Quasi-exponential discounting*

Stephen L. Cheung†
Kieran MacGibbon
Arquette Milin-Byrne
Agnieszka Tymula

The University of Sydney

30 July 2024

Abstract

Alternatives to the standard model of time preference typically relax the assumption of an exponential discount function while retaining the framework of discounted utility. We report novel behavioural data inconsistent with this approach. Illustrating this, we estimate highly significant "present bias", despite our data exhibiting stationarity. The paradox is resolved by relaxing discounted utility itself to allow discounting to be context dependent. We propose quasi-exponential discounting (QED), a fixed penalty applied to all episodes of delay, as a particularly simple model of this type and show that it provides an excellent approximation to the best fit to our data.

Keywords: time preference, present bias, stationarity, relative discounting.

JEL classification: C91, D15, D90.

---

*We thank Paul Glimcher, Hyundam Je and Xueting Wang for discussions on related projects, and participants at the 2024 conference on Foundations of Utility and Risk for their feedback. We acknowledge support from the Australian Research Council Centre of Excellence for Children and Families over the Life Course (CE200100025). This study was approved by The University of Sydney Human Research Ethics Committee, under protocol number 2023/491 (CI Cheung). All authors declare that they have no no relevant or material financial interests that relate to the research described in this paper.

†Corresponding author. School of Economics, Social Sciences Building A02, The University of Sydney NSW 2006, Australia (Stephen.Cheung@sydney.edu.au).

1 Introduction

Choices involving trade-offs between alternatives at different points in time are central to many branches of economics. In evaluating such trade-offs, the standard model of *exponentially discounted utility* (Samuelson, 1937) posits that a chooser will discount each alternative at a constant exponential rate according to the length of time that passes between the moment of decision and the receipt of that alternative, and make the choice that has the greatest discounted utility. In particular, since this model is time consistent (Strotz, 1955), the preference between delayed alternatives will depend on the interval of time that separates them, but not the time separating the alternatives from the moment of decision.

While behavioural economists and psychologists largely agree that exponential discounting is descriptively inadequate (Frederick et al., 2002), they do not agree on a preferred alternative model. In economics, the dominant behavioural model is quasi-hyperbolic discounting (Laibson, 1997; O'Donoghue and Rabin, 1999), which draws a sharp distinction between choices that involve the prospect of an immediate reward and ones that do not: it is only in the presence of an immediately-available reward that preferences may reverse. On the other hand, psychologists favour models of hyperbolic (Mazur, 1984) or generalised hyperbolic discounting (Loewenstein and Prelec, 1992; Myerson and Green, 1995), which allow for continuously declining discount rates.

Reflecting these differences in the modelling of discounting, behavioural economists and psychologists also differ in the protocols they use to measure time preferences. In economics, identification of the "present bias" parameter β of the quasi-hyperbolic model depends crucially on varying the absence or presence of a "front-end delay" separating the moment of decision from the sooner reward. But since the model implies that all delayed rewards are discounted at a constant rate, variation in the "back-end delay" between sooner and later rewards is not essential to estimate the long-run discounting parameter δ. By contrast in psychology, identifying the shape of a hyperbolic discount function requires more extensive variation in back-end delays, but does not rely on varying front-end delay since the model does not involve any discontinuity at the present. As a result of these distinct approaches, many datasets can only be used to estimate one of these two popular classes of models, and there is a surprising paucity of studies that allow for a comparison between economists' and psychologists' preferred models.

Although the quasi-hyperbolic and (generalised) hyperbolic models relax the assumption of an exponential discount function, they retain the framework of discounted utility. According to this framework, future rewards are weighted by a discount function $D(t) \leqslant 1$ which measures the weight attached to utility t periods in the future relative to the present which has weight $D(0) = 1$. When choosing at time 0 between rewards at times t and $t+k$, both are discounted back to time 0 and the discount from t to $t+k$ is found by *comparing the discounting that occurs from* 0 *to* $t+k$ *to that which occurs from* 0 *to* t. When $D(t)$ is exponential this detail is immaterial because the exponential function satisfies stationarity (Fishburn and Rubinstein, 1982): the discount from t to $t+k$ is independent of the front-end delay preceding t. However, the exponential model is the only stationary model of discounted utility. As an alternative to generalising the discount function, Scholten and Read (2006) and Ok and Masatlioglu (2007) generalise discounted utility itself, by measuring discounting relative to time t instead of time 0. It will be seen that this approach will be necessary to adequately describe the data we report in this paper.

In this paper, we report the results of a study that we originally designed to allow for a comparison between the quasi-hyperbolic and (generalised) hyperbolic models of discounted utility. Over 300 participants, recruited from a general UK sample, took part in an incentivised study in which we asked them to state the amount on a specified sooner date (the "sooner equivalent") that they would consider equally desirable as a fixed amount on a specified later date. Our design both varies the front-end delay (three levels: zero, one, and seven days) to identify the present bias component of a quasi-hyperbolic model and provide for a test of stationarity, and also includes substantial variation in back-end delays (five levels: seven, 14, 30, 90, and 180 days) to identify the shape of long-run discounting.

Our results are unexpected. First, we do not see any effect of front-end delay: at each back-end delay, participants report sooner equivalents that are indistinguishable across the three levels of front-end delay. Thus not only is there no present bias, but more generally the data satisfy stationarity. Under discounted utility, only the exponential discount function has these features. However, we find that an exponential function describes the data poorly: it is too shallow to capture the steep discounting of short intervals, and too steep at longer intervals. In fact, it performs so poorly that a quasi-hyperbolic function – which allows for steeper discounting specifically where a sooner reward is available immediately – both improves the model fit (as measured by the Bayesian information criterion, BIC) and detects significant "present bias" ($\beta = 0.833$, less than 1 with $p < 0.0001$), despite there clearly not being present bias in the data.

The explanation for this spurious finding of "present bias" is simple: the β parameter of quasi-hyperbolic discounting effectively "decouples" the intercept of the discount function, better matching the shape of long-run discounting, but by definition it can only do so for choices with no front-end delay. Since behaviour does not differ across the three levels of front-end delay, the performance of the model would be further improved by allowing this "decoupling" to occur even where there is a front-end delay. This amounts to a model in which, in addition to long-run discounting captured by δ, *all choices are subject to an added penalty β to wait for a later reward, not only ones that do not involve a front-end delay*. We confirm this conjecture, and dub the resulting model "quasi-exponential discounting" (QED).

Clearly, our quasi-exponential proposal falls outside the framework of discounted utility: there is no discount function, measured from time 0, that can generate this prediction. However it falls within a more general framework of "relative discounting" (Scholten and Read, 2006; Ok and Masatlioglu, 2007) in which discounting is context dependent. In particular, let the discount fraction $\Delta(t,k)$ measure the amount of discounting that occurs over an interval of length k, starting at time t. Whereas under discounted utility this discount fraction is derived from the discount function as $\Delta(t,k) = D(t+k)/D(t)$, in relative discounting it is treated as a primitive and there need not be any discount function that generates it. Models of this form have been proposed to account for subadditive discounting, the finding that an interval is discounted more heavily when it is divided into subintervals than when it is left undivided (Read, 2001; Scholten and Read, 2006). Specific examples of relative discounting models, suggested previously in the literature, include the subadditive model proposed by Read (2001, equation 16), and the "as soon as possible" (ASAP) model proposed by Kable and Glimcher (2010).

In principle, relative discounting opens considerable degrees of freedom in the modelling of $\Delta(t,k)$. However, our finding of stationarity imposes a disciplining constraint: $\Delta(t,k)$ should depend only on the back-end delay length k, and not on t. Both the subadditive and ASAP models have this feature, as does our quasi-exponential model. We suggest that a natural way to generate such models is to take existing models of the discount function, defined over time measured from 0, and recast them as discount fractions defined over the interval length k. Seen in this light, ASAP represents a recasting of the Mazur (1984) simple hyperbolic discount function, while our quasi-exponential model is of course a recasting of the quasi-hyperbolic function. Recasting the generalised hyperbolic function of Loewenstein and Prelec (1992) in the same manner yields a model we refer to as "generalised ASAP".

