mathbf{z} \in \mathcal{C}_d$, and $\{[a_i, b_i)\}_{i=1}^{\infty}$ with $[a_1, b_1] \supset [a_2, b_2] \supset \ldots$ and $\cap_{i=1}^{\infty}[a_i, b_i)$ either empty or a single point, if $\mathbf{y} \succcurlyeq_d \langle x \rangle[a_i, b_i)\mathbf{z}$ for all i, or if $\langle x \rangle[a_i, b_i)\mathbf{y} \succcurlyeq_d \mathbf{z}$ for all i, then $\mathbf{y} \succcurlyeq_d \mathbf{z}$.

Axiom 6 (*Time invariance*) For all $\mathbf{x}, \mathbf{y} \in \mathcal{C}$, $\mathbf{x} \succcurlyeq_0 \mathbf{y}$ if and only if $\mathbf{x}_d \succcurlyeq_d \mathbf{y}_d$.

Axioms 1–5 are essentially intertemporal analogues of the axioms for subjective expected utility (Kopylov 2010). Axiom 6 is a necessary condition for discounted utility, and is consumption stream version of the condition presented by Halevy (2015). Under Axioms 1–6, the following condition is the characterising property of exponential discounting:

Definition (*Time consistency*) For all $0 \leqslant a \leqslant b < \infty$, $\mathbf{x}, \mathbf{y}, \mathbf{z} \in \mathscr{C}$ and $d, d' \leqslant a$ we have $\mathbf{x}[a, b)\mathbf{z} \succcurlyeq_d \mathbf{y}[a, b)\mathbf{z}$ if and only if $\mathbf{x}[a, b)\mathbf{z} \succcurlyeq_{d'} \mathbf{y}[a, b)\mathbf{z}$.

Consider two streams that are identical, except on some interval $[a, b)$. Time consistency requires that, at any time before this interval, the decision-maker does not reverse previously expressed preferences. Under discounted utility, time consistency holds if and only if exponential discounting holds. Hence, discount functions exhibiting decreasing impatience must, in the context of discounted utility, violate time consistency. Our axiomatisation is phrased in terms of the time consistency properties that remain in the more general model. The following is the consumption stream version of the condition introduced in Sect. 4:

Definition (*Time consistency in $[\sigma, \tau)$*) For all $0 \leqslant a \leqslant b < \infty$, $\mathbf{x}, \mathbf{y}, \mathbf{z} \in \mathscr{C}$ and $0 \leqslant \sigma \leqslant a - d \leqslant b - d < \tau$ and $0 \leqslant \sigma \leqslant a - d' \leqslant b - d' < \tau$, we have $\mathbf{x}[a, b)\mathbf{z} \succcurlyeq_d \mathbf{y}[a, b)\mathbf{z}$ if and only if $\mathbf{x}[a, b)\mathbf{z} \succcurlyeq_{d'} \mathbf{y}[a, b)\mathbf{z}$.

We use this restricted version of time consistency to define the key axiom for trichotomic discounting:

Axiom (*Trichotomic time consistency*) For all $0 \leqslant \sigma \leqslant \tau < \infty$, at least one of: time consistency in $[0, \sigma)$, time consistency in $[\sigma, \tau)$, or time consistency in $[\tau, \infty)$ holds.

The interpretation is similar to the axiom introduced in Sect. 4. The main result of this section characterises an integral version of trichotomic discounting model in the consumption streams framework:

