Management Science

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

Beta-Delta or Delta-Tau? A Reformulation of Quasi-Hyperbolic Discounting

Han Bleichrodt, Rogier J. D. Potter van Loon, Drazen Prelec

To cite this article:
Han Bleichrodt, Rogier J. D. Potter van Loon, Drazen Prelec (2022) Beta-Delta or Delta-Tau? A Reformulation of Quasi-Hyperbolic Discounting. Management Science 68(8):6326-6335. https://doi.org/10.1287/mnsc.2022.4453

MANAGEMENT SCIENCE
Vol. 68, No. 8, August 2022, pp. 6326–6335
ISSN 0025-1909 (print), ISSN 1526-5501 (online)

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

Beta-Delta or Delta-Tau? A Reformulation of Quasi-Hyperbolic Discounting

Han Bleichrodt,[a,*] Rogier J. D. Potter van Loon,[a] Drazen Prelec[b]

[a] Erasmus School of Economics, Erasmus University Rotterdam, 3000 DR Rotterdam, the Netherlands; [b] Department of Economics and Department of Brain and Cognitive Sciences, Sloan School, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139
*Corresponding author

Contact: bleichrodt@ese.eur.nl, https://orcid.org/0000-0002-2700-412X (HB); r.pottervanloon@gmail.com, https://orcid.org/0000-0002-6679-0338 (RJDPvL); dprelec@mit.edu, https://orcid.org/0000-0002-9507-5368 (DP)

Received: September 9, 2021
Revised: February 15, 2022
Accepted: March 15, 2022
Published Online in Articles in Advance: June 7, 2022

https://doi.org/10.1287/mnsc.2022.4453

Abstract. This paper introduces the index $\tau = \ln\beta/\ln\delta$ as a measure of time inconsistency and vulnerability to self-control problems in the quasi-hyperbolic, beta-delta (β, δ) discounting model. We provide a preference foundation for τ and, consequently, a revealed preference definition of failed self-control. The τ index is independent of utility and has an intuitive interpretation as the maximum number of future selves who can disagree with the current self with respect to uniform deviations from an intertemporal plan. The index is also computable for continuous discount functions after an appropriate mapping of functions onto the (β, δ) family. The τ index thus provides a common yardstick for comparing temporal inconsistency across different functional forms.

History: Accepted by Manel Baucells, behavioral economics and decision analysis.
Supplemental Material: The data files are available at https://doi.org/10.1287/mnsc.2022.4453.

Keywords: quasi-hyperbolic discounting • intertemporal choice • time inconsistency

1. Introduction

Because of its simplicity and tractability, the quasi-hyperbolic β-δ model (Phelps and Pollak 1968, Laibson 1997) is the most widely used representation of nonexponential time preference. The model adds an additional discount β to the standard exponential (compound) discounting model, with discount factor δ, yielding

$$U(x_0, t_0; \ldots; x_n, t_n) = u(x_0) + \beta \sum_{i=1}^{n} u(x_i)\delta^{t_i}. \qquad (1.1)$$

In Equation (1.1), $u(x_i)$ is the utility of the outcome received at time point t_i (with $u(0) = 0$) and $t_0 = 0$. The "present-bias" β-parameter reflects (inversely) the additional weight assigned to immediate consequences, creating a wedge between the preferences of the current self and those of future selves. This wedge can lead to time-inconsistent preferences and costly precommitment strategies by sophisticated agents, or actual plan reversals and money-pumping of naïve ones (Strotz 1955, Pollak 1968, O'Donoghue and Rabin 1999).

The literature generally interprets β as a measure of time inconsistency and reduced self-control. For example, DellaVigna (2009, p. 318) refers to β as " … capturing the self-control problems."[1] However, such an interpretation is not entirely correct, because the discount factor δ also plays a role. To see this, consider two quasi-hyperbolic agents with the same value of $\beta = 1/2$, but different δ-s, say $\delta_1 = 1/3$ and $\delta_2 = 2/3$. Imagine that both will face a choice tomorrow ($t = 1$) between receiving 25 utility units immediately or waiting one more day ($t = 2$) for 60 utility units. Given these parameters, both agents know that tomorrow they will choose 25 (and give up the 60). From the perspective of today ($t = 0$), this is a good result for the impatient agent, but a poor result for the more patient one who would prefer to wait for 60.[2] Counterintuitively, the more patient agent is the one who fears the decision of their future self and is in the market for self-control. They would like to tie their hands and precommit to 60.

An index of time inconsistency should therefore include the δ parameter and β. Our example suggests that if β is held constant, inconsistency will increase with patience, δ. At the same time, it is evident that if δ is held constant, inconsistency will increase with present bias (smaller β). Qualitatively, inconsistency is created by a contrast between short- and long-run impatience—exacerbated by short-run impatience but reduced by long-run impatience. The index of time inconsistency that we derive expresses this contrast precisely as $\tau = \ln\beta/\ln\delta$.

The preference criterion for defining τ will be the maximum number of future selves who might fail to implement $t = 0$ preferences if presented with a temptation to deviate from the plan set by the $t = 0$ self. We first consider this for two timed outcomes, as in the introductory example. The plan in that case is just to wait for a larger outcome (e.g., 60) fixed at some future time point; the temptation is to sacrifice this large outcome for a smaller outcome (e.g., 25) that might arise at some intervening moment. We count the number of selves that would yield to temptation by choosing the smaller outcome, and, critically, *thereby conflict with the preferences of the $t = 0$ self.*

This criterion generates an index $\tau = \ln\beta/\ln\delta$ as the maximum number of future selves that can disagree with the current self. Section 2 provides graphical intuition for the result, which is followed by a formal proof in Section 3. The τ value is computed from (ordinal) preferences and does not involve the utility function.

Using τ we may rewrite the quasi-hyperbolic model as

$$U(x_0, t_0; \ldots; x_n, t_n) = u(x_0) + \sum_{i=1}^{n} u(x_i)\delta^{t_i+\tau}. \tag{1.2}$$

The discount function is $\delta^{t+\tau}$, where τ is a virtual extra delay that the agent adds to all outcomes not received immediately.[3] In effect, τ translates present bias into time units, equal to the maximum interval during which the agent is vulnerable to self-control problems. During this "vulnerable period," an option that is disliked by the current self could be chosen by a future self.

Going beyond the simple single outcome setting, we show in Section 4 that τ remains an upper bound on disagreement for more complex outcome streams. The plan in the more general case is an outcome stream selected by the $t = 0$ self, whereas temptations take the form of alternative streams of outcomes, offered at $t = 1$, $t = 2, \ldots$ The alternative streams are identical relative to their starting point, $t = 1$, $t = 2, \ldots$. Again, we consider the maximum number of future selves that would select the alternative stream when this choice contradicts the preferences of time zero self. We show that the single outcome setting is the worst-case scenario for self-control problems.

