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

Han Bleichrodt,[1] Rogier J.D. Potter van Loon[2],

Drazen Prelec[3], and Peter P. Wakker[4]

April, 2017

Abstract

This paper introduces an index $\tau = \ln \beta / \ln \delta$ for the popular quasi-hyperbolic (β, δ) discounting, to replace β as a measure of time inconsistency and, thus, of vulnerability to self-control problems. Pros are that, unlike β, τ is directly based on revealed preference, has a preference foundation, reflects pure time attitude, and is independent of utility and of δ. We prove that the number of future selves who can disagree with the current self is at most τ. This suggests expressing $\beta\delta^t$ as $\delta^{\tau+t}$. We illustrate the (δ, τ) parametrization by reanalyzing data in Tanaka et al. (2010).

[1] bleichrodt@ese.eur.nl, Dept. of Economics, H13-27, Erasmus University, P.O. Box 1738, 3000 DR Rotterdam, the Netherlands

[2] rottervanloon@ese.eur.nl, Dept. of Economics, Erasmus University, P.O. Box 1738, 3000 DR Rotterdam, the Netherlands

[3] dprelec@mit.edu, E62-540, Sloan School, MIT, Cambridge MA 02139

[4] wakker@ese.eur.nl, Econometric Institute, Erasmus University, P.O. Box 1738, Rotterdam, 3000 DR, the Netherlands

1. Introduction

Because of its simplicity and tractability, the quasi-hyperbolic β-δ model (Phelps & Pollak 1968, Laibson 1997) is the most widely used representation of decreasingly impatient ("present biased") time preferences. The model takes the standard exponential (compound) discounting equation as starting point, with discount factor δ. It then assigns an additional discount β to all future time points, yielding:

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

where $u(x_i)$ is the utility of the outcome received at time point t_i (with $u(0) = 0$) and $t_0 = 0$. The β-parameter (inversely) reflects the additional weight assigned to immediate consequences, creating a wedge between the preferences of the current self and future selves. A smaller β implies a greater utility loss for the future selves. Decreasing impatience (DI) can lead to time-inconsistent preferences and costly pre-commitment strategies by sophisticated agents, or actual plan reversals and money-pumping of naïve ones (Strotz 1956, Pollak 1968, O'Donoghue & Rabin 1999).

In empirical work, the quasi-hyperbolic model often serves as a diagnostic instrument, revealing that one or other group is more deviant relative to the exponential, time-consistent norm. The standard view in the literature, accepted without debate, is that β provides the appropriate diagnostic measure of deviance, assuming that someone with a smaller β will be less time-consistent and therefore more vulnerable to self-control problems. For example, in an influential review, DellaVigna (2009, p.318) refers to β as "…*capturing the self-control problems*." We are, however, not aware of a revealed preference basis for this common interpretation of β. Further, β interacts with utility (see Figure 1 below) and in this sense is not a pure index of time preference.

Using Prelec's (2004) definition of comparative time inconsistency, we introduce a new measure of time inconsistency, the ratio $\tau = \ln \beta / \ln \delta$. We provide a revealed preference foundation and show that τ reflects pure time preference; i.e., it is not affected by utility. The parameter τ is present-bias expressed in time units, as a "virtual extra delay" for all outcomes that are not immediate. As we show, τ is also the maximum temporal "vulnerable period," during which an option that is disliked by the current self would be chosen if made available to a future self.

Prelec (2004) already showed that two sets of preferences fall in the same DI class if and only if the associated discount functions are related by a power transformation. For quasi-hyperbolic discounting, this implies that the two discount functions have the same ratio $\tau = \ln\beta / \ln\delta$, which already suggests that τ captures the essence of DI. The familiar β-parameter is precisely the compounded discount factor $\beta = \delta^{\tau}$ applied to this virtual delay. Given this result, we see that the effect of β on DI is confounded by δ. Example 2.1 will show that, correspondingly, the effect of β on DI is affected by utility.

The question of vulnerability is important in drawing policy inferences from laboratory or econometric parameter estimates of (1.1). Such estimates can shed light on the causes of self-harming behavior and focus attention on specific remedies. 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. However, if smokers are time-inconsistent then other interpretations of their behavior become available, such as sophisticated fatalism ("I believe I cannot stop smoking hence I might as well smoke now"), emotional choice (Fudenberg & Levine 2011), 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, e.g., by imposing penalties or banning certain goods altogether (Gruber & Köszegi 2001).