Theorem 5 *The following statements are equivalent:*

1. *Axioms 1–6 (weak ordering, common outcome independence, interval monotonicity, weak outcome separability, strong monotone continuity, and time invariance) and the trichotomic time consistency axiom hold.*
2. *Dynamic preferences $\mathscr{R}$ are represented by a dynamic model $\mathscr{V}$ such that, for all $V_d \in \mathscr{V}$:*

$$V_d(\mathbf{x}) = \int_d^\infty D(t - d)u(\mathbf{x}(t))dt,$$

with $u : X \to \mathbb{R}$ a utility function for outcomes, and $D : T \to \mathbb{R}$ a strictly decreasing and continuous discount function. Furthermore, there exists $S, M \in T$ and constants $\delta_S, \delta_M, \delta_L \in (0, 1)$ such that:

$$D(t) = \begin{cases} \delta_S^t & \text{if } 0 \leqslant t < S, \\ \beta \delta_M^t & \text{if } S \leqslant t < M, \\ \beta \gamma \delta_L^t & \text{if } M \leqslant t < \infty, \end{cases}$$

where $\beta = \left(\frac{\delta_S}{\delta_M}\right)^S$ and $\gamma = \left(\frac{\delta_M}{\delta_L}\right)^M$.

The representation (D, u) in Theorem 5 can be replaced with $(\tilde{D}, \tilde{u})$ if and only if $\tilde{D} = D$ and $\tilde{u} = au + b$ for some $a > 0$ and $b \in \mathbb{R}$. If it happens that $\delta_S = \delta_M = \delta_L$, then S and M can be chosen arbitrarily. If $\delta_S = \delta_M$ and $\delta_M \neq \delta_L$, then M is uniquely defined, but S can be chosen arbitrarily in the interval $[0, M]$. If $\delta_S \neq \delta_M$ and $\delta_M = \delta_L$, then S is uniquely defined, but M can be chosen arbitrarily in the interval $[S, \infty)$. If $\delta_S \neq \delta_M$ and $\delta_M \neq \delta_L$, then S and M are uniquely defined.

Appendix B: Proofs

Proof of Theorem 1 Suppose that preferences exhibit the static inverse-S pattern, so there exists $x < y$, $s < t$, and $0 < \Delta_1 < \Delta_2$, such that: $(s, x) \sim_0 (t, y)$, $(s + \Delta_1, x) \succcurlyeq_0 (t + \Delta_1, y)$, and $(s + \Delta_2, x) \preccurlyeq_0 (t + \Delta_2, y)$. Under discounted utility, these are equivalent to $D(s)u(x) = D(t)u(y)$, $D(s + \Delta_1)u(x) \geqslant D(t + \Delta_1)u(y)$, and $D(s + \Delta_2)u(x) \leqslant D(t + \Delta_2)u(y)$. Fix a decision time $d_3 > \Delta_2$. Now define $\tilde{s} = s + d_3$, $\tilde{t} = t + d_3$, $d_2 = d_3 - \Delta_1$, and $d_1 = d_3 - \Delta_2$. Substituting these definitions, we have $D(\tilde{s} - d_3)u(x) = D(\tilde{t} - d_3)u(y)$, $D(\tilde{s} - d_2)u(x) \geqslant D(\tilde{t} - d_2)u(y)$, and $D(\tilde{s} - d_1)u(x) \leqslant D(\tilde{t} - d_1)u(y)$. These are equivalent to $(\tilde{s}, x) \sim_{d_3} (\tilde{t}, y)$, $(\tilde{s}, x) \succcurlyeq_{d_2} (\tilde{t}, y)$, and $(\tilde{s}, x) \preccurlyeq_{d_3} (\tilde{t}, y)$, hence preferences exhibit the dynamic inverse-S pattern. The converse implication is established similarly. □