Similarly, considering general discounting functions, the worst-case scenario arises within the (β, δ) model. That is, one can map a general discount function onto the (β, δ) family and prove that number of deviating selves for the general function cannot exceed $\tau = \ln\beta/\ln\delta$.

Measuring τ is straightforward using existing methods for measuring quasi-hyperbolic discounting (Abdellaoui et al. 2010). In Section 6, we reanalyze the data of Tanaka et al. (2010) and show that the (δ, τ) parametrization has statistical advantages in that τ and δ are not correlated, whereas β and δ are strongly correlated on that data set. This leads to clearer interpretations of the impact of demographic variables.

2. Visual Derivation of τ

Before turning to formal results, we provide a graphical intuition for τ. We define agent A as more vulnerable to self-control problems than agent B if A disagrees in preference with a greater number of his future selves. The panels in Figure 1 show how to count the number of disagreeing selves within the $\beta - \delta$ model.

Example 2.1. Figure 1 displays scenarios of potential intertemporal conflict, in the form introduced by Ainslie (1975). The x axis shows calendar time in months. Outcomes are given in utility units. The y axis denotes discounted utility (DU) in a logarithmic scale, which linearizes discounted utility for the $\beta - \delta$ model.

As in our introductory example, there is a higher, more remote outcome with utility 60, scheduled for 12 months from now (time zero). The agent computes quasi-hyperbolic discounted utility (Equation (1.1)) with $\beta = 0.5$ and $\delta = 0.84$. The solid line in each panel indicates how the discounted utility of 60 increases as the agent moves closer to it. The line is linear because of logarithmic scaling on the y axis.

At time zero, the agent anticipates that a smaller outcome delivering 25 utility units will become available at some point before the 12 months are up. The dashed lines in Figure 1(a)–(c) indicate the discounted utility of this smaller outcome if that outcome becomes available at one of three different time points. The panels address two questions: (a) whether the smaller outcome will be selected, and (b) whether this is consistent with the preferences of the time zero self.

Figure 1(a) displays a scenario where the smaller option becomes available at month 9 and will be selected by the month 9 self because its immediate utility is greater than the utility of 60, discounted by the remaining 3 months. This is unattractive from the perspective of any earlier self, as the dashed line is below the solid line.

If the smaller option becomes available at month 7 (Figure 1(b)), it will again be selected, but this is no longer a problem for the month zero self who is indifferent between 25 at month 7 and 60 at month 12. If the smaller option becomes available after month 11 (Figure 1(c)), it will be rejected, which is consistent with the preferences of month zero self.

Conflict between month zero self and future selves can only arise between months 7 through 11. Month 7 marks the transition from unconflicted impatience (the agent always chooses the smaller sooner outcome) to vulnerability, and month 11 marks the transition from vulnerability to unconflicted patience (the agent always chooses the larger later outcome).

From the perspective of month zero self, the longer the vulnerable period, the greater the chance that their future self will make the wrong choice. From Figure 1(d) we can deduce that the duration of the vulnerable period is equal

Figure 1. (Color online) Intuition Behind τ

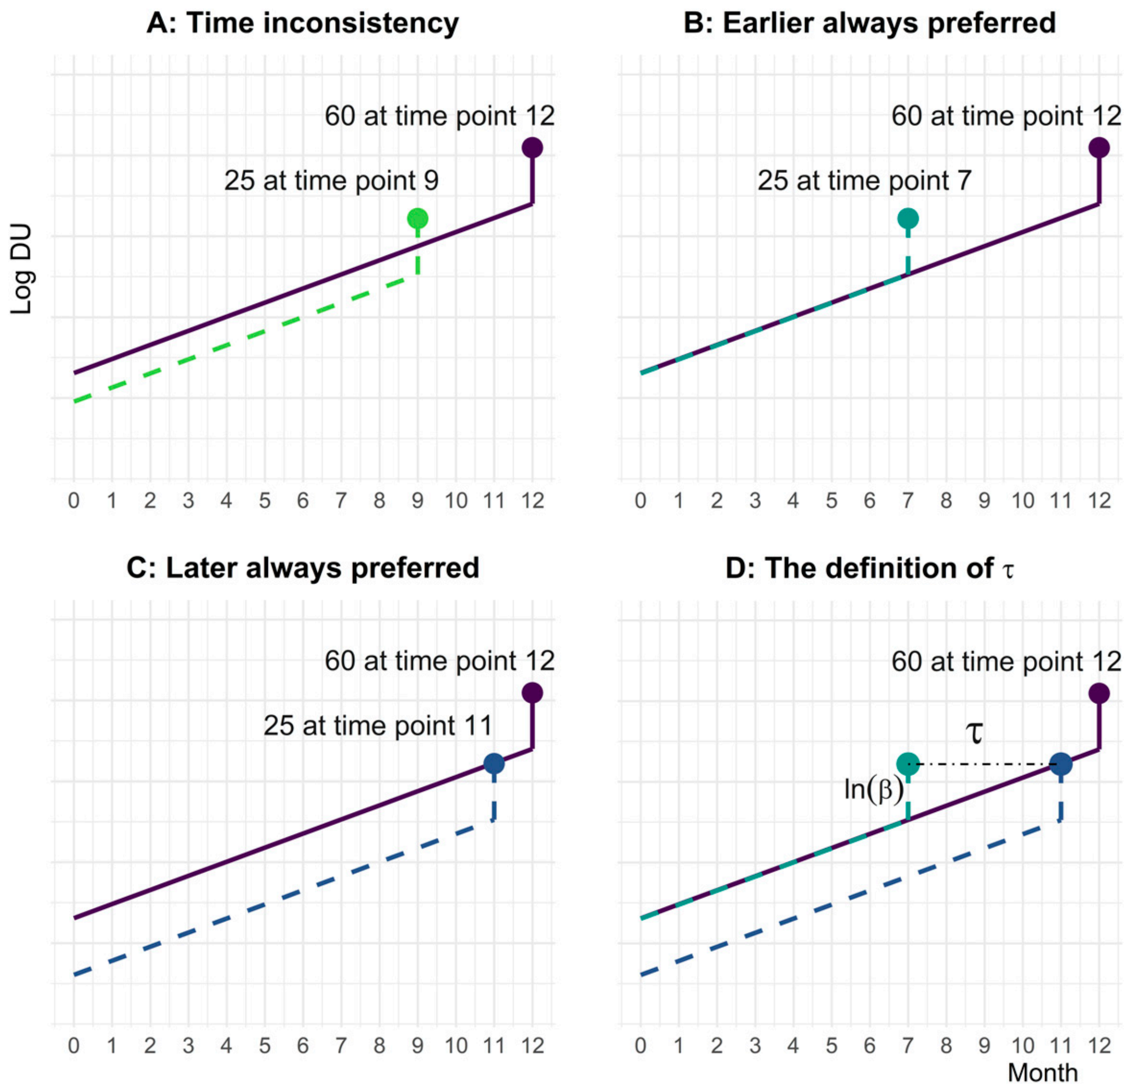

to the length of one side of a right triangle defined by the vertical step, $\ln(\beta)$, and the slope of the log DU line, $\ln(\delta)$. It is therefore equal to $\frac{\ln(\beta)}{\ln(\delta)} = \tau$. This shows that τ is the length of the vulnerable period and therefore an index of time inconsistent preferences.