Overall, we find that the joint best-fitting models for our data, as measured by BIC, are the subadditive model of Read (2001) and the generalised ASAP model. However, our quasi-exponential model provides an excellent approximation to the best models, while enjoying the analytical simplicity and ease of interpretation that it shares with its quasi-hyperbolic cousin. Indeed, just as quasi-hyperbolic discounting represents a *minimal relaxation of the exponential discount function* that can account for diminishing impatience, quasi-exponential discounting is a *minimal relaxation of discounted utility* that can account for subadditivity. More generally, our findings illustrate how the study of discounting behaviour can benefit from looking beyond the shape of the discount function to consider parsimonious extensions to the discounted utility framework itself.

The paper proceeds as follows: Section 2 elaborates on the conceptual framework to our study, Section 3 sets out our research methodology, Section 4 reports results, and Section 5 provides a discussion.

2 Framework

In this Section we set out the conceptual framework underpinning our paper, and clarify the distinction between discounted utility and relative discounting models. We also define the functional forms for each of the models of discounting that we estimate in our data.

In our study, participants are asked at time 0 to specify an amount x which, if received at time $t \geqslant 0$, they would consider equally desirable as receiving y at time $t+k > t$. The value of y is fixed, and the front-end delay t and back-end delay k are varied across trials. The "sooner equivalent" x is constrained to lie within $[0, y]$, and a chooser who is more impatient will report a lower value of x.

Assuming discounted utility, the condition for indifference can be written as:

$$D(t)\,v(x) = D(t+k)\,v(y) \tag{1}$$

where $D(t)$ is the discount function, which measures the weight assigned to utility at time t relative to utility at time 0 (with $D(0) = 1$), and $v(x)$ is the instantaneous utility function.

More generally, following Scholten and Read (2006, equation 6) and Ok and Masatlioglu (2007, equation 1), we can write the indifference equation as:

$$v(x) = \Delta(t,k)\,v(y) \tag{2}$$

where $\Delta(t,k)$ is the discount fraction (or relative discount factor), which measures the amount of discounting that occurs over an interval of length k, starting at time t.

Ok and Masatlioglu (2007) derive a representation theorem for time preferences of the type described by equation 2. They show that the only preferences in this class that obey transitivity are ones that take the form of discounted utility as in equation 1; we define parametric exemplars of such models in Section 2.1. Ok and Masatlioglu (2007) further show that the only relative discounting models that obey stationarity are ones in which the discount fraction is a function of k alone; we define exemplars of such models in Section 2.2. Finally, the only model that is both stationary and transitive is the standard model of exponentially discounted utility.

2.1 Discounted utility models

Under discounted utility, $\Delta(t,k)$ is derived from the discount function as $\Delta(t,k) = D(t+k)/D(t)$. Thus all discounting is measured from time 0, and the discounting that occurs between t and $t+k$ is found by comparing $D(t+k)$, the discounting that occurs between 0 and $t+k$, to $D(t)$, the discounting which occurs between 0 and t.

We next define discount functions for four popular discounted utility models that we consider, and their corresponding discount fractions. First, the standard exponential discount function (Samuelson, 1937) assumes a constant rate of discounting in every period:

$$D(t) = \delta^t \quad \Rightarrow \quad \Delta(t,k) = \delta^k \tag{3}$$

Second, the quasi-hyperbolic discount function (Laibson, 1997; O'Donoghue and Rabin, 1999) extends the exponential model to capture dynamic inconsistency by applying an additional fixed discount $\beta \leqslant 1$, interpreted as present bias, to all rewards that are not available immediately:

$$D(t) = \begin{cases} 1 & t = 0 \\ \beta\delta^t & t > 0 \end{cases} \quad \Rightarrow \quad \Delta(t,k) = \begin{cases} \beta\delta^k & t = 0 \\ \delta^k & t > 0 \end{cases} \tag{4}$$

Next, the simple hyperbolic discount function (Mazur, 1984) is used in psychology to account for diminishing impatience (the finding that the rate of discounting decreases as the interval length increases):

$$D(t) = \frac{1}{1+\alpha t} \quad \Rightarrow \quad \Delta(t,k) = \frac{1+\alpha t}{1+\alpha(t+k)} \tag{5}$$

Finally, the generalised hyperbolic discount function (Loewenstein and Prelec, 1992) extends the simple hyperbolic model by introducing an additional parameter to allow the discount function to decline more or less rapidly than predicted by the simple hyperbolic model:[1]

[1] An equivalent specification is proposed by Myerson and Green (1995, equation 4). In this paper, we estimate the parametrisation proposed by Loewenstein and Prelec (1992).

$$D(t) = \frac{1}{(1+\alpha t)^{\gamma/\alpha}} \quad \Rightarrow \quad \Delta(t,k) = \left(\frac{1+\alpha t}{1+\alpha(t+k)}\right)^{\gamma/\alpha} \tag{6}$$

From inspection of equations 3–6, it is clear that the exponential discount function is the only one for which $\Delta(t,k)$ does not depend upon the front-end delay t. Indeed, within the framework of discounted utility, the exponential model is the only one to have this property of stationarity.

2.2 Models of stationary relative discounting

Relative discounting generalises discounted utility because the discount fraction $\Delta(t,k)$ need not be derived from a discount function, and may instead be treated as a primitive. We consider such models because it turns out that our data satisfy stationarity, and we are interested in stationary models that are not necessarily exponential. We thus restrict attention to models of stationary relative discounting in which the discount fraction is a function of k but not t, in other words where the discounting that occurs over an interval of length k does not depend on the timing of the onset of that interval, t.

We next define discount fractions for the four models of stationary relative discounting that we consider. Each of these models is closely related to one of the discounted utility models defined in Section 2.2 above. First, the subadditive model proposed by Read (2001, equation 16) extends the discount fraction for the exponential model in equation 3 by replacing k with a power transformation:

$$\Delta(t,k) = \delta^{k^{\vartheta}} \tag{7}$$

This model was proposed to accommodate the finding of subadditivity in discounting. One interpretation of the model is that the parameter ϑ captures non-linear time perception, such that discounting is exponential with respect to the subjectively-perceived interval length, k^{ϑ}.

The remaining models are obtained by taking the *discount functions* in equations 4–6, originally defined over time measured from 0, and recasting them as *discount fractions* defined directly over interval lengths k. This amounts to a "shifting of the origin" such that instead of being measured from time 0, discounting is measured from the start of the interval being discounted at time t.

Recasting the quasi-hyperbolic discount function from equation 4 in this manner produces the model that we refer to as "quasi-exponential discounting" (QED):

$$\Delta(t,k) = \begin{cases} 1 & k = 0 \\ \beta\delta^k & k > 0 \end{cases} \tag{8}$$

This has the interpretation that, in addition to long-run exponential discounting through δ, there is *an additional fixed penalty for waiting* β *that is applied to all episodes of delay* (and not only ones that start in the present, as in quasi-hyperbolic discounting).

Next, recasting the simple hyperbolic discount function in equation 5 yields the "as soon as possible" (ASAP) model proposed by Kable and Glimcher (2010, equation 3):

$$\Delta(t,k) = \frac{1}{1+\alpha k} \tag{9}$$

Finally, recasting the generalised hyperbolic discount function in equation 6 yields a model that we refer to as "generalised ASAP":

$$\Delta(t,k) = \frac{1}{(1+\alpha k)^{\gamma/\alpha}} \tag{10}$$

This is also a simplified (two-parameter) version of the "discounting by intervals" model of Scholten and Read (2006), which has a total of four parameters.