Proof of Theorem 2 Observe that, if time consistency in $[\sigma, \tau)$ holds for $0 \leqslant \sigma < \tau < \infty$, then time consistency in $[\tilde{\sigma}, \tilde{\tau})$ holds for all $\sigma \leqslant \tilde{\sigma} < \tilde{\tau} < \tau$. Let $\sigma^* = \sup\{\sigma \in \mathbb{R}_+ : \text{time consistency in } [0, \sigma) \text{ holds}\}$. and let $\tau^* = \inf\{\tau \in \mathbb{R}_+ : \text{time consistency in } [\tau, \infty) \text{ holds}\}$. If $\sigma^* = \tau^*$, then CQH discounting holds (Theorem 4.2.1, Pan et al. 2015). If $\sigma^* > \tau^*$, it can be shown that time consistency holds and an exponential discounting representation exists. We consider $\sigma^* < \tau^*$ in what follows. In this case, time consistency in $[\sigma^*, \tau^*)$ must hold. If not, there would exist $\sigma^* < \sigma < \tau < \tau^*$ such that neither time consistency in $[0, \sigma)$, nor time consistency in $[\sigma, \tau)$, nor time consistency in $[\tau, \infty)$ hold. This cannot occur under trichotomic time consistency. Preferences $\{\succcurlyeq_d\}_{d \in \mathbb{R}_+}$ therefore simultaneously satisfy time consistency in $[0, \sigma^*)$, time consistency in $[\sigma^*, \tau^*)$, and time consistency in $[\tau^*, \infty)$. Define $S := \sigma^*$ and $M := \tau^*$. By, Theorem 4, applied in each interval:

$$
D(t) = \begin{cases} \alpha \delta_S^t & \text{if } 0 \leqslant t < S, \\ \beta \delta_M^t & \text{if } S \leqslant t < M, \\ \beta \gamma \delta_L^t & \text{if } M \leqslant t < \infty, \end{cases}
$$

holds, with $\alpha_1, \beta, \beta\gamma > 0$ and $\delta_S, \delta_M, \delta_L \in (0, 1)$. Recall that $D(0) = 1$, and D is continuous, so $\alpha_1 = 1$, $\delta_S^S = \beta\delta_M^S$ and $\delta_M^M = \gamma\delta_L^M$. The discount factors $\delta_S, \delta_M, \delta_L$ and utility u inherit the uniqueness properties of discounted utility's multiplicative form in the timed outcome framework. □

Proof of Theorem 3 Assume that preferences exhibit the dynamic inverse-S pattern, so that there exists $x < y$, and $d_1 < d_2 < d_3 < s < t$, such that: $(s, x) \sim_{d_3} (t, y)$, $(s, x) \succcurlyeq_{d_2} (t, y)$, and $(s, x) \preccurlyeq_{d_1} (t, y)$. If preferences are determined by risk, with

a constant hazard rate, then these are equivalent to: $(e^{-\lambda(s-d_3)}, x) \sim^r (e^{-\lambda(t-d_3)}, y)$, $(e^{-\lambda(s-d_2)}, x) \succcurlyeq^r (e^{-\lambda(t-d_2)}, y)$, and $(e^{-\lambda(s-d_1)}, x) \preccurlyeq^r (e^{-\lambda(t-d_1)}, y)$. Define $p = e^{-\lambda(s-d_3)}$, $q = e^{-\lambda(t-d_3)}$, $\alpha = e^{-\lambda(d_3-d_2)}$ and $\beta = e^{-\lambda(d_3-d_1)}$ and it follows that risk preferences exhibit the inverse-S pattern. For the converse implication, suppose that risk preferences exhibit the inverse-S pattern, so that there exists $x < y$, $0 < q < p \leqslant 1$, and $0 < \beta < \alpha < 1$, such that: $(p, x) \sim^r (q, y)$, $(\alpha p, x) \prec^r (\alpha q, y)$, and $(\beta p, x) \succcurlyeq^r (\beta q, y)$. Fix a decision time $d_3 > -\frac{\ln(\beta)}{\lambda}$. Then define $s = \frac{\lambda d_3 - \ln(p)}{\lambda}$, $t = \frac{\lambda d_3 - \ln(q)}{\lambda}$, $d_2 = \frac{\lambda d_3 + \ln(\alpha)}{\lambda}$ and $d_1 = \frac{\lambda d_3 + \ln(\beta)}{\lambda}$. Then, time preferences, determined by risk preferences, will exhibit the dynamic inverse-S pattern for these particular values.   □