In summary, the time until the receipt of the later outcome (60 at month 12 in Figure 1) can be subdivided into three distinct periods, depicted in Figure 2[4]:

(1) An initial period of *unconflicted impatience*, where the future self prefers the lower reward and the time zero self agrees;

(2) A *vulnerable* period of conflicted impatience, where the future self prefers the lower reward, but the time zero self disagrees; and

(3) A final period of *unconflicted patience*, where both the time zero and the future self prefer to wait for the later reward.

The lengths of the periods of unconflicted impatience and unconflicted patience depend on the utility difference between the outcomes and may equal zero. Changing the outcome utilities will shift the starting point of the vulnerable period but will not affect its duration, which, with the β and δ in our example, remains fixed at 4 (as long as all three periods are nonzero). Because τ is independent of utility, we can define it for any type of outcomes, including qualitative health states, environmental quality, or consumption.

3. Measure of Decreasing Impatience for Timed Outcomes

We consider a continuous strict preference relation $\succ_t$ over timed outcomes, where the subscript t indicates the time point at which the decision is made: $(l, s) \succ_t (h, r)$ indicates that at time point $t \leq s$ the agent

Figure 2. Vulnerable period for Figure 1, $l < h$, $t < s$

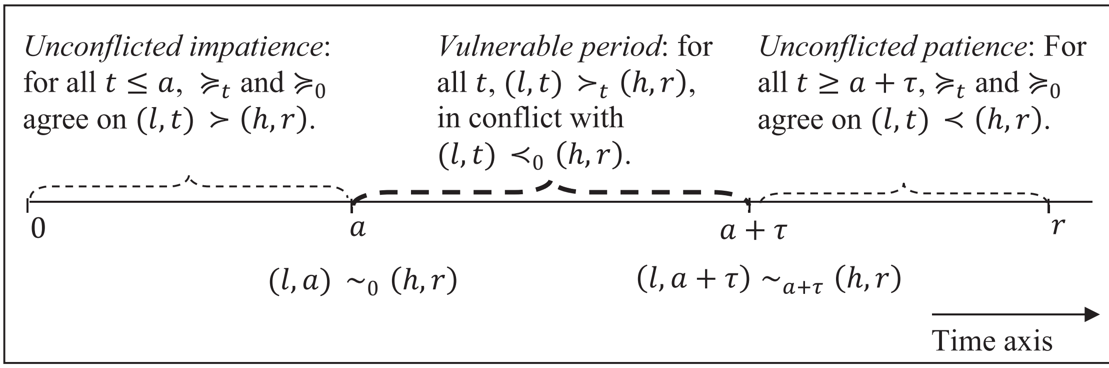

(strictly) prefers (l,s) to (h,r). Time points designate calendar time, with $s-t$ the time between s and t. We assume *time invariance*: only the differences between consumption time and decision time matter and not their location in time. Nontrivial choices between timed outcomes (l,s) and (h,r) always involve low and high outcomes $l < h$, and soon and remote consumption times $s < r$.

Under quasi-hyperbolic discounting, time inconsistency can only arise because of the immediacy effect and occurs at time points at which the lower outcome is received immediately (i.e., at $t = s$). Current (time-zero) selves who prefer the higher remote reward (h,r) may fear that when time point s arrives, their future time s self will choose the lower sooner reward (l,s):

$$(l,s) \prec_0 (h,r) \text{ and } (l,s) \succ_s (h,r) \text{ with } 0 < s. \tag{3.1}$$

For ease of presentation, we assume in the main text that preferences are nondegenerate:

$$(l,0) \prec_0 (h,\varepsilon) \text{ for some } 0 < \varepsilon \text{ and } (l,b) \succ_0 (h,r) \text{ for some } b > 0. \tag{3.2}$$

In words, if h comes soon enough, it can offset an immediate l and if r is remote enough then a sufficiently early l can still offset (h,r) even without the immediacy effect (implying $a > 0$ in Figure 2). By nondegeneracy, the period of unconflicted impatience, the vulnerable period, and the period of unconflicted patience (the three periods shown in Figure 2) are all nonempty. The appendix shows that if Equation (3.2) does not hold, the vulnerable period may be truncated at $t = 0$ or $t = r$.

We measure time inconsistency by counting the time points t besides s that are vulnerable to inconsistencies. Formally, we call time point t *vulnerable* if Equation (3.1) holds with t instead of s. Given $(l,s) \prec_0 (h,r) \prec_0 (l,b)$, there exists by continuity a unique time point $a > 0$ such that

$$(l,a) \sim_0 (h,r). \tag{3.3}$$

In Equation (3.3), a is the time point at which (l,t) without the immediacy effect is equivalent to (h,r). In the example of Section 2, $a = 7$. Because of the immediacy effect, $(l,a) \succ_a (h,r)$. As $(l,r) \prec_r (h,r)$, by continuity, there must exist an a' such that

$$(l,a') \sim_{a'} (h,r). \tag{3.4}$$

We define $\tau = a' - a$. Then $a' = a + \tau$ is the time point at which the immediacy effect "loses its bite": (l,t) with the immediacy effect is equivalent to (h,r). In the example of Section 2, $a + \tau$ equals 11. In between a and $a + \tau$, conflicts between the current and future selves can arise. Let us summarize.

Theorem 3.1. *Assume Equations* (3.1) and (3.2). *Then the vulnerable period is* $(a,\ a+\tau)$ *with* a *as in Equation* (3.3).

Theorem 3.1 shows that τ has a natural interpretation as the length of the vulnerable period. The larger τ is, the more an agent is exposed to dynamic inconsistencies. In the quasi-hyperbolic discount model, τ is constant and does not depend on $l, h,$ and r. This motivates our proposal to use τ as an index of time inconsistency.

4. Outcome Streams

We now extend our result to the general setting of outcome streams. The key observation is that moving from single outcomes to outcome streams cannot increase the length of the vulnerable period. In the quasi-hyperbolic model, timed outcomes are the worst-case scenario for self-control.