3 Methodology

Our measurement of time preference consisted of 15 matching tasks. In each task, participants were asked to state the amount which, if received on a specified sooner date, they would consider equally desirable as receiving GBP £30 on a specified later date. Sooner rewards were offered with front-end delays (t) of zero, one, or seven days from the date of the study. This variation in sooner reward dates facilitates the identification of present bias and violations of stationarity. Later rewards were offered with back-end delays (k) of seven, 14, 30, 90, or 180 days after the sooner date. Variation in back-end delay lengths facilitates identification of the shape of long-run discounting. We crossed every combination of the three front- and five back-end delay lengths to create a total of 15 items. The order in which these items were presented was randomised at an individual level.

Figure 1 illustrates the presentation of a sample item. A stylised calendar was used to visualise the full 187-day timespan of the study, with the sooner and later reward dates for the current trial highlighted as coloured boxes. In each trial, participants reported their sooner equivalent of the later £30 by placing a marker on a slider. In its initial state the slider appeared without any marker being visible, such that participants had to first click to make the marker appear, and then drag it to the desired position. The sooner equivalent was displayed above the marker, and the value updated dynamically in response to movements in its position. Participants could adjust the position of the marker without time limit, and once satisfied with their choice submitted it by clicking "Next". In the Instructions (see Appendix A), the concept of a sooner equivalent was explained as an amount such that a participant would find it "very hard to choose" between receiving the amount they specify sooner, or £30 later. In the task interface, this was summarised using the prompt "To me, £____ **in** t **days** is as good as £30 **in** $t+k$ **days**".

To induce truthful reporting, if a participant was rewarded for one of these trials, the amount was determined using the Becker et al. (1964) mechanism. A random value z was drawn from a uniform distribution between £0.10 and £30 (in increments of £0.10), representing the offer of a sooner amount, and this

Figure 1: Interface for a sample time preference trial.

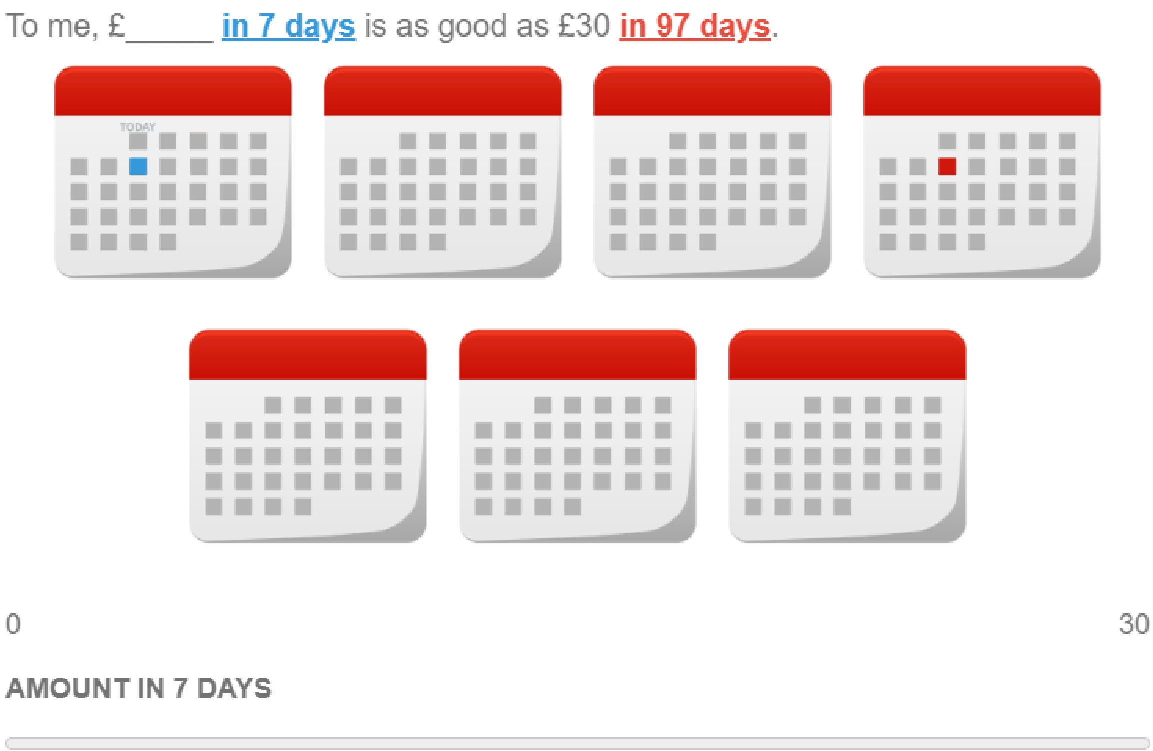

was compared to the participant's response in the chosen trial. If z was greater than (or equal to) the participant's sooner equivalent, the participant received z on the sooner date; otherwise they received £30 on the later date. Participants were instructed that it was in their best interest to report their preferences truthfully, and this was illustrated by means of an example showing how misreporting one's sooner equivalent could result in receiving a less desired reward.[2]

A total of 302 UK-resident participants (55.6% males; mean age 44.0 years, $SD = 13.8$) completed the study, which was programmed in Qualtrics, online through Prolific (prolific.com) on the morning of 10 November 2023, GMT. To be eligible for the study, participants were required to be at least 18 years of age, reside in the UK, and have an approval rating on Prolific of at least 95%. After giving informed consent, participants were required to pass two attention checks before continuing. In addition to the incentivised time preference and exponential growth items, participants completed an unincentivised subjective time perception task (Bradford et al., 2019) which will be described in Section 4.4, the Cognitive Reflection Test (Frederick, 2005), the Barratt Impulsiveness Scale - Brief (Steinberg et al., 2013), the "Big five" financial literacy items (Lusardi, 2011), and exit questions covering clarity of the instructions and trust in payments. On a five-point Likert scale, 80% of participants rated the clarity of the

[2]In addition to the 15 time preference items, participants completed a block of 15 items measuring exponential growth bias (Stango and Zinman, 2009) for a separate project, and the order of presentation of the two blocks was randomised at an individual level. In exponential growth bias trials, participants were asked to estimate what amount they would need to deposit, at a specified compound interest rate, to achieve a specified savings goal after some number of years. As well as receiving a flat fee for completing the study, each participant had a one-in-ten chance to receive a bonus determined by their response to one randomly-selected item, which could either be a time preference or exponential growth trial. Exponential growth items were incentivised using a quadratic scoring rule, following Levy and Tasoff (2017).

instructions as four or above, while 79% reported their trust in the researchers to pay as described in the instructions as four or above. The median time taken to complete the study as a whole was 29 minutes.

Each participant received base compensation of £4 in addition to a one-in-ten chance to receive a bonus of up to £30 based on their choice in one randomly-selected trial. Sooner rewards for time preference trials with no front-end delay were sent within 30 minutes of the final participant completing the study, and within four hours of the first participant commencing.

4 Results

We set out our results as follows. Section 4.1 documents the absence of present bias in the aggregate. While participants report lower sooner equivalents when later rewards have longer back-end delays, there is no discernable response to the front-end delay. This indicates that a model that explains the data should have the property of stationarity. Section 4.2 reports representative-agent estimates of discounted utility. The standard exponential model is the only discounted utility model to exhibit stationarity, yet it performs poorly as it struggles to explain choices at both short and long back-end delays. A quasi-hyperbolic model offers added flexibility, but only for choices without front-end delay. We find that it improves on the exponential model (as measured by BIC) and even detects highly-significant "present bias" ($\beta = 0.833$, $p < 0.001$), despite there being no present bias in the data. From the class of discounted utility models, the generalised hyperbolic function performs best, although it predicts greater patience with increasing front-end delays, a pattern we do not see in our data.

Section 4.3 reports estimates of stationary relative discounting. The subadditive model of Read (2001) and generalised ASAP (equation 10) are the joint best models, however our quasi-exponential model provides an excellent approximation to these models. Section 4.4 introduces the subjective time perception data. We find that while time perception is indeed nonlinear, it cannot fully account for deviations from exponential discounting. Finally, Section 4.5 examines which models explain choices best at an individual level, showing that the vast majority of individual participants are stationary but not exponential (and thus violate discounted utility).