This paper begins with a geometric derivation of index τ (Figure 1) followed by behavioral foundations for τ. The usefulness of the new index is illustrated by a re-analysis of inter-temporal choice data from Tanaka et al. (2010). We find that τ exhibits less correlation with impatience δ than β does, and also suggests differences in relations between time attitudes and demographic variables.

2. A visual argument

Informally, we say that person A is more vulnerable to self-control problems than person B if A disagrees in preference with a greater number of future selves. In the (β, δ) model using β and δ as in Eq. 1.1, disagreement is promoted by lower values of β, but also by higher values of δ, as we will see. That is, these parameters interact here. If we rewrite $\beta\delta^{t}$ as $\delta^{t-\tau}$ then only τ affects disagreement and δ no longer does, removing the interaction. Before stating our formal results, we present a diagram that conveys the gist of the argument.

EXAMPLE 2.1. Figure 1 displays an inter-temporal dilemma in a representation due to Ainslie (1975). Outcomes are described in utility units. DU, displayed on the y-axis, denotes discounted utility. Outcomes can be equated with money if utility is linear, and DU on the y-axis then is present value. The dilemma that we consider involves two outcomes, a higher, more remote outcome of 60 and a lower, sooner outcome of 25.

From the vantage time point zero, 60 at month 12 is preferred to 25 at month 9 because the DU of the former, indicated by the solid line, exceeds the DU of the latter, indicated by the dashed line, at $t = 0$. At month $s = 9$, however, 25 is preferred if offered immediately, due to the discontinuous jump of the dashed line there. Hence the preferences of the month zero self are inconsistent with the preferences of the month 9 self. They are also inconsistent with any other self between months $a = 7$ and $a + \tau = 11$. Month 7 marks the transition from unconflicted impatience to vulnerability, and month 11 marks the transition from vulnerability to unconflicted patience. The triangle in bold shows that the vulnerable period exactly equals the height $-\ln(\beta)$ divided by the slope of the line, $-\ln(\delta)$, i.e., τ.