Let $\boldsymbol{x} = (x_1, t_1, \ldots, x_n, t_n)$ denote an outcome stream that gives money amount $x_j \geq 0$ at time point t_j, $j = 1, \ldots, n$, and nothing otherwise. Our notation implicitly assumes that $t_1 < \cdots < t_n$. Like in Section 3, we analyze preference reversals of the type:

$$(x_1, s_1; \ldots; x_n, s_n) \succ_t (y_1, r_1; \ldots; y_m, r_m) \text{ and}$$
$$(x_1, s_1; \ldots; x_n, s_n) \prec_{t'} (y_1, r_1; \ldots; y_m, r_m).$$

Under quasi-hyperbolic discounting, these preference reversals can only occur because of the immediacy effect, and we assume, without loss of generality, that it favors x. Consequently, $t = s_1$, and either $s_1 < r_1$ and $x_1 > 0$, or $s_1 = r_1$ and $x_1 > y_1$. The right-hand preference also holds if we replace t' by zero because then the immediacy effect for x is weakened:

$$(x_1, s_1; \ldots; x_n, s_n) \succ_{s_1} (y_1, r_1; \ldots; y_m, r_m) \text{ and}$$
$$(x_1, s_1; \ldots; x_n, s_n) \prec_0 (y_1, r_1; \ldots; y_m, r_m). \tag{4.1}$$

Like in Section 3, we measure the degree of time inconsistency by checking how much we can change the timing of outcomes such that the previous preference reversal is preserved. For $\varepsilon \in \mathbb{R}$, let $x^{\uparrow\varepsilon}$ denote the shift $(x_1, t_1 + \varepsilon; \ldots; x_n, t_n + \varepsilon)$ of x, where $t_1 + \varepsilon \geq 0$. Shifts $x^{\uparrow\varepsilon}$ preserve the preference reversal if

$$(x_1, s_1; \ldots; x_n, s_n)^{\uparrow\varepsilon} \succ_{s_1+\varepsilon} (y_1, r_1; \ldots; y_m, r_m) \text{ and}$$
$$(x_1, s_1; \ldots; x_n, s_n)^{\uparrow\varepsilon} \prec_0 (y_1, r_1; \ldots; y_m, r_m). \tag{4.2}$$

In keeping with Section 3, we call such $s_1 + \varepsilon$ *vulnerable*.

Theorem 4.1. *Under Equation (4.1), τ is the maximum length of the vulnerable period.*

Theorem 4.1 shows that the vulnerable period for outcome streams cannot exceed that for timed outcomes. Intuitively, the immediacy effect only affects the first outcome for outcome streams, whereas it affects all outcomes for timed outcomes. As the maximum length τ is reached in choices between timed outcomes (see Theorem 3.1), the maximum length of the vulnerable period is exactly equal to τ.

5. Simple Extension to General Discounting

Thus far, we have only considered the quasi-hyperbolic model. Apart from being familiar and tractable, it is also the only discounting model in which τ is constant and does not depend on the outcomes or time periods:

Observation 5.1. Quasi-hyperbolic discounting is the only nonconstant rate discounting model in which τ is constant, independent of the outcomes and time points.

In light of Observation 5.1, one may wonder what our results imply for more general discounting models with a continuous (decreasingly impatient) discount function $\varphi(t)$. In this section, we address this question.

With general discounting the utility of outcome stream $\boldsymbol{x} = (x_0,\ t_0; x_1, t_1; \ldots; x_n, t_n)$ with $t_0 = 0$ is equal to $u(x_0) + \sum_{i=1}^{n} \varphi(t_i) u(x_i)$. This equation may be expressed in a $\beta - \delta$ form but with a time-varying present bias parameter $\beta(t)$,

$$u(x_0) + \sum_{i=1}^{n} \beta(t_i) \delta_\varphi^{t_i} u(x_i), \tag{5.1}$$

and δ_φ defined as the asymptotic (long-run) discount factor $\delta_\varphi = e^{\lim_{t\to\infty} \frac{\varphi'(t)}{\varphi(t)}}$. If $\delta_\varphi = 1$, then the vulnerable period is unbounded. If $\delta_\varphi \neq 1$, then the length of the vulnerable period is determined by δ_φ and the limit of $\beta(t)$. This follows from $\beta(t)$ nonincreasing in t and τ decreasing in β (for given δ). For large enough t, $\beta(t)$ will vary by less than ϵ, and consequently, $\frac{\ln(\beta(t))}{\ln(\delta_\varphi)}$ will vary by less than $\frac{\ln(\beta(t)+\epsilon)}{\ln(\delta_\varphi)}$. Define the long run $\beta_\varphi = \lim_{t\to\infty} \beta(t)$. Then, by our results, the maximal length of the vulnerable period is $\tau_\varphi = \frac{\ln(\beta_\varphi)}{\ln(\delta_\varphi)}$.

For illustration, consider the discount function that blends fast (δ_F) and slow (δ_S) compound discounting,

$$\varphi(t) = (1-\beta)\delta_F^t + \beta\delta_S^t, \tag{5.2}$$

where the long run discount factor is evidently $\delta_\varphi = \delta_S$, yielding

$$\beta(t_i) = (1-\beta)\left(\frac{\delta_F}{\delta_S}\right)^t + \beta, \tag{5.3}$$

with long run $\beta_\varphi = \beta$, as $\delta_F < \delta_S$. The $\beta - \delta$ counterpart for this function is $\beta\delta_S^t$, and the maximum length of the vulnerable period is $\tau_\varphi = \frac{\ln(\beta)}{\ln(\delta_S)}$. Equation (5.2) can thus be interpreted as a continuous model for $\beta - \delta$ discounting, with the fast component representing present bias.

Several continuous discount functions yield $\delta_\varphi = 1$ for generic parameter values, which would imply an unbounded maximum vulnerability interval. Examples are the generalized hyperbola of Loewenstein and Prelec (1992), $\varphi(t) = (1+at)^{-\frac{b}{a}}$, and the constant relative decreasing impatience (CRDI) function, $\varphi(t) = \exp(-bt^a)$, of Ebert and Prelec (2007) and Bleichrodt et al. (2009). This is because the maximum value is determined by the asymptotic discount rate, which is zero for both functions.

By restricting outcomes to a finite interval $[0, T]$, we can put an upper bound on τ relative to a decision horizon T. Instead of computing asymptotic β, δ values, we define them at the horizon T. For example, for the CRDI function, we have $\delta_\varphi = e^{\frac{\varphi'(T)}{\varphi(T)}} = e^{-ab^aT^{a-1}}$ and $\beta_\varphi = \beta(T) = e^{-(1-a)b^aT^a}$ (applying Equation (5.1) with above δ_φ inserted). The τ upper bound with horizon T is therefore a constant fraction of T:

$$\tau = \frac{\ln\beta_\varphi}{\ln\delta_\varphi} = \frac{1-a}{a}T. \tag{5.4}$$

For example, the median estimate of a (for gains) in Abdellaoui et al. (2010) is 0.81. Then the vulnerability interval would be capped at 0.23 of the decision horizon (i.e., at slightly less than one-fourth of future selves who might contradict the preferences of the time zero self).[5]

6. Empirical Illustrations

At high levels of impatience (small δ), τ will be short, and preference reversals are rare. Low discount rates (i.e., values of δ closer to one) are often seen as reflecting more rationality but are also more likely to promote preference reversals. Dynamic inconsistencies are most likely for high values of δ and small values of β. The limiting case is where $\delta = 1$, but $\beta < 1$. Although in the literature these agents are often seen as close to rational (especially if β is also close to one), our index suggests that they are not. In fact, they only distinguish between the present and the future and assign the same weight to all future outcomes. To us, this seems rather irrational, and we believe our index correctly classifies such agents as irrational.