4.1 Absence of present bias

Figure 2 depicts participants' mean sooner equivalents of a later £30 in each of the 15 trials, with lower sooner equivalents indicating more impatient choices. The shading of bars is used to denote different levels of front-end delay, while the bars are grouped by back-end delay and the error bars represent ±1 standard error of the mean. The figure shows that while participants respond as expected to back-end delay – by assigning lower sooner equivalents to more delayed rewards – there is little aggregate response to the front-end delay. If there is present bias in this data, the leftmost and darkest bar in each group should be lower than the other two. Instead there is no such systematic pattern, and all differences between the three levels of front-end delay are minor, as can be seen from the overlapping error bars.

Figure 2: Mean sooner equivalents, by front- and back-end delay length.

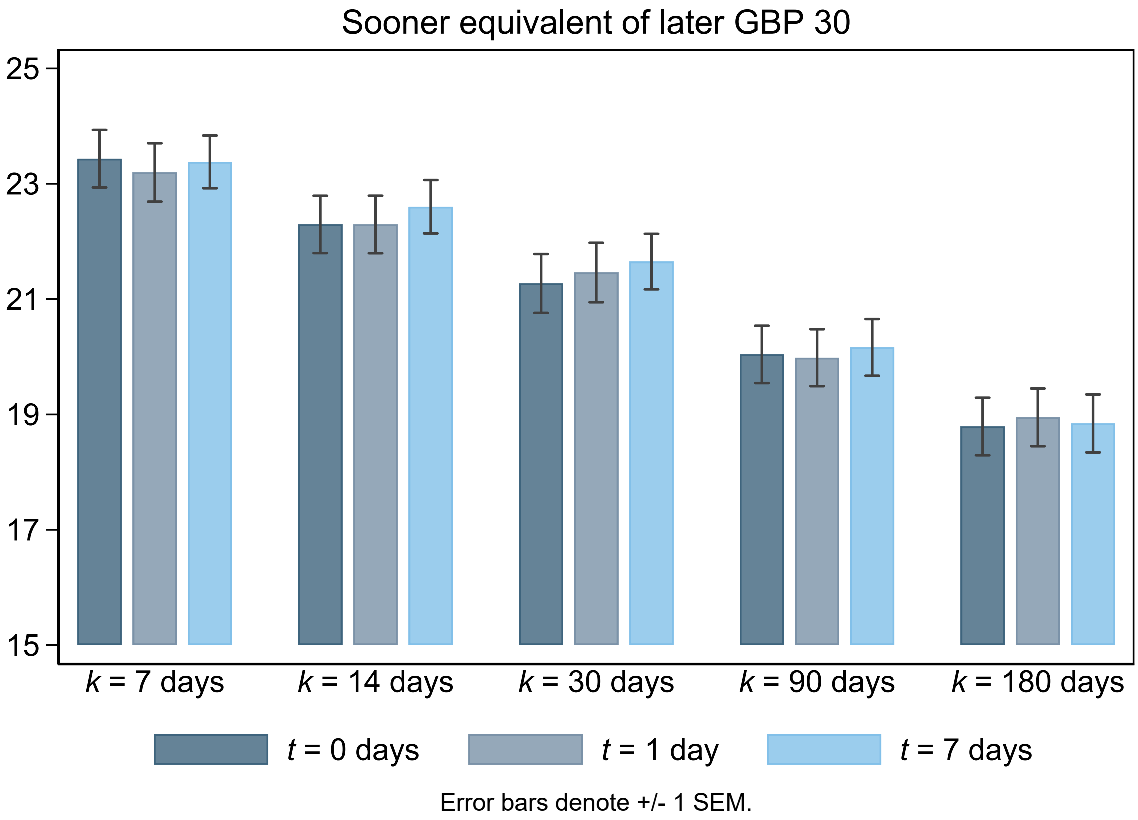

To examine effects of front- and back-end delay on behaviour without committing to a specific model of discounting, we conduct a two-way repeated-measures analysis of variance to examine the effects of t (zero, one, and seven days) and k (seven, 14, 30, 90, and 180 days) as within-subjects factors on participants' sooner equivalents, with degrees of freedom corrected using Box's conservative epsilon. This reveals a significant main effect of back-end delay ($F = 176.41$, $p < 0.0001$), but no significant main effect of front-end delay ($F = 0.75$, $p = 0.3877$). Moreover, there is also no significant interaction between t and k ($F = 0.26$, $p = 0.6086$), confirming that the effect of back-end delay on participants' sooner equivalents is similar at each level of the front-end delay. These findings indicate that the data are best described by a model in which discounting is sensitive to the length of the interval over which a reward is deferred (k), but not whether a sooner reward is available immediately or later (t).

4.2 Estimates of discounted utility

In this section we report representative agent estimates of the four models of discounted utility defined in Section 2.1, pooling the data of all 302 participants. We assume a linear instantaneous utility function $v(x)$,[3] and estimate the parameter(s) of each model by nonlinear least squares, with standard errors

[3]In the context of intertemporal choice, the curvature of instantaneous utility captures the preference of a decision-maker to smooth reward streams over time. Therefore, to relax the assumption of linear utility requires a more complex design in which participants make choices over *bundles of rewards* (delivered on both the sooner and later dates) that vary in the opportunities

clustered at the level of individual participants. Table 1 reports point estimates and robust 95% confidence intervals of the parameters for each of the discounted utility models defined in equations 3 to 6. Figure 3 plots the resulting predicted values for the discounted value of a delayed £30 as a function of back-end delay on the horizontal axis, with each panel corresponding to a different level of the front-end delay. The mean sooner equivalents and their standard errors are denoted by black markers.

We begin with the standard model of exponentially discounted utility. Since this model satisfies stationarity, its predicted values are identical across all three panels of Figure 3. The estimated daily discount factor of $\delta = 0.9962$ corresponds to a daily discount rate of 0.38%, which in turn compounds to an annualised discount rate of 302%. However, it can be seen in Figure 3 that the exponential function greatly overstates the discounting of the longest back-end delay of 180 days: whereas the mean sooner equivalent at 180 days is around £19, the prediction of the exponential model is closer to £15. The exponential function is thus both too flat to account for the initial sharp decline in sooner equivalents over short intervals, and too steep to capture the flattening out that occurs at longer intervals. This highlights the central theme of this section: there is no discounted utility model that simultaneously captures the two most salient features of our data, namely stationarity and the shape of long-run discounting.

The model of quasi-hyperbolic discounted utility was proposed to capture the phenomenon of "present bias": steeper discounting specifically in the presence of a potentially immediate reward. The model thus allows for more pronounced discounting in the left panel of Figure 3 ($t = 0$), but makes identical predictions for both the centre ($t = 1$) and right ($t = 7$) panels. Of course, our data do not exhibit any present bias, and we see very similar patterns of behaviour across all three panels. It is thus striking, and quite unexpected, that we estimate highly significant "present bias", with $\beta = 0.833$ and a 95% confidence interval of $(0.809, 0.857)$. Moreover, the BIC value of the quasi-hyperbolic model improves substantially upon that of the exponential model, as seen in the bottom row of Table 1.

The explanation for this remarkable finding can be seen from closer examination of Figure 3: the quasi-hyperbolic present bias parameter β allows for a "decoupling" of the intercept of the fitted curve, better describing choices at short back-end delays – but only in the left panel where there is no front-end delay. At the same time, the estimate of $\beta < 1$ is accompanied by a larger estimate of δ (less long-run discounting as compared to the exponential model), which better describes choices at longer back-end delays in the centre and right panels where β is not activated. In short, whereas the parameter β was intended to model the response to front-end delay (of which there is none), it actually functions to describe the shape of long-run discounting. Given that behaviour is in fact similar across all three panels, in Section 4.3 we estimate the quasi-exponential model in which β is allowed to act upon all choices, both with and without front-end delay. We show that this further improves upon the quasi-hyperbolic model, and in fact closely approximates our best-fitting models.