Figure 1. Present value of four timed outcomes

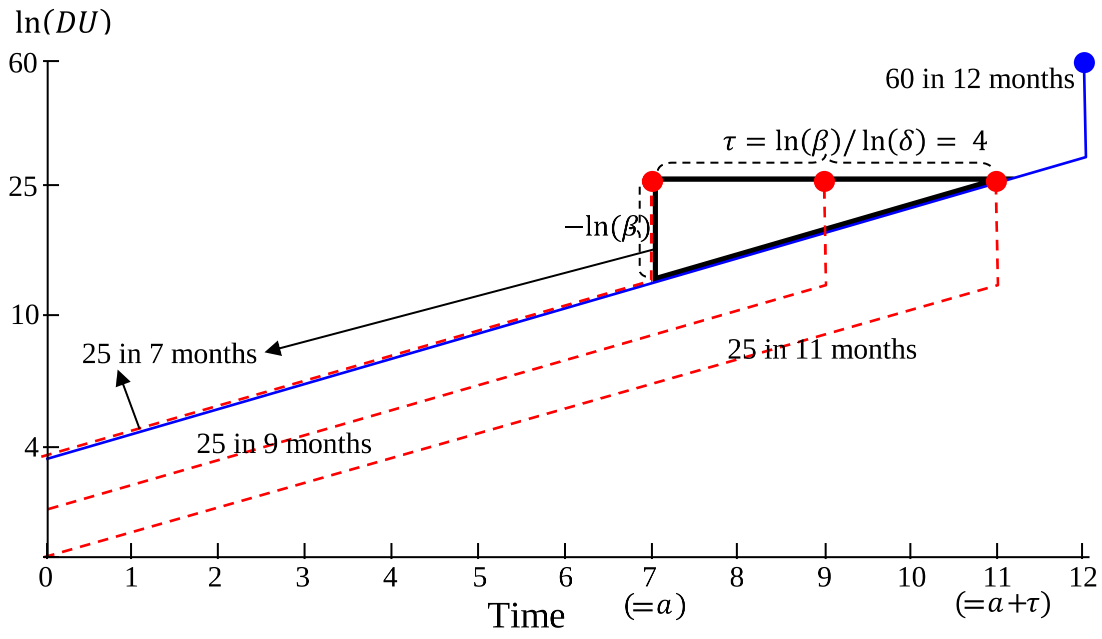

Four timed outcomes (60,12), (25,7), (25,9), and (25,11) are indicated by dots. The x-axis refers to calendar time in months relative to a starting date, which for convenience we set at zero. A line emanating from a dot indicates the DU of the corresponding timed outcome at preceding time points. The DU of various options at the considered decision time is shown on the y-axis with logarithmic scaling. For example, at time $t = 4$, the DU of receiving 25 in $s = 9$ months slightly exceeds 4. We chose $\beta = 0.5$, β = 0.5, δ = .84, and utility values to obtain convenient numbers.

□

In the example, we can infer τ from the preferences without knowing utility U. All we need for our analysis is that the utility of 60 exceeds that of 25. In particular, unlike preceding studies that estimated β (e.g. Augenblick, Niederle, & Sprenger 2015, footnote 19 and Eqs. 3 and 5), we need no commitment to a parametric utility family. Thus, outcomes need not be monetary and may refer to qualitative health states or consumption, for instance. As Figure 1 shows, β measures a utility loss which, of course, is affected by the utility function.

Whether β or τ is more relevant in applications depends on the context. If the size of the commitment necessary to overcome procastination is important, then β is relevant. If the period of commitment necessary to overcome procastination is important, then τ is relevant. We do not claim that one index is always preferable over the other. However, τ has advantages over β: τ reflects pure time preference—which is the topic of this paper—, it has a preference foundation, and it requires no knowledge about utility.

If the date of the later reward is held fixed (60 at $t = 12$ in the Figure), then the intervening time divides into three distinct periods, depicted in Figure 2:[5]

(1) An initial period of un-conflicted impatience, where the future self prefers the lower reward and this is OK with the current self;

(2) A *vulnerable* period of conflicted impatience, where the future self prefers the lower reward, but this is not OK with the current self;

(3) A final period of un-conflicted patience, where both the present and the future self prefer to wait for the later reward.

[5] 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.

Figure 2. Vulnerable period for Figure 1

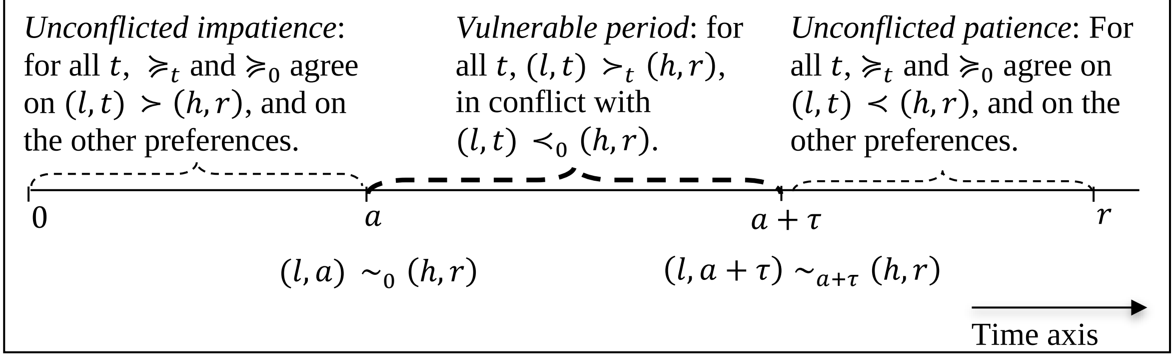

3. **A measure of decreasing impatience for timed outcomes**

The next two theorems formalize our claims. The first focuses on choices between timed outcomes (single outcomes) as in Figure 1, and the second, in §4, considers choices between general outcome streams.