Empirical estimates in the literature are inconclusive about the relationship between β and δ (Table 1 shows some estimates). The meta-analysis of Imai et al. (2021) finds an overall β between 0.95 and 0.97, whereas Dean and Ortoleva (2019) find a positive correlation between the discounting of a present and a future amount and between two future amounts. Both papers suggest that the combination of a low β and a high δ is unlikely.

To provide an example of how τ might affect our interpretation of individual differences, we reanalyze Tanaka et al. (2010). The data are individual (risk and) time preferences and demographic variables. One hundred eight-one subjects answered 15 time preference questions by choosing between a money amount now and a larger amount in the future (three days to three months), with real incentives. The average payment was wages for about six to nine days.

Table 2 displays the Spearman correlation matrix of parameter estimates when computed for each individual on the basis of the 15 questions, using the constraints $0 \le \beta \le 1$, $0 < \delta < 1$, and $\tau \ge 0$ and ignoring demographics (we display Spearman rank correlations to mitigate the impact of outliers). Estimates of β and δ are strongly negatively correlated, suggesting that the two parameters are tapping a common individual difference variable: impatience. In contrast, estimates of δ and τ are not significantly related. This supports the hypothesis that impatience and decreasing impatience are to some extent distinct psychological constructs, captured by the δ and τ parameters.

Table 1. Selection of Empirical Studies That Yield Estimates of β, δ, and τ

	β	δ	τ
Fang and Silverman (2009)	0.34	0.87	7.7
Abdellaoui et al. (2013)	0.92	0.93	1.1
	0.97	0.62	0.1
Courtemanche et al. (2015)	0.80	0.75	0.8
Fang and Wang (2015)	0.34	0.88	8.4
Augenblick, Niederle, Sprenger (2015)	0.89	0.99	11.6
DellaVigna et al. (2017)	0.58	0.88	4.3
Blow et al. (2021)	0.84	0.96	4.3

Note. Discount factors δ are monthly, and τ is (thus) also expressed in months.

Table 2. Spearman Correlations for Coefficients Computed Separately for Each Individual Using Data from Tanaka et al. (2010)

	β	δ	τ
β	1.00		
δ	−0.53*	1.00	
τ	−0.45*	−0.13	1.00

Note. The correlation between β and δ is strongly negative ($z = 4.23$, $p < 0.001$) and different ($z = 2.93$, $p = 0.003$) from the nonsignificant correlation between δ and τ.

Using less correlated variables reduces the problem of multicollinearity in econometric analyses and tends to improve the coefficient estimates of individual predictors. In particular, it reduces their standard errors. To explore if the (δ, τ) parametrization yields different interpretation of the demographic variables, we repeated Tanaka et al.'s (2010) group-level analysis with nonlinear least squares regression on (β, δ) and on (δ, τ) (see Tanaka et al. 2010 for details, function (2), p. 567).

Table 3 shows the regression results for the two models. The meaning of most variables is self-evident. "Trusted agent" equals one for subjects who stored the money earned in the experiments, and "Risk payment" is the amount of money the subject received in the elicitation of the risk preferences. All variables are standardized.

The average β is 0.65 and is smaller than one ($p < 0.001$); β is not significantly related to any demographic variable. The average *daily* discount factor δ is equal to 0.9929, which implies 0.0742 on an annual basis, and is higher (indicating more patience) for subjects with higher age, education, income, and money won in the risk part.

Replacing (β, δ) by (δ, τ) appears to increase statistical power, as the p values for τ are lower than those for β with respect to six of nine demographic variables. Likewise, the standard errors for δ are generally also lower with (δ, τ) than with (β, δ). Although β is not significantly related to any demographic variable, τ is associated with two variables that are plausible proxies for wealth: people in the (richer) South of Vietnam and people who received more money in the first part of the experiment, which measured risk preferences. These individuals have on average higher δ and τ, indicating greater patience combined with greater *decreasing* impatience. According to our interpretation

Table 3. Regression Results for the (β, δ) and the (δ, τ) Models Using Data from Tanaka et al. (2010)

	Tanaka et al. (β, δ)		New framework (τ, δ)	
	β (%)	δ (%)	τ	δ (%)
Constant $(\beta_0, \delta_0, \tau_0)$	64.85	99.20	85.10	99.29
	1.88	0.07	15.87	0.07
	0.00	0.00	0.00	0.00
Chinese	−0.71	0.04	14.51	0.03
	1.67	0.07	9.38	0.02
	0.67	0.58	0.12	0.19
Trusted agent	−0.71	0.03	−1.58	−0.03
	1.32	0.04	2.21	0.02
	0.59	0.47	0.48	0.12
Age	1.18	0.18*	7.86	0.12*
	1.94	0.07	4.82	0.04
	0.54	0.02	0.10	0.00
Female	0.64	0.05	−3.30	0.00
	1.87	0.07	4.22	0.03
	0.73	0.44	0.44	0.93
Education	−3.33	0.15*	4.49	0.01
	2.03	0.07	5.36	0.03
	0.10	0.03	0.40	0.75
Income	1.08	0.10*	1.95	0.04*
	1.18	0.03	6.19	0.02
	0.36	0.00	0.75	0.02
Distance to market	2.49	0.02	−0.75	0.04
	2.13	0.07	9.19	0.04
	0.25	0.82	0.94	0.20
South Vietnam	−2.67	0.08	30.84*	0.18*
	2.28	0.08	15.18	0.09
	0.25	0.30	0.04	0.05
Risk payment	−1.75	0.15*	23.23*	0.14*
	2.14	0.08	10.94	0.06
	0.42	0.05	0.04	0.02
No. of observations	5,340		5,340	

Notes. The units for τ are days, and δ is computed on a daily basis and reported in percentages. For each variable the table shows (from top to bottom) the coefficient, the standard error, and the p value.
*Significance at the 5% level.

of τ, this suggests that wealthier individuals would be more in the market for self-control mechanisms. The improvement in explanatory power likely reflects the fact that τ and δ measure distinct psychological constructs, whereas β and δ both draw on two aspects of a single construct: impatience.