Finally, the models of simple hyperbolic and generalised hyperbolic discounted utility were proposed to capture patterns of diminishing impatience (not limited to present bias) that are incompatible with

they offer for smoothing. The consensus from studies using such designs is that when measured in this way, the instantaneous utility function is very close to linear, and moreover its curvature is quite distinct from that of a Bernoulli utility function measured in the domain of risk (Andreoni and Sprenger, 2012; Abdellaoui et al., 2013; Cheung, 2020). In particular, Cheung (2020) finds the effect upon estimated annual discount rates of adjusting for the curvature of instantaneous utility, as opposed to simply imposing linear utility, to be little more than one percentage point, whereas the alternative of measuring utility in the domain of risk results in a sizeable overcorrection.

Table 1: Estimates of discounted utility models

	(3)	(4)	(5)	(6)
	Exponential	Quasi-hyperbolic	Simple hyperbolic	Generalised hyperbolic
δ	0.9962	0.9968		
	(0.9958, 0.9966)	(0.9964, 0.9971)		
β		0.8330		
		(0.8093, 0.8567)		
α			0.0058	0.4523
			(0.0049, 0.0067)	(0.3194, 0.5852)
γ				0.0631
				(0.0451, 0.0812)
R^2	0.8168	0.8267	0.8249	0.8515
BIC	2,749	2,504	2,542	1,805

Notes: All models are estimated by nonlinear least squares, using 4,530 observations from 302 participants. Column numbers correspond to equations in Section 2.1. Robust 95% confidence intervals, with standard errors clustered at the participant level, are in parentheses.

Figure 3: Fitted models of discounted utility

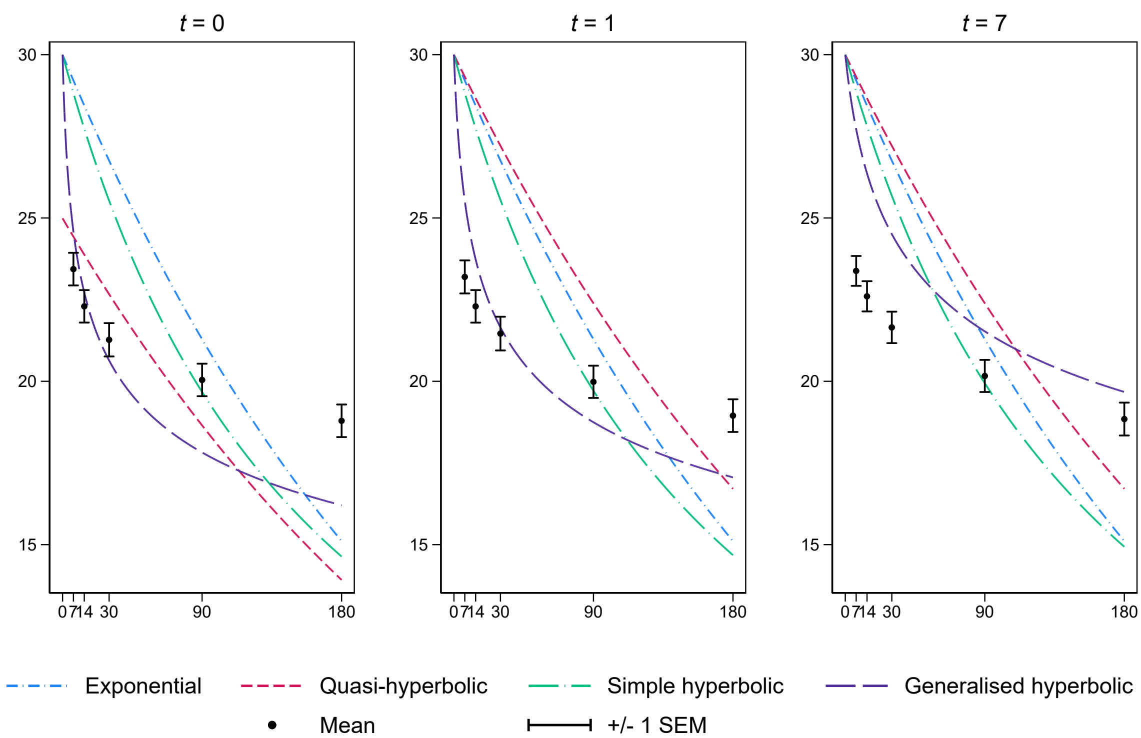

exponential discounting. These models make distinct predictions in each of the three panels of Figure 3 (although for the simple hyperbolic function the differences are slight). It can be seen that (as with the exponential function), the shape of the simple hyperbolic function is too inflexible to describe choices at both short and long back-end delays. The generalised hyperbolic function, which features an additional curvature parameter, is found to be the best-fitting discounted utility model by a considerable margin. However it fails to capture the other key feature of our data, namely stationarity. As can be seen by comparing the panels of Figure 3, the model predicts more patient choices with increasing front-end delay, a pattern not evident in our data.

4.3 Estimates of stationary relative discounting

Table 2 reports point estimates and robust 95% confidence intervals of the parameters (from standard errors clustered at the level of individual participants) for each of the models of stationary relative discounting defined in equations 7 to 10 of Section 2.2. Figure 4 plots the resulting predicted values for the discounted value of a delayed £30, using the same format as for the discounted utility models in Figure 3. Since we restrict attention to relative discounting models that have the property of stationarity, each of these models predicts an unchanged pattern of behaviour across all three panels of Figure 4. These models thus allow discounting to be stationary without having the very particular shape of an exponential discount function; however to do so it is necessary to relax the assumption of discounted utility.

It can be seen in Figure 4 that the single-parameter ASAP model of Kable and Glimcher (2010) performs poorly because, as was the case for the exponential discount function, its shape is too inflexible to capture the observed extent of diminishing impatience with increasing interval length. Moreover, comparing its BIC in Table 2 to that of the corresponding model of discounted utility in Table 1 reveals no discernable improvement. This is because the rather flat shape of the function at small values of k, combined with our short front-end delays of one and seven days, implies that the effect of "shifting the origin" is quite limited. By contrast, Kable and Glimcher (2010) first proposed ASAP in the context of a dataset in which the front-end delay was 60 days, such that the effect of this shift is more pronounced.

The remaining models have two parameters, which suffice to capture the observed patterns of behaviour. Comparing Tables 2 and 1, the BIC values of these models improve considerably upon any discounted utility model. In particular, the subadditive model proposed by Read (2001) and generalised ASAP are the joint best models: their predictions are essentially indistinguishable, and track the aggregate behaviour closely. Moreover, generalised ASAP is the relative discounting counterpart to the best-fitting discounted utility model. Comparing it to generalised hyperbolic discounted utility thus isolates the effect of relaxing the assumption of discounted utility while holding the functional form constant.

Finally, the quasi-exponential model retains an underlying exponential shape which makes it less flexible than either the subadditive or generalised ASAP models. To account for the steep discounting of shorter intervals, this model instead imposes a fixed multiplicative penalty β for delay. To account for stationarity, it applies this penalty to all instances of delay and not only ones that begin in the present – this is what distinguishes it from quasi-hyperbolic discounted utility. The predictions of this model do not track the data precisely, yet they clearly capture its essential qualitative features. As measured by the BIC, the

Table 2: Estimates of stationary relative discounting models

	(7)	(8)	(9)	(10)
	Subadditive	Quasi-exponential	ASAP	Generalised ASAP
δ	0.8371	0.9988		
	(0.8058, 0.8684)	(0.9986, 0.9991)		
ϑ	0.1840			
	(0.1490, 0.2190)			
β		0.7598		
		(0.7290, 0.7906)		
α			0.0057	6.7079
			(0.0049, 0.0066)	(−0.3477, 13.7635)
γ				0.4297
				(0.0368, 0.8226)
R^2	0.8603	0.8599	0.8249	0.8603
BIC	1,528	1,542	2,542	1,529

Notes: All models are estimated by nonlinear least squares, using 4,530 observations from 302 participants. Column numbers correspond to equations in Section 2.2. Robust 95% confidence intervals, with standard errors clustered at the participant level, are in parentheses.