Proof of Theorem 4 For given $0 \leqslant \sigma < \tau < \infty$, define *stationarity in* $[\sigma, \tau)$ by the condition: for all $(s, x), (t, y) \in \mathbb{R}_+^2$ and Δ with $\sigma \leqslant s, s + \Delta, t, t + \Delta \leqslant \tau$, we have $(s, x) \succcurlyeq_0 (t, x)$ if and only if $(s + \Delta, x) \succcurlyeq_0 (t + \Delta, y)$. Under discounted utility, time consistency in $[\sigma, \tau)$ is equivalent to stationarity in $[\sigma, \tau)$. Suppose that local time consistency holds, so there exists $0 \leqslant \sigma < \tau < \infty$, such that time consistency in (hence stationarity in) $[\sigma, \tau)$ holds. We will show that D is an exponential function on $[\sigma, \tau)$. If σ and τ are sufficiently close (if $\sigma > \frac{\tau}{2}$) then, for all $\sigma \leqslant s, t \leqslant \tau$, we have that $s + t$ will lie outside of $[\sigma, \tau)$. Hence, establishing the standard Cauchy functional equation for exponentials, $D(s + t) = D(s)D(t)$, is not immediately clear. To accomplish this, define a function $\tilde{D} : [\sigma, \tau) \cup [2\sigma, 2\tau) \to \mathbb{R}$ such that $\tilde{D}|_{[\sigma, \tau)} = D$, and $\tilde{D}(r) = \{D(s)D(t) : r = s + t, \sigma \leqslant s < t \leqslant \tau\}$ for all $r \in [2\sigma, 2\tau)$.

We now confirm that $\tilde{D}$ is well-defined. To this end, it suffices to show that $\sigma \leqslant s, s', t, t' < \tau$ and $s + t = s' + t'$ jointly imply $D(s)D(t) = D(s')D(t')$. Suppose not, so there exists $\sigma \leqslant s, s', t, t' < \tau$ with $\tilde{D}(s)\tilde{D}(t) \neq \tilde{D}(s')\tilde{D}(t')$. Let $s' = s - \varepsilon$, so that $\frac{\tilde{D}(s'+\varepsilon)}{\tilde{D}(s')} \neq \frac{\tilde{D}((t+\varepsilon)}{\tilde{D}(t)}$. Given such s', t, and discounted utility's continuity and monotonicity properties, there exists x, y such that $(s', x) \sim_0 (t, y)$, equivalent to $\tilde{D}(s')u(x) = \tilde{D}(t)u(y)$. Then, we have that $\tilde{D}(s+\varepsilon)u(x) \neq \tilde{D}(t+\varepsilon)u(y)$, equivalent to $(s' + \varepsilon, x) \nsim_0 (t + \varepsilon, y)$, contradicting stationarity in $[\sigma, \tau)$. Hence, the function $\tilde{D}$ is well-defined and, by construction, $\tilde{D}(s + t) = \tilde{D}(s)\tilde{D}(t)$ for all $t, s \in [\sigma, \tau)$. By Aczel and Skof (2007, Note 4, p. 315), there exists constants $\alpha > 0$ and $\delta \in (0, 1)$ such that $\tilde{D}(t) = D(t) = \alpha\delta^t$, for all $t = [\sigma, \tau)$.   □

Proof of Theorem 5 By Theorem 5.1 of Webb (2016), which is based on Kopylov (2010), Axioms 1–6 imply that dynamic preferences $\mathscr{R}$ are represented by a dynamic model $\mathscr{V}$ such that $V_d(\mathbf{x}) = \int_d^\infty D(t - d)u(\mathbf{x}(t))\mathrm{d}t$, for all $V_d \in \mathscr{V}$, with $u : X \to \mathbb{R}$ a utility function for outcomes, and $D : T \to \mathbb{R}$ a strictly decreasing and continuous discount function. Furthermore, the discount function satisfies $D(0) = 1$ and $\lim_{t \to \infty} D(t) = 0$.