7. Related Work

The τ index is a revealed preference measure of the deviation from dynamic consistency, the normative requirement that preferences of the current self and those of future selves agree. The revealed preference literature has proposed other indices of deviations from rationality. Afriat (1973) and Varian (1990) measured these deviations by the amount of wealth that had to be taken away from the agent so that they behaved in line with their criterion of rationality. Apesteguia and Ballester (2015) showed that these measures fail to satisfy some desiderata and, instead, proposed the *swaps index*, which measures how many options should be removed from the choice set such that an agent's choices are consistent with a transitive and complete preference relation. In removing options, their index resembles what τ does in terms of future selves who violate dynamic consistency. In a sense, τ asks the question "how many timed outcomes should be removed from the choice set so that the observed preferences are dynamically consistent and can be approximated by an exponential discount function?" In our example of Section 2, the answer is "all outcomes that give 25 between time point 7 and 11." To study the similarity between τ and the swaps index and to explore the extent to which τ satisfies the conditions of Apesteguia and Ballester (2015) is an interesting topic for future research.

Ericson and Noor (2015) measure discounting using a delay function. Their approach is very general in that they give up the assumption that time and money are evaluated separately. They do not propose an index of decreasing impatience, however.

More closely related to the current measure is the approach in Prelec (1989; 2004). The idea is to assess decreasing impatience using two-stage decision problems where a precommitment option is available in the first stage. Time inconsistency is defined as the tendency by the "first stage self" to select a dominated outcome in order to prevent the "second stage self" from choosing something even worse. The main result is that one agent exhibits more preference for dominated options than another agent if and only if one log-discount function is a convex transformation of the other one. For the quasi-hyperbolic family, relative convexity delivers the same ordering of time inconsistency as does the comparison of τ ratios (Prelec 2004). However, the preference for dominated options seems less intuitive than the window of vulnerability proposed here, and the former criterion does not extend to general outcome streams.

Attema et al. (2010) and Rohde (2019) also defined indices of decreasing impatience, which can be measured independent of utility. Like Prelec (2004), their measures do not translate directly into the number of future selves disagreeing with the current self, nor do they apply to more general outcome streams. For this reason, they may be more relevant to experimental elicitation and measurement of time preference at the individual level. In contrast, the τ index offers an intuitive measure of vulnerability to self-control failure outside of the laboratory and is inspired by psychological analyses of self-control, notably the seminal work of Ainslie (1975). It also speaks more directly to the empirical economics literature, which has embraced the (β, δ) parametrization for theoretical tractability and convenience.

The τ index also provides a new diagnostic instrument for practical decision analysis, which aim is to help people make better decisions. Time inconsistencies are seen as undesirable and to have no place in prescriptive analyses. Measuring τ, which is easy as we showed in Section 5 and can be done, for example, using the method in Abdellaoui et al. (2010), can help identify clients that are particularly prone to time inconsistencies.

8. Discussion

The quasi-hyperbolic model is sometimes used to assess whether a variable influences departures from the exponential, time-consistent norm. The standard assumption has been that someone with a smaller β will be less time consistent and therefore more vulnerable to self-control problems. In this paper we have criticized this view. However, because the proposed (δ, τ) parametrization will fit the data just as well as the standard (β, δ), one might ask what is gained by choosing one rather than the other form?

The interpretive issue at stake here is normative and relevant to policy. Econometric parameter estimates can clarify the causes of self-harming activities and focus attention on specific remedies. In this context, it is important to know whether harmful behavior is due to simple time preference or failed self-control. For example, if cigarette smokers care little about the future, then their choices may be consistent with the exponential model and in that sense may be rational. Smoking may be a type of optimal self-medication, as implied by Becker and Murphy (1988). However, if smokers are time inconsistent, then other interpretations of their actions become available, such as sophisticated fatalism ("I believe I cannot stop smoking hence I might as well smoke now") or naïve optimism leading to procrastination ("I believe I will quit tomorrow and therefore I can smoke now") (O'Donoghue and Rabin 1999). Because the preferences of different temporal selves are already in conflict, the policy maker may feel justified in acting paternalistically on behalf of one self against another, for example, by imposing penalties or banning certain goods altogether (Gruber and Kőszegi 2001).

Of course, whether impatience and self-control are distinct psychological dimensions is an empirical rather than a modeling question. It is possible that individuals who care little about the future as measured by δ also exhibit more time-inconsistent preferences, as measured by τ. In that case, by combining impatience and self-conflict into a cardinal index of *present-bias intensity*, the β parameter may provide a robust indicator of impulsive, low-quality decision making. Its charm may reside precisely in the (con)fusion of two characteristics that should in theory be kept apart—the pure lack of concern about the future per se, and the sense that one's future actions will not match current intentions. The theoretical point we emphasize here is that such aggregation of two conceptually distinct dimensions—impatience and time inconsistency—into a single number should be done with eyes open.

9. Conclusion

We have shown that time inconsistency within the popular quasi-hyperbolic discounting is captured by writing the model as $f(t) = \delta^{t+\tau}$ rather than by the usual $f(t) = \beta\delta^t$ (for $t > 0$, with $f(0) = 1$). This reformulation separates impatience (measured by δ) and decreasing impatience (measured by τ). The τ index may be extended to continuous functions, providing a common time-denominated yardstick for comparing different parametrizations and functional forms. The index is based on revealed preference and has a natural interpretation as the period of vulnerability to dynamic inconsistencies and hence to self-control problems. Being easy to measure, τ is a useful new tool for prescriptive decision analysis and policy.

Acknowledgments

Helpful comments were received from the department editor Manel Baucells, an associate editor, two anonymous reviewers, and from Jose Apesteguia, Adam Sanjurjo, and Peter P. Wakker.

Appendix

In the case of $\delta = \beta = 1$, the agent is perfectly rational with no DI or vulnerability. Then all our results follow. In the case of $\delta = 1$ but $\beta < 1$, vulnerability periods can be infinite, and all our results follow again. We assume in this appendix that $\delta < 1$ so that $\ln(\delta)$ is well defined and can appear in denominators.