Figure 4: Fitted models of stationary relative discounting

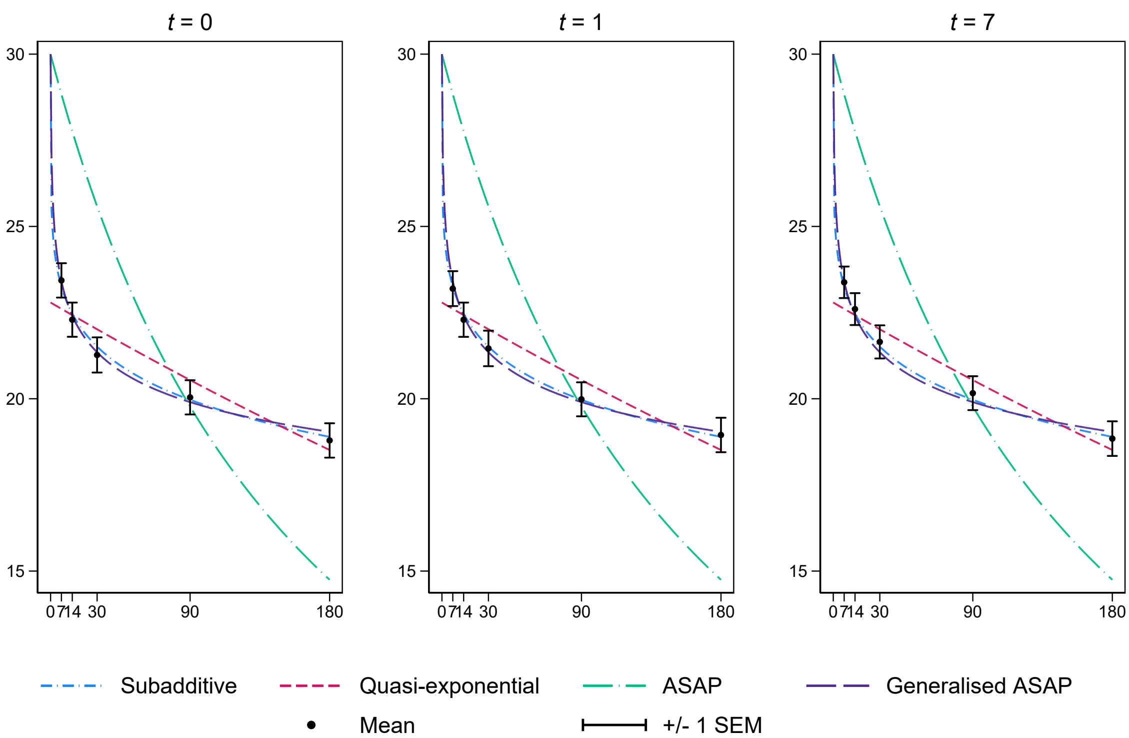

cost of the quasi-exponential approximation relative to the best-fitting models is minor when compared to the cost of incorrectly maintaining the assumption of discounted utility.

4.4 Subjective time perception

As noted in Section 2.2, one interpretation of the power parameter ϑ of subadditive discounting in equation 7 is in terms of nonlinear time perception: if people discount exponentially with respect to subjective delay k^{ϑ}, they would behave as described by a subadditive model. It has similarly been argued that if time perception follows a logarithmic relation then exponential discounting would generate behaviour as described by the generalised hyperbolic function in equation 6 (Takahashi, 2005; Bradford et al., 2019).[4] If these conjectures are correct, accounting for subjective time perception would negate evidence of non-exponential discounting, and re-establish exponentially discounted utility as the preferred model (Zauberman et al., 2009; Bradford et al., 2019). In this section we introduce our measure of time perception to examine whether this is indeed the case.

Our unincentivised self-report measure of time perception was based on that of Bradford et al. (2019). In each trial of this task, participants were asked to place a marker on a slider to indicate how near or far in the future they perceived different dates to be. The dates were presented in random order, with each shown on a separate screen, and participants knew in advance that they would vary from seven days to 25 years in the future.[5] The left-most end of the scale was labelled “Very near” and the right-most end “Very far”. There were no other values identified on the scale, and initially no marker was visible. Responses were coded on a linear scale from one to 100, however these values were not visible to participants.

Following Bradford et al. (2019), we account for individual differences in the propensity to make use of the full width of the scale by assuming that each participant has an objective perception of the shortest presented interval (in our case, seven days). For intervals longer than seven days, we divide each participant’s reported perception of that interval by their own reported perception of the seven-day interval, and multiply this normalised value by seven to express it in units of subjectively-perceived days. Thus, letting $s(k)$ denote the value between one and 100 coded in response to an interval of length k, we define the subjectively-perceived length of that interval in days, $\kappa(k)$, as:

$$\kappa(k) = \frac{7s(k)}{s(7)} \tag{11}$$

Figure 5 reports the means of $\kappa(k)$ and their standard errors for interval lengths up to 180 days. The dashed diagonal corresponds to an objective perception of time, while the curves represent fitted values of the Weber-Fechner (logarithmic) and Stevens’ power laws, in each case normalised such that the shortest interval of seven days is objectively perceived. It can be seen that time perception is nonlinear and

[4] In psychophysics, a logarithmic relation between a stimulus and its perception is known as the Weber-Fechner law, while a power relation is known as Stevens’ power law.

[5] 25 years was the longest investment horizon considered in the exponential growth bias trials. For this paper, we only analyse the responses recorded for intervals from seven to 180 days in length.

concave: the 180-day interval is on average perceived as 87 subjective days in length, and the estimated value of the Stevens' power law parameter is 0.791, with 95% confidence interval $(0.754, 0.828)$.[6]

While time perception is indeed nonlinear, it cannot fully account for the deviations from exponential discounting we observe in our data. To see this, note firstly that the estimated power parameter $\vartheta = 0.184$ for the subadditive model in Table 2 is substantially smaller than the estimated parameter of Stevens' power law. In other words, if ϑ is interpreted as reflecting the effect of time perception then the amount of time compression required to explain behaviour in the discounting task is much greater than what is implied by participants' responses to the time perception task.

Secondly, Table 3 reports estimates of the exponential, subadditive and quasi-exponential models in terms of subjective time, in which the objective interval lengths k are simply replaced by the participant-specific values of $\kappa(k)$ as defined by equation 11. Thus note that these estimates do not impose the functional form of either the Weber-Fechner or Stevens' power laws, nor do they rely upon the fitted curves in Figure 5. The estimates reveal that the non-exponential components of the subadditive and quasi-exponential models (respectively, ϑ and β) remain significantly smaller than the value of one at which these models reduce to exponential discounting. Moreover, both models improve substantially upon the exponential model as measured by the BIC. Thus our data do not support the proposition that deviations from exponential discounting can be fully explained by a nonlinear perception of time.

4.5 Individual-level estimates

In this Section, we estimate a selection of models at an individual level and use the estimates to identify the proportions of participants whose behaviour is best described by different classes of discounting models. In our sample of 302 participants, there are seven who report sooner equivalents of zero in all 15 trials (and would thus accept any BDM offer of a sooner amount) and ten who always report 30 (and would thus reject any offer in favour of a later £30); these 17 participants thus exhibit no variation in their responses. We are able to estimate models at an individual level for the remaining 285 participants.