The preference relation is subscripted by the time point at which the decision is made: $(l,s) \succ_t (h,r)$ indicates that at time point $t \le s$ the person prefers (l,s) to (h,r). In this paper, time points designate calendar time, with for instance $s-t$ the time interval between s and t. Throughout this paper, we make the common assumption of age independence; i.e., only the differences between consumption time and decision time matter. Nontrivial choices between timed outcomes (l,s) and (h,r) always involve low and high outcomes $l < h$, and soon and remote consumption times $s < r$. With quasi-hyperbolic discounting, time inconsistency can only arise due to the immediacy effect, where the lower outcome is immediate ($t = s$). Then the time-zero self may fear that, against his wish, the time-s self will choose the lower reward:

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

For ease of presentation, in the main text we focus on non-degenerate preferences (that allow the exact measurement of the degree of impatience):

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

That is, h is sufficiently big to compete with an immediate l if h comes soon enough (implying $a+\tau < r$ in Figure 2), and r is sufficiently remote to ensure that a sufficiently early l without the immediacy effect can still compete with (h,r)

(implying $a > 0$ in Figure 2). This way the three periods in §2 are nonempty. The appendix shows that in other cases, where Eq. 3.2 does not hold, the vulnerable period may be truncated at $t = 0$ or $t = r$.

We quantify the degree of decreasing impatience by inspecting which early time points t besides s are vulnerable to inconsistencies. Formally, we call time point t *vulnerable* if Eq. 3.1 holds with t instead of s. The proof of the following theorem will show that, given Eqs. 3.1 and 3.2, there exists a unique time point $a > 0$ such that

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

In Eq. 3.3, a is the time point at which (l, t) without the immediacy effect is equivalent to (h, r). Then $a + \tau$ is the time point at which (l, t) with the immediacy effect is equivalent to (h, r). In between these two time points, conflicts can arise.

THEOREM 3.1. Assume Eqs. 3.1 and 3.2. Then the vulnerable period is $(a, a + \tau)$ with a as in Eq. 3.3. □

Proofs are in the appendix. Theorem 3.1 shows that τ has a natural interpretation as the length of the vulnerable period. The larger τ is, the more a decision maker is prone to dynamic inconsistencies. An extreme case occurs if $\delta = 1$ and $\beta < 1$. Then the vulnerability period can be infinitely large. If the present self has a conflict with one future self, then it indeed has a conflict with all future selves. This combination of parameters is unlikely to occur in practice because $\delta = 1$ reflects a high level of rationality, higher than $\beta = 1$, let be that it would go together with the highly irrational $\beta < 1$.

4. Streams of outcomes

Section 3 considered preferences between two timed outcomes. This section extends our result to the general setting of outcome streams. The key observation is that moving from single outcomes to streams of outcomes cannot increase the length of the vulnerable period in the quasi-hyperbolic model. In this sense, timed outcomes are the worst-case scenario for self-control.

Let $x = (x_1, t_1, \dots, x_n, t_n)$ denote an income stream that gives money amount $x_j \geq 0$ at time point t_j, $j = 1, \dots, n$, and nothing otherwise. Implicit is $t_1 < \cdots < t_n$.

For $\varepsilon \in \mathbb{R}$, $x^{\uparrow\varepsilon}$ denotes the shift $(x_1, t_1 + \varepsilon, \ldots, x_n, t_n + \varepsilon)$ of x, where $t_1 + \varepsilon \geq 0$. We again analyze a general preference reversal:

$(x_1, s_1, \ldots, x_n, s_n) \succ_t (y_1, r_1, \ldots, y_m, r_m)$ 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, the immediacy effect must be the cause of the preference reversal and we assume, without loss of generality, that this favors x. Hence $t = s_1$, and either $s_1 < r_1$ and $x_1 > 0$, or $s_1 = r_1$ and $x_1 > y_1$. The preference with t' is maintained if we replace t' by 0 because then the immediacy effect for x is weakened:

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

We investigate the degree of time inconsistency by considering which shifts $x^{\uparrow\varepsilon}$ preserve the preference reversal:

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

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

THEOREM 4.1. Under Eq. 4.1, τ is the maximum length of the vulnerable periods. □

Theorem 4.1 shows that for preferences over n outcomes, vulnerable periods can be shorter than τ, but not longer. Intuitively, for timed outcomes the immediacy effect maximally affects the whole prospect, whereas for n-tuples it only affects the first outcome and, hence, is weaker. 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. Empirical illustration

To illustrate how replacing β with τ may affect our interpretation of time inconsistencies, we reanalyzed the data from Tanaka et al. (2010), on individual risk and time preferences in Vietnam. These preferences were related to demographic and economic variables, seeking links between economic success and preferences. 181 subjects answered 15 time preference questions by choosing between a money amount now and a larger amount in the future (3 days to 3 months), with real incentives. The average payment was about 6-9 days' wage. We repeated the authors' group-level analysis with non-linear least squares (see Tanaka et al. for details) for (β, δ) and extended it to (δ, τ). To ensure both impatience and decreasing impatience, as characteristic of the (β, δ) model, we truncated $0 \leq \beta \leq 1$, $0 < \delta < 1$, and $\tau \geq 0$.