Suppose time consistency in $[\sigma, \tau)$ holds, for some $0 < \sigma < \tau < \infty$. For a given $\varepsilon \in [0, \tau - \sigma)$, define a function $D_\varepsilon : [\sigma + \varepsilon, \tau] \to \mathbb{R}$ such that $D_\varepsilon(t) = D(\varepsilon)D(t - \varepsilon)$ for all $t \in [\sigma + \varepsilon, \tau)$. It will be shown that D and D_ε coincide on $[\sigma + \varepsilon, \tau)$. For all $\mathbf{x}[a, b)\mathbf{z}$ and $\mathbf{y}[a, b)\mathbf{z}$, with $a, b, a - \varepsilon, b - \varepsilon \in [\sigma, \tau)$, we have:

$$\mathbf{x}[a, b)\mathbf{z} \succcurlyeq_0 \mathbf{y}[a, b)\mathbf{z} \; \Leftrightarrow \; \int_a^b D(t)u(\mathbf{x}(t))\mathrm{d}t \geqslant \int_a^b D(t)u(\mathbf{y}(t))\mathrm{d}t.$$

By time consistency in $[\sigma, \tau)$, the above holds if and only if $\mathbf{x}[a, b)\mathbf{z} \succcurlyeq_\varepsilon \mathbf{y}[a, b)\mathbf{z}$, equivalent to:

$$\int_a^b D_\varepsilon(t)u(\mathbf{x}(t))\mathrm{d}t \geqslant \int_a^b D_\varepsilon(t)u(\mathbf{y}(t))\mathrm{d}t.$$

Consider the set of consumption streams:

$$\mathscr{C}_\varepsilon := \{\mathbf{x}[a, b)\mathbf{z} : \mathbf{x}, \mathbf{z} \in \mathscr{C}, a, b, a - \varepsilon, b - \varepsilon \in [\sigma, \tau)\}.$$

Preferences $\succsim_0$ over $\mathscr{C}_\varepsilon$ are equivalently represented by both (D, u) and (D_ε, u). The set $\mathscr{C}_\varepsilon$ is sufficiently rich that the uniqueness results of Kopylov's theorem hold. Because $D_\varepsilon(0) = 1$ and $\lim_{t\to\infty} D_\varepsilon(t) = 0$, we have $D = D_\varepsilon$ on $[\sigma + \varepsilon, \tau)$. Recall that ε was chosen arbitrarily from $[0, \tau - \sigma)$. Hence, D satisfies $D(s + \varepsilon)D(t) = D(s)D(t + \varepsilon)$ for all $s, t, s + \varepsilon, t + \varepsilon \in [\sigma, \tau)$. Now proceed as in the proof of Theorem 4. Define a (well-defined) function $\tilde{D} : [\sigma, \tau) \cup [2\sigma, 2\tau) \to \mathbb{R}$ such that $\tilde{D}|_{[\sigma,\tau)} = D$, and $\tilde{D}(r) = \{D(s)D(t) : \sigma \leqslant s < t < \tau, s+t = r\}$ for all $r \in [2\sigma, 2\tau)$. By Aczel and Skof (2007, Note 4, p.315), there exists constants $\alpha > 0$ and $\delta \in (0, 1)$ such that $\tilde{D}(t) = D(t) = \alpha\delta^t$, for all $t = [\sigma, \tau)$.