We report individual estimates for two discounted utility models and two models of stationary relative discounting. These are the standard exponential model (equation 3), the quasi-hyperbolic model popular in behavioural economics (equation 4), the model of subadditive discounting as the best-performing aggregate model (equation 7), and our quasi-exponential model that provides a simple approximation to the best-fitting models (equation 8). Since each of these models is built upon the underlying foundation of an exponential function, the exercise can be interpreted as exploring which extensions of the exponential form best describe discounting behaviour at an individual level.

Table 4 reports summary statistics of individual estimates of the parameters of each of these models. The median estimate of the quasi-hyperbolic β parameter is 0.9345, indicating only slight present bias: there are 23 participants (8% of the individual estimation sample) for whom the quasi-hyperbolic β differs

[6] For comparison, Bradford et al. (2019, Figure 2) find that a six-month interval is perceived to be around 95 subjective days in length. Figure 5 also shows that the 14 and 30 day intervals are very slightly overweighted, however this tendency is considerably less pronounced than in Bradford et al. (2019).

Figure 5: Subjective perceptions of time

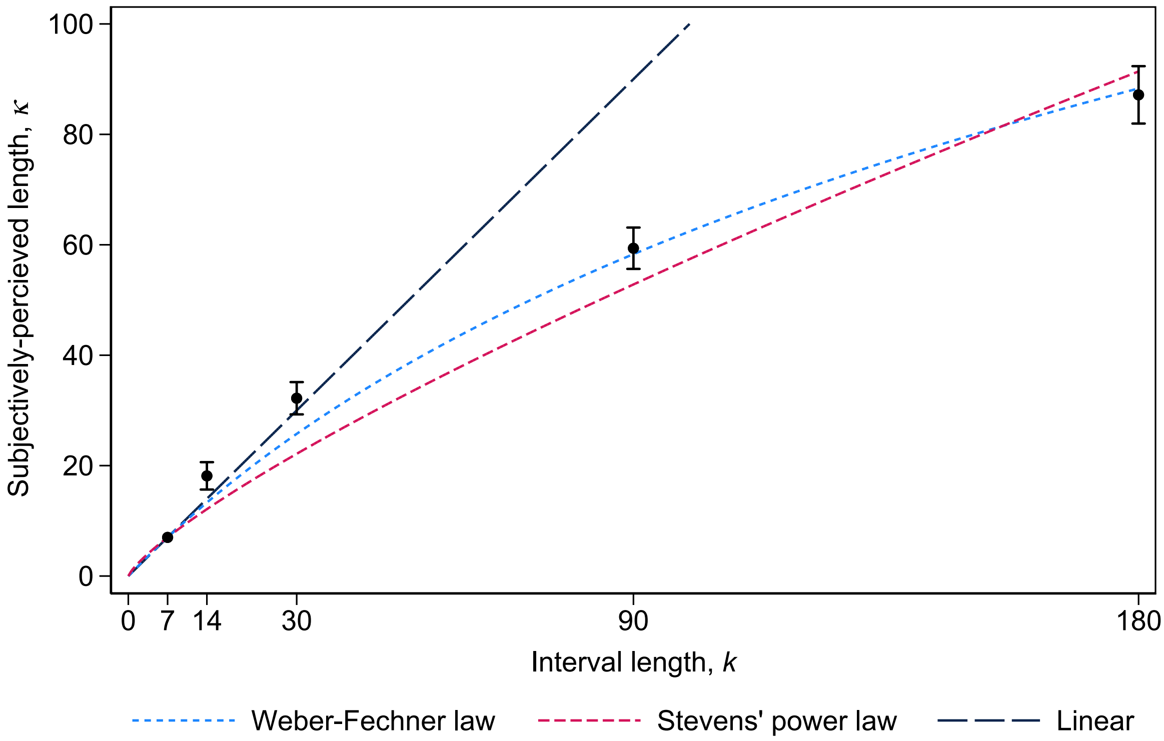

Table 3: Discounting in subjective time

	(3)	(7)	(8)
	Exponential	Subadditive	Quasi-exponential
δ	0.9938	0.7699	0.9993
	(0.9929, 0.9948)	(0.7127, 0.8272)	(0.9989, 0.9998)
ϑ		0.0918	
		(0.0189, 0.1647)	
β			0.7259
			(0.6928, 0.7590)
R^2	0.7897	0.8565	0.8567
BIC	3373	1651	1644

Notes: All models are estimated by nonlinear least squares, using 4,530 observations from 302 participants. Column numbers correspond to equations in Sections 2.1 and 2.2. Robust 95% confidence intervals, with standard errors clustered at the participant level, are in parentheses.

significantly from 1 at the 5% level. Our quasi-exponential model instead applies β to all episodes of delay and not only ones that begin in the present. The median estimate of this quasi-exponential β is 0.8753, and there are 218 participants (76%) for whom it differs significantly from 1.

Figure 6 reports the proportions of participants whose choices are best explained by each of the models, as selected by the BIC of their individual model estimates. There are 171 participants (61% of the individual estimation sample) whose choices are best described by a subadditive model, 71 (25%) for whom the quasi-exponential model is best, and 35 (12%) for whom the standard exponential model is best. Quite remarkably, there are only 6 individuals (2%) whose choices are best described by the model of present-biased quasi-hyperbolic discounted utility. These findings highlight that the vast majority of individuals exhibit stationarity, despite the fact that only a minority of them are best described by the exponential model that is the *only* stationary model of discounted utility.

5 Discussion

Our first finding is that the behaviour of our participants is stationary. Indeed, not only is aggregate behaviour stationary, so too are the choices of the vast majority of individuals. Since stationarity is a form of consistency, a potential concern is that this might not reflect time preference but rather a preference for consistency in responding to choice stimuli. We think this unlikely for two reasons. First, our interface (Figure 1) presents temporal trade-offs directly in terms of days until the sooner and later reward dates, and does not spell out the back-end delay k. For example, participants face one trial involving rewards today and in 14 days, and another involving rewards in seven and in 21 days. As such, it is not clear that a participant who wishes to respond consistently "ought to" exhibit stationarity by giving the same response in both trials. Second, since we present trials in a random order, this further complicates the task of deliberately responding in a consistent (and thus stationary) manner.

Within the framework of discounted utility, stationarity is synonymous with an exponential discount function. Our next finding is that exponential discounting describes the data poorly. As measured by BIC, it is the worst-performing model of discounted utility in Table 1. Thus, evidence of stationarity should not automatically be taken as support for exponential discounting. This point bears reiterating: the exponential model performs worse, *in explaining a stationary dataset*, than alternatives originally put forward to characterise non-stationarity. This is because non-exponential discount functions (such as the generalised hyperbola of Loewenstein and Prelec, 1992) better describe the shape of long-run discounting, despite failing to capture the property of stationarity. These are two distinct features of our data, and there is no discounted utility model that can simultaneously account for both. The tension arises from our use of extensive variation in back-end delays: by focusing on designs with more limited variation in k, researchers in economics may neglect, or even turn a blind eye to, this point.