A.1. Measure of Decreasing Impatience

Lemma A.1. *In the quasi-hyperbolic model, $\succcurlyeq^*$ exhibits more decreasing impatience than $\succcurlyeq$ if and only $\tau^* = \frac{\ln(\beta^*)}{\ln(\delta^*)} \geq \tau = \frac{\ln(\beta)}{\ln(\delta)}$.*

Proof. According to proposition 1 in Prelec (2004), quasi-hyperbolic discount function φ^* reflects more decreasing impatience than quasi-hyperbolic discount function φ if and only if $\ln(\varphi^*)$ is a convex transformation of $\ln(\varphi)$. Both $\ln(\varphi^*)$ and $\ln(\varphi)$ take a value of zero at $t = 0$, and we describe only their values at $t \neq 0$. Expressing $\ln(\varphi^*)$ as a function of $\ln(\varphi)$ gives

$$\ln(\varphi^*(t)) = \ln(\beta^*) - \frac{\ln(\delta^*)}{\ln(\delta)}\ln(\beta) + \frac{\ln(\delta^*)}{\ln(\delta)}\ln(\varphi(t)).$$

Given a value of zero at zero, this transformation of $\ln(\varphi(t))$ is convex if and only if $\ln(\beta^*) - \frac{\ln(\delta^*)}{\ln(\delta)}\ln(\beta) \leq 0$, which holds if and only if $\frac{\ln(\beta^*)}{\ln(\delta^*)} \geq \frac{\ln(\beta)}{\ln(\delta)}$. □

Prelec (2004) showed that two discount functions are related through a power transformation if they have the same (in our notation) τ, which suggests that τ may serve as an index of decreasing impatience. However, Prelec did not provide an ordering result as in Lemma A.1 and gave no derivations. In particular, he did not handle the discontinuity at $t = 0$.

A.2. Proof of Theorem 3.1

By continuity and impatience a exists and is between b in Equation (3.2) and r. For all $t \le a$, $\succcurlyeq_t$ and $\succcurlyeq_0$ agree on $(l,t) \succ (h,r)$ and on the preferences between (l,t') and (h,r) for all other t' (the latter do not involve the immediacy effect). No vulnerability arises.

Equation (3.3) implies $\beta\delta^a u(l) = \beta\delta^r u(h)$, implying $\delta^{a+\tau}u(l) = \beta\delta^r u(h)$, $u(l) = \beta\delta^{r-a-\tau}u(h)$, and

$$(l, a+\tau) \sim_{a+\tau} (h,r). \tag{A.1}$$

Here $a + \tau < r - \varepsilon$ by the left-hand side of Equation (3.2) (which is equivalent to $(l, r-\varepsilon) \prec_{r-\varepsilon} (h,r)$). For all $t \ge a+\tau$, $\succcurlyeq_t$ and $\succcurlyeq_0$ agree on $(l,t) \prec (h,r)$ and on the preferences between (l,t') and (h,r) for all other t' (the latter do not involve the immediacy effect). No vulnerability arises.

For all s with $a < s < a+\tau$, we have Equation (3.1), where the left-hand side follows from Equation (3.3) and the right-hand side from Eq. A.1. Hence these s are vulnerable. □

Theorem 3.1 without Equation (3.2).

We now consider the general case, without Equation (3.2) (Figure 2). Let $-\infty < a < r$ be the unique real number solving the equation

$$\beta\delta^a u(l) = \beta\delta^r u(h).$$

It implies

$$\delta^{a+\tau}u(l) = \beta\delta^r u(h).$$

We consider a number of cases.

Case 1: $a \le -\tau$. Then (h,r) is always preferred, even if l is received immediately at time point $t = 0$. This is the trivial case where Equation (3.1) never arises.

Case 2: $-\tau \le a < 0$. Now the vulnerable period is not $(a,\ a+\tau)$ but rather its truncation at 0, being $(0, a+\tau)$. This case was excluded in the main text by the second preference in Equation (3.2).

Case 3: $0 \le a \le r-\tau$. This is the case of Theorem 3.1, with vulnerable period $(a,\ a+\tau)$.

Case 4: $r - \tau < a < r$. The vulnerable period is not $(a,\ a+\tau)$ but rather its truncation at r, being (a, r). This case was excluded in the main text by the first preference in Equation (3.2).

A.3. Proof of Theorem 4.1

We write $QH(y)$ for the quasi-hyperbolic discounted utility of y at time point s_1. Consider the shifts ε such that $(x_1, s_1; \ldots; x_n, s_n)^{\uparrow\varepsilon} \succ_{s_1+\varepsilon} (y_1, r_1; \ldots; y_m, r_m)$. This implies $U(x_1) + \beta\sum_{j=2}^n \delta^{s_j - s_1} U(x_j) > \delta^{-\varepsilon}\ QH(y)$. Because $\beta < 1$, it is also true that, $U(x_1) + \sum_{j=2}^n \delta^{s_j-s_1} U(x_j) > \delta^{-\varepsilon}\ QH(y)$. Because $\beta\delta^{-\tau} = 1$,

$$\beta\delta^{-\tau}U(x_1) + \beta\sum_{j=2}^n \delta^{s_j - s_1 + \tau} U(x_j) > \delta^{-\varepsilon}\ QH(y),$$

or (multiplying by $\delta^{s_1+\varepsilon}$)

$$\beta\delta^{s_1+\varepsilon-\tau}U(x_1) + \beta\sum_{j=2}^n \delta^{s_j+\varepsilon-\tau} U(x_j) > \delta^{s_1}\ QH(y).$$

The last inequality concerns preference $\succ_0$ and the shift $x^{\uparrow\varepsilon-\tau}$. It gives a preference opposite to the second one in Equation (4.2) for the shift ε. Hence, Equation (4.2) cannot hold for both a shift by ε, and a shift by $\varepsilon - \tau$. If for any ε, the outcomes of x are shifted more to the present than τ, the previous inequalities become stronger and favor x more. This implies that the set of vulnerable shifts cannot contain shifts that are further apart than τ. Numerical examples show that the length of the vulnerable period can indeed be less than τ for some x, y. We saw in Section 3 that for timed outcomes the vulnerable period has length τ, and hence its maximum length is τ.

A.4. Proof of Observation 5.1

Assume a general discounted utility model $u(x_0) + \sum_{i=1}^n \varphi(t_i)\, u(x_t)$. Suppose τ is constant. Let l, h, a, t be such that $(l,a) \sim_0 (h,t)$. Then by the definition of τ: $(l, a+\tau) \sim_{a+\tau} (h,t)$. For some $e \in \mathbb{R}$, let h' be such that $(l, a+e) \sim_0 (h', t+e)$. Because τ is constant, we also have $(l, a+\tau+e) \sim_{a+\tau+e} (h', t+e)$. Hence, $\varphi(t-a-\tau) = \frac{\varphi(t)}{\varphi(a)} = \frac{\varphi(t*e)}{\varphi(a+e)}$. Consequently, $\ln(\varphi(t)) - \ln(\varphi(a)) = \ln(\varphi(t+e)) - \ln(\varphi(a+e))$, which implies that $\ln(\varphi)$ is linear: $\ln(\varphi) = a + bt$. Thus, $\varphi = e^a e^{bt} = \beta e^{\delta t}$ is quasi-hyperbolic. □

Endnotes

[1] Other examples suggesting that β measures time inconsistency are DellaVigna and Malmendier (2004, p. 360), Fang and Silverman (2009, p. 1048), Meier and Sprenger (2010, p. 200), Kaur et al. (2015, p. 1240), Augenblick et al. (2015, p.1092), Augenblick and Rabin (2019, table 1), Yoon (2020, p. 5851), Lockwood (2020, table 1), Heidhues and Strack (2021, p. 2595), and Imai et al. (2021, pp. 1788–1789).