Now assume that trichotomic time consistency holds. Let $\sigma^* = \sup\{\sigma \in \mathbb{R}_+ : \text{time consistency in } [0, \sigma) \text{ holds}\}$. and let $\tau^* = \inf\{\tau \in \mathbb{R}_+ : \text{time consistency in } [\tau, \infty) \text{ holds}\}$. If $\sigma^* = \tau^*$, then CQH discounting holds (Webb 2016). If $\sigma^* > \tau^*$, then exponential discounting holds. We consider $\sigma^* < \tau^*$ in what follows. In this case, time consistency in $[\sigma^*, \tau^*)$ must hold. If not, there would exist $\sigma^* < \sigma < \tau < \tau^*$ such that neither time consistency in $[0, \sigma)$, nor time consistency in $[\sigma, \tau)$, nor time consistency in $[\tau, \infty)$ hold. This cannot occur under trichotomic time consistency. Preferences $\{\succcurlyeq_d\}_{d\in T}$ therefore simultaneously satisfy time consistency in $[0, \sigma^*)$, time consistency in $[\sigma^*, \tau^*)$, and time consistency in $[\tau^*, \infty)$. Define $S := \sigma^*$ and $M := \tau^*$. By, the above analysis, applied in each interval:

$$D(t) = \begin{cases} \alpha\delta_S^t & \text{if } 0 \leqslant t < S, \\ \beta\delta_M^t & \text{if } S \leqslant t < M, \\ \beta\gamma\delta_L^t & \text{if } M \leqslant t < \infty, \end{cases}$$

holds, with $\alpha_1, \beta, \beta\gamma > 0$ and $\delta_S, \delta_M, \delta_L \in (0, 1)$. Recall that $D(0) = 1$, and D is continuous, so $\alpha_1 = 1, \delta_S^S = \beta\delta_M^S$ and $\delta_M^M = \gamma\delta_L^M$. The discount factors are unique—a property inherited from the discounted utility representation. □

References

Abdellaoui, M., Attema, A. E., & Bleichrodt, H. (2010). Intertemporal tradeoffs for gains and losses: An experimental measurementof discounted utility. *The Economic Journal*, *120*(545), 845–866.

Abdellaoui, M., Bleichrodt, H., & l'Haridon, O. (2013). Sign-dependence in intertemporal choice. *Journal of Risk and Uncertainty*, *47*(3), 225–253.

Aczel, J., & Skof, F. (2007). Local Pexider and Cauchy equations. *Aequationes Mathematicae*, *73*, 311–320.

Anchugina, N. (2017). A simple framework for the axiomatization of exponential and quasi-hyperbolic discounting. *Theory and Decision*, *82*(2), 185–210.

Asheim, G. B. (1997). Individual and collective time-consistency. *Review of Economic Studies*, *64*, 427–443.

Attema, A. E., Bleichrodt, H., Rohde, K. I. M., & Wakker, P. P. (2010). Time-tradeoff sequences for analyzing discounting and time inconsistency. *Management Science*, *56*, 2015–2030.

Attema, A. E., Bleichrodt, H., Gao, Y., Huang, Z., & Wakker, P. P. (2016). Measuring discounting with measuring utility. *American Economic Review*, *106*(6), 1476–1494.

Barro, R. (1999). Laibson meets Ramsey in the neoclassical growth model. *Quarterly Journal of Economics*, *114*, 1125–1152.

Bleichrodt, H., Rohde, K., & Wakker, P. (2009). Non-hyperbolic time inconsistency. *Games and Economic Behavior*, *66*, 27–38.

Chesson, H., & Viscusi, W. Kip. (2003). Commonalities in time and ambiguity aversion for long-term risks. *Theory and Decision*, *54*, 57–71.

Dhami, S. (2016). *The foundations of behavioural economic analysis*. Oxford: Oxford University Press.

Diamond, P., & Koszegi, B. (2003). Quasi-hyperbolic discounting and retirement. *Journal of Public Economics*, *87*, 1839–1872.

Ebert, J., & Prelec, D. (2007). The fragility of time: Time-insensitivity and valuation of the near and far future. *Management Science*, *53*, 1423–1438.

Fishburn, P. C., & Rubinstein, A. (1982). Time preference. *International Economic Review*, *23*, 677–694.