As a particularly striking illustration of these concerns, we report a spurious estimate of quasi-hyperbolic "present bias", despite there being no present bias in the data. We emphasise that the magnitude of this estimate: $\beta = 0.833$ with 95% confidence interval $(0.809, 0.857)$, smaller than 1 with $p < 0.0001$, is entirely plausible and in line with the amount of present bias that a behavioural economist would find

Table 4: Individual-level estimates of discounting parameters

	p10	p25	p50	p75	p90
Exponential, δ	0.9304	0.9927	0.9971	0.9987	0.9995
Quasi-hyperbolic, δ	0.9459	0.9934	0.9975	0.9989	0.9996
Quasi-hyperbolic, β	0.6515	0.8119	0.9345	0.9896	1.0061
Subadditive, δ	0.3341	0.7020	0.9395	0.9893	0.9979
Subadditive, ϑ	−0.1374	0.1287	0.3495	0.6476	0.9286
Quasi-exponential, δ	0.9948	0.9972	0.9989	0.9997	1.0012
Quasi-exponential, β	0.3933	0.6530	0.8753	0.9708	0.9986

Figure 6: Best-fitting individual models

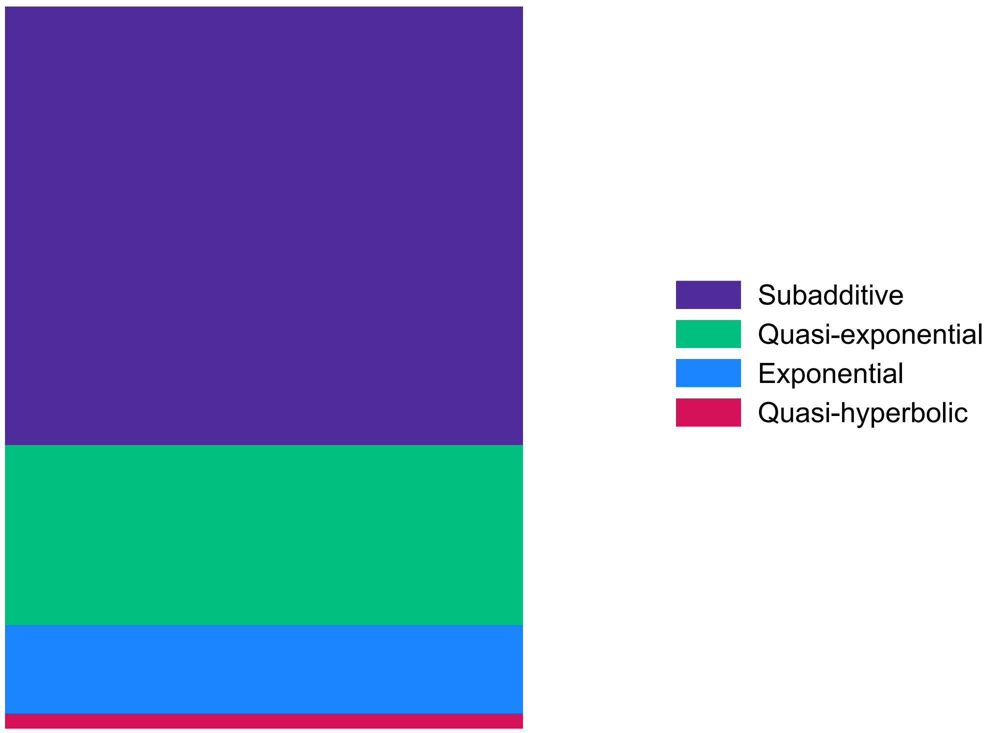

reasonable. However, as discussed in Section 4.2, it is found because, in the absence of actual present bias, the estimate of β instead models the non-exponential shape of long-run discounting.

We highlight two key methodological lessons from this finding. The first is the importance of performing model-free descriptive analysis and visualisation of the data, as reported in Section 4.1, before proceeding to structural modelling. Second, we would argue that behavioural economists should be more modest in claiming to be able to "identify" model parameters through experimental design. An oft-repeated assertion in the literature is that estimates of the quasi-hyperbolic present bias parameter β are "identified" by exogenous variation in the front-end delay. Our result illustrates how such claims may be misleading: our estimate of $\beta = 0.833$ is well-identified in this sense, yet it would be fallacious to interpret it as evidence of "present bias". By contrast, our estimate of the β parameter in QED is identified by assumption of the functional form. Nonetheless, we would argue that the latter model represents, both empirically and descriptively, a more accurate interpretation of behaviour.

The bulk of empirical research generalising the standard model of discounting focuses on relaxing the assumption of an exponential discount function. By contrast, in order to account for both stationarity and the shape of long-run discounting, we find it is necessary to relax discounted utility itself, invoking a relative discounting framework in which discounting is measured from time t instead of time 0. For our purposes, relative discounting allows us to explore models that are stationary yet not exponential. Models of relative discounting have also been proposed to account for the finding of subadditivity, which is similarly incompatible with any model of discounted utility (Read, 2001).

Of course, this added flexibility comes at a cost: because there is no single discount function that governs all temporal trade-offs, relative discounting models admit violations of transitivity, albeit only ones induced by the passage of time (Ok and Masatlioglu, 2007). Since the axioms of relative discounting entail that static preferences over outcomes are transitive (and indeed stable over time), any non-transitivity must arise solely from the treatment of time. For example, Ok and Masatlioglu (2007), outline an application of their framework to an alternating-offers bargaining game. They posit that an individual may be indifferent between an immediate agreement and a short delay, yet strictly prefer the immediate agreement to a sequence of such short delays. They consider such preferences to be descriptively plausible, and contend that a theory of discounting ought to be flexible enough to accommodate them.

Our results raise the question of why decision-makers would encode delay in relative terms. We speculate that this may be attributable to efficient-coding mechanisms, similar to those that shape the utility function (see Glimcher, 2022, for a review). Given that our brains encode value with limited precision, they likely utilise a relative coding specific to the current decision context. Applying a single discount function across all contexts would, in a noisy brain, make it harder to discriminate between alternatives within any given context, and increase the likelihood of error. While economists have not traditionally thought about discounting in terms of noisy perception, our main findings – that discounting is stationary yet non-exponential – are broadly consistent with the model set out in Vieider (2021).

Our final contribution is to propose QED as a simple and intuitive model of stationary relative discounting. QED relaxes discounted utility in an analogous manner to how quasi-hyperbolic discounting relaxes the discount function. Whereas quasi-hyperbolic discounting applies β specifically in the presence of an immediate reward (and is thus a non-stationary model of discounted utility), QED applies β to all instances of delay (and is thus a stationary model of "non-discounted utility"). Although QED does not quite match the explanatory power of the best-fitting models, we find that the cost of the approximation is minor compared to that of any model of discounted utility. At the same time, QED enjoys certain attractions over its more flexible counterparts, foremost among which are the intuitive ease of interpretation and analytical tractability that it shares with exponential and quasi-hyperbolic discounted utility. Indeed, QED may be the simplest and most parsimonious model yet proposed that can account for subadditive discounting. We hope that in future it may prove useful not only in empirical research, but also in theoretical work exploring the behavioural implications of moving beyond discounted utility.

References

Abdellaoui, M., Bleichrodt, H., l'Haridon, O., Paraschiv, C., 2013. Is there one unifying concept of utility? An experimental comparison of utility under risk and utility over time. Management Science 59, 2153–2169.

Andreoni, J., Sprenger, C., 2012. Estimating time preferences from convex budgets. American Economic Review 102, 3333–3356.

Becker, G.M., DeGroot, M.H., Marschak, J., 1964. Measuring utility by a single-response sequential method. Behavioral Science 9, 226–232.

Bradford, W.D., Dolan, P., Galizzi, M.M., 2019. Looking ahead: Subjective time perception and individual discounting. Journal of Risk and Uncertainty 58, 43–69.

Cheung, S.L., 2020. Eliciting utility curvature in time preference. Experimental Economics 23, 493–525.

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

Frederick, S., 2005. Cognitive reflection and decision making. Journal of Economic Perspectives 19, 25–42.

Frederick, S., Loewenstein, G., O'Donoghue, T., 2002. Time discounting and time preference: A critical review. Journal of Economic Literature 40, 351–401.

Glimcher, P.W., 2022. Efficiently irrational: Deciphering the riddle of human choice. Trends in Cognitive Sciences 26, 669–687.

Kable, J.W., Glimcher, P.W., 2010. An "as soon as possible" effect in human intertemporal decision making: Behavioral evidence and neural mechanisms. Journal of Neurophysiology 103, 2513–2531.

Laibson, D., 1997. Golden eggs and hyperbolic discounting. Quarterly Journal of Economics 112, 443–478.

Levy, M.R., Tasoff, J., 2017. Exponential-growth bias and overconfidence. Journal of Economic Psychology 58, 1–14.

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

Lusardi, A., 2011. Americans' financial capability. Working Paper 17103, National Bureau of Economic Research.

Mazur, J.E