[2] The $\delta_2 = 2/3$ agent today prefers 60, as $60 \times (1/2) \times (4/9) = 13.33 > 25 \times (1/2) \times (2/3) = 8.33$, but knows that tomorrow they will choose 25, as $25 > 60 \times (1/2) \times (2/3) = 20$. The $\delta_1 = 1/3$ agent also knows that tomorrow they will choose 25, but this is also the preferred choice today, as $25 \times (1/2) \times (1/3) = 4.17 > 60 \times (1/2) \times (1/9) = 3.33$.

[3] This follows from $\beta\delta^t = e^{\ln\beta + t\ln\delta} = e^{\left(t+\frac{\ln\beta}{\ln\delta}\right)ln\delta} = e^{(t+\tau)\ln\delta} = e^{\ln\delta^{t+\tau}} = \delta^{t+\tau}$.

[4] Exceptions arise for degenerate cases where the period gets truncated at $t = 0$ or $t = r$ and one or both of the other periods are absent.

[5] Note that τ is an upper bound on the vulnerability interval. This is true for continuous functions and for $\beta - \delta$. However, with $\beta - \delta$, the upper bound can be tight with the right combination of utilities and dates, as in Figure 1. With continuous functions, the bound will in general not be tight.

References

Abdellaoui M, Attema AE, Bleichrodt H (2010) Intertemporal tradeoffs for gains and losses: An experimental measurement of discounted utility. *Econom. J.* 120(545):845–866.

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

Afriat SN (1973) On a system of inequalities in demand analysis: An extension of the classical method. *Internat. Econom. Rev.* 14(2):460–472.

Ainslie G (1975) Specious reward: A behavioral theory of impulsiveness and impulse control. *Psych. Bull.* 82(4):463–496.

Apesteguia J, Ballester MA (2015) A measure of rationality and welfare. *J. Political Econom.* 123(6):1278–1310.

Attema AE, Bleichrodt H, Rohde KIM, Wakker PP (2010) Time-tradeoff sequences for analyzing discounting and time inconsistency. *Management Sci.* 56(11):2015–2030.

Augenblick N, Niederle M, Sprenger C (2015) Working over time: Dynamic inconsistency in real effort tasks. *Quart. J. Econom.* 130(3):1067–1115.

Augenblick N, Rabin M (2019) An experiment on time preference and misprediction in unpleasant tasks. *Rev. Econom. Stud.* 86(3):941–975.

Becker GS, Murphy KM (1988) A theory of rational addiction. *J. Political Econom.* 96(4):675–700.

Bleichrodt H, Rohde KIM, Wakker PP (2009) Non-hyperbolic time inconsistency. *Games Econom. Behav.* 66(1):27–38.

Blow L, Browning M, Crawford I (2021) Non-parametric analysis of time-inconsistent preferences. *Rev. Econom. Stud.* 88(6):2687–2734.

Courtemanche C, Heutel G, McAlvanah P (2015) Impatience, incentives and obesity. *Econom. J.* 125(582):1–31.

Dean M, Ortoleva P (2019) The empirical relationship between nonstandard economic behaviors. *Proc. Natl. Acad. Sci. USA* 116(33): 16262–16267.

DellaVigna S (2009) Psychology and economics: Evidence from the field. *J. Econom. Literature* 47(2):315–372.

DellaVigna S, Lindner A, Reizer B, Schmieder JF (2017) Reference-dependent job search: Evidence from Hungary. *Quart. J. Econom.* 132(4):1969–2018.

DellaVigna S, Malmendier U (2004) Contract design and self-control: Theory and evidence. *Quart. J. Econom.* 119(2):353–402.

Ebert JEJ, Prelec D (2007) The fragility of time: Time-insensitivity and valuation of the near and far future." *Management Sci.* 53(9): 1423–1438.

Ericson KM, Noor J (2015) *Delay Functions as the Foundation of Time Preference: Testing for Separable Discounted Utility.* National Bureau of Economic Research, Cambridge, MA.

Fang H, Silverman D (2009) Time-inconsistency and welfare program participation: Evidence from the NLSY. *Internat. Econom. Rev.* 50(4):1043–1077.

Fang H, Wang Y (2015) Estimating dynamic discrete choice models with hyperbolic discounting, with an application to mammography decisions. *Internat. Econom. Rev.* 56(2):565–596.

Gruber J, Kőszegi B (2001) Is addiction "rational"? Theory and evidence. *Quart. J. Econom.* 116(4):1261–1303.

Gruber J, Kőszegi B (2004) Tax incidence when individuals are time-inconsistent: The case of cigarette excise taxes. *J. Public Econom.* 88(9–10):1959–1987.

Heidhues P, Strack P (2021) Identifying present bias from the timing of choices. *Amer. Econom. Rev.* 111(8):2594–2622.

Imai T, Rutter TA, Camerer CF (2021) Meta-analysis of present-bias estimation using convex time budgets. *Econom. J.* 131(636):1788–1814.

Kaur S, Kremer M, Mullainathan S (2015) Self-control at work. *J. Political Econom.* 123(6):1227–1277.

Laibson DI (1997) Golden eggs and hyperbolic discounting. *Quart. J. Econom.* 112(2):443–477.

Lockwood BB (2020) Optimal income taxation with present bias. *Amer. Econom. J. Econom. Policy* 12(4):298–327.

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

Meier S, Sprenger C (2010) Present-biased preferences and credit card borrowing. *Amer. Econom. J. Appl. Econom.* 2(1):193–210.

O'Donoghue T, Rabin M (1999) Doing it now or later. *Amer. Econom. Rev.* 89(1):103–124.

Phelps ES, Pollak RA (1968) On second-best national saving and game-equilibrium growth. *Rev. Econom. Stud.* 35(2):185–199.

Pollak RA (1968) Consistent planning. *Rev. Econom. Stud.* 35(2):201–208.

Prelec D (1989) Decreasing impatience: Definition and consequences. Working paper, Division of Research, Harvard Business School, Cambridge, MA.

Prelec D (2004) Decreasing impatience: A criterion for non-stationary time preference and "hyperbolic" discounting. *Scandinavian J. Econom.* 106(3):511–532.

Rohde KIM (2019) Measuring decreasing and increasing impatience. *Management Sci.* 65(4):1700–1716.

Strotz RH (1955) Myopia and inconsistency in dynamic utility maximization. *Rev. Econom. Stud.* 23(3):165–180.

Tanaka T, Camerer CF, Nguyen Q (2010) Risk and time preferences: Linking experimental and household survey data from Vietnam. *Amer. Econom. Rev.* 100(1):557–571.

Varian HR (1990) Goodness-of-fit in optimizing models. *J. Econometrics* 46(1–2):125–140.

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