Table 1 shows the estimation results, with conventional labeling of variables. The dummy variable "Trusted agent" is equal to 1 for subjects who stored the money earned during the experiments, and the variable "Risk payment" is equal to the amount of money the subject received during the elicitation of the risk preferences. People in South Vietnam generally have a higher income than their Northern Vietnamese counterparts. All variables (including the dummies) have been standardized.

The average β is .65, smaller than 1 ($p < 0.001$); β is unrelated to the demographic variables. The average daily discount factor δ is equal to .992, and is higher (more patience) for subjects with higher age, education, income, and money won in the risk part.

The third and fourth columns show the estimates for the (τ, δ) model. The p-values are usually lower for τ than for β, suggesting that τ delivers more statistical power. Whereas β is unrelated to any of the demographic variables, τ is associated with two variables that can serve as proxies for wealth: people in the (richer) South of Vietnam and people who received more money in the first part of the experiment have a higher τ (more decreasing impatience).

A comparison between the second and fourth columns shows that the standard errors for δ are higher in the (β, δ) framework and that we have more statistical power in the (τ, δ) framework. Again, p-values are generally lower in the (τ, δ) framework. This may be the result of collinearity between β and δ: both measure impatience (β in the short run and δ in the long run), and will thus be collinear when estimated jointly.

	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
# Observations	5,340		5,340	

Table 1. Regression results for the (β, δ) and the (τ, δ) models using the data from Tanaka et al. (2010). For each variable the table shows (from top to bottom) the coefficient, the standard error, and the p-value. Asterisks indicate significance at the 5% level.

To check for collinearity between β and δ, we computed parameter estimates for each individual on the basis of the 15 questions in the survey. Table 2 gives the correlation matrix of these individual estimates. We display Spearman rank order correlations to reduce the impact of outliers; an asterisk indicates significance. The parameters β and δ are strongly negatively correlated ($z = 4.23$, $p < 0.001$), suggesting that they tap into a common individual difference variable — impatience.

Toubia et al. (2013) found a similar negative correlation between β and δ. In contrast, the parameters δ and τ are not significantly correlated.

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

Table 2. Spearman correlations. The correlation between δ and τ is weaker than the correlation between β and δ (z=2.93, p = 0.003).

Of course, whether impatience and self-control are distinct psychological dimensions is an empirical rather than a modeling question. It is certainly possible that individuals who care little about the future as measured by their δ will also exhibit more time-inconsistent preferences, as measured by τ. In that case, the coefficient β, which merges impatience and time inconsistency into a single index of intertemporal misbehavior may be useful in empirical work, just as combining, say, verbal and mathematical ability into a single summary cognitive aptitude index may be useful in certain applications. The theoretical point we underline here is that such aggregation of two conceptually distinct dimensions into a single number should be done with eyes open. In contrast, the traditional approach of estimating β and δ and then interpreting β as self control may suggest a relationship between impatience and self-control when no such relationship exists.

6. Conclusion

We have examined time inconsistency for the popular quasi-hyperbolic discounting model by writing $f(t) = \delta^{t+\tau}$ rather than the standard $f(t) = \beta\delta^t$ (for $t > 0$, with $f(0) = 1$). This reformulation leads to a cleaner separation between impatience (measured by δ) and decreasing impatience (measured by τ), both empirically and theoretically. We provided a revealed preference basis for our new index of time inconsistency, τ, whereas for β none is known as yet to our best knowledge. This may be because β interacts with utility whereas τ reflects pure time preference. Our index has a natural interpretation in time units, as the perceived time penalty of any delay beyond the present. It is also the period of vulnerability to dynamic inconsistencies and, hence, to self-control problems.