Gigliotti, G., & Sopher, B. (2004). Analysis of intertemporal choice: A new framework and experimental results. *Theory and Decision*, *55*, 209–233.

Halevy, Y. (2008). Strotz meets Allais: Diminishing impatience and the certainty effect. *American Economic Review*, *98*(3), 1145–1162.

Halevy, Y. (2015). Time consistency: Stationarity and time invariance. *Econometrica*, *83*(1), 335–352.

Harvey, C. M. (1995). Proportional discounting of future costs and benefits. *Mathematics of Operations Research*, *20*, 381–399.

Harvey, C. M. (1986). Value functions for infinite-period planning. *Management Science*, *32*, 1123–1139.

Hayashi, T. (2003). Quasi-stationary cardinal utility and present bias. *Journal of Economic Theory*, *112*, 343–352.

Kopylov, I. (2010). Simple axioms for countably additive subjective probability. *Journal of Mathematical Economics*, *46*, 867–876.

Laibson, D. (1997). Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics*, *112*(2), 443–477.

Loewenstein, G., & Prelec, D. (1991). Decision making over time and under uncertainty: A common approach. *Management Science*, *37*(7), 770–786.

Loewenstein, G., & Prelec, D. (1992). Anomalies in intertemporal choice: Evidence and interpretation. *Quarterly Journal of Economics*, *107*, 573–597.

Luttmer, E., & Mariotti, T. (2003). Subjective discounting in an exchange economy. *Journal of Political Economy*, *111*(5), 959–989.

Mazur, J. (1987). An adjusting procedure for studying delayed reinforcement. In J. Mazur, M. Commons, J. Nevin, & H. Rachlin (Eds.), *Quantitative analyses of behavior 5: The effect of delay and of intervening events on reinforcement value* (pp. 55–73). Hillsdale, NJ: Erlbaum.

McClure, S. M., Ericson, K. M., Laibson, D. I., Loewenstein, G., & Cohen, J. D. (2007). Time discounting for primary rewards. *Journal of Neuroscience*, *27*(21), 5796–5804.

Olea, J. L. M., & Strzalecki, T. (2014). Axiomatization and measurement of quasi-hyperbolic discounting. *Quarterly Journal of Economics*, *116*, 121–160.

Onay, S., & Onculer, A. (2007). Intertemporal choice under timing risk: An experimental approach. *Journal of Risk and Uncertainty*, *34*, 99–121.

Pan, J., Webb, C. S., & Zank, H. (2015). An extension of quasi-hyperbolic discounting to continuous time. *Games and Economic Behavior*, *89*, 43–55.

Prelec, D. (1989). Decreasing impatience: Definition and consequences. Harvard Business School working paper.

Prelec, D. (2004). A criterion for non-stationary time preference and 'hyperbolic' discounting. *Scandinavian Journal of Economics*, *106*, 511–532.

Read, D., Airoldi, M., & Loewenstein, G. (2005). Intertemporal tradeoff priced in interest rates and amounts: A study of method variance. Working Paper.

Rohde, K. I. M. (2018). Measuring decreasing and increasing impatience. *Management Science*. https://doi.org/10.1287/mnsc.2017.3015.

Sayman, S., & Onculer, A. (2009). An investigation of time inconsistency. *Management Science*, *55*(3), 470–482.

Saito, K. (2015). A relationship between risk and time preferences. Working paper, California Institute of Technology.

Takeuchi, K. (2011). Non-parametric test of time consistency: Present bias and future bias. *Games and Economic Behavior*, *71*, 456–478.

Thaler, R. H. (1981). Some empirical evidence of dynamic inconsistency. *Economics Letters*, *8*, 201–207.

Wakker, P. P. (2010). *Prospect theory*. Cambridge: Cambridge University Press.

Webb, C. S. (2016). Continuous quasi-hyperbolic discounting. *Journal of Mathematical Economics*, *64*, 99–106.