Appendix

In the case of $\delta = \beta = 1$, the decision maker 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 measure of decreasing impatience

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

PROOF. All functions considered take value 0 at $t = 0$, and we describe only their values at $t \neq 0$. Substitution shows that

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

Given value 0 at 0, this transformation of $\ln\big(d(t)\big)$ 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 1, and gave no derivations. In particular, he did not handle the discontinuity at $t = 0$.

Proof of Theorem 3.1

By continuity and impatience a exists, and is between b in Eq. 3.2 and r. For all $t \leq a$, $\succcurlyeq_t$ 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.

Eq. 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 Eq. 3.2 (which is equivalent to $(l, r-\varepsilon) \prec_{r-\varepsilon} (h, r)$). For all $t \geq a+\tau$, $\succcurlyeq_t$ 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 Eq. 3.1, where the left-hand side follows from Eq. 3.3 and the right-hand side from Eq. A.1. Hence these s are vulnerable. □

Theorem 3.1 without Eq. (3.2)

We now consider the general case, without Eq. 3.2; see 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. $a \leq -\tau$. Then (h, r) is always preferred, even if l is received immediately at time point $t = 0$. This is the trivial case where Eq. 3.1 never arises.
CASE 2. $-\tau \leq a < 0$. 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 Eq. 3.2.
CASE 3. $0 \leq a \leq r - \tau$. 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 Eq. 3.2.

Proof of Theorem 4.1

We write $QH(y)$ for the quasi-hyperbolic discounted utility of y at time point s_1.

$(x_1, s_1, \ldots, x_n, s_n)^{\uparrow\varepsilon} \succ_{s_1+\varepsilon} (y_1, r_1, \ldots, y_m, r_m)$ 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 Eq. 4.2 for the shift ε. Hence Eq. 4.2 cannot hold for both a shift by ε and a shift by $\varepsilon + \tau$. For shifts that exceed $\varepsilon + \tau$, the above inequalities become stronger and favor x more. This implies that the set of vulnerable shifts cannot contain shifts that are further than τ apart. Numerical examples show that the length of the vulnerable period can indeed be less than τ for some x, y. We saw in §3 that for timed outcomes the vulnerable period has length τ and, hence, the maximum length is τ. □

References

Ainslie, George (1975) "Specious Reward: A Behavioral Theory of Impulsiveness and Impulse Control," *Psychological Bulletin* 82, 463–496.

Augenblick, Ned, Muriel Niederle, & Charles Sprenger (2015) Working over Time: Dynamic Inconsistency in Real Effort Tasks," *Quarterly Journal of Economics* 130, 1067–1115.

DellaVigna, Stefano (2009) "Psychology and Economics: Evidence from the Field," *Journal of Economic Literature* 47, 315–372.

Fudenberg, Drew & David K. Levine (2011) "Risk, Delay, and Convex Self-Control Costs," *American Economic Journal: Microeconomics* 3, 34–68.

Gruber, Jonathan & Botond Kőszegi (2004) "Tax Incidence when Individuals are Time-Inconsistent: The Case of Cigarette Excise Taxes," *Journal of Public Economics*, 88, 1959–1987.

Laibson, David I. (1997) "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics* 112, 443–477.

O'Donoghue, Ted & Matthew Rabin (1999) "Doing It now or later," *American Economic Review* 89, 103–124.

Phelps, Edmund S. & Robert A. Pollak (1968) "On Second-Best National Saving and Game-Equilibrium Growth," *Review of Economic Studies* 35, 185–199.

Pollak, Robert A. (1968) "Consistent Planning," *Review of Economic Studies* 35, 201–208.

Prelec, Drazen (2004) "Decreasing Impatience: A Criterion for Non-Stationary Time Preference and "Hyperbolic" Discounting," *Scandinavian Journal of Economics* 106, 511–532.

Strotz, Robert H. (1955) "Myopia and Inconsistency in Dynamic Utility Maximization," *Review of Economic Studies* 23 (Issue 3, June 1956), 165–180.

Tanaka, Tomomi, Colin F. Camerer, & Quang Nguyen (2010) "Risk and Time Preferences: Linking Experimental and Household Survey Data from Vietnam," *American Economic Review* 100, 557–571.

Toubia, Olivier, Eric Johnson, Theodoros Evgeniou, & Philippe Delquié (2013), "Dynamic Experiments for Estimating Preferences: An Adaptive Method of Eliciting Time and Risk Parameters," *Management Science* 59, 613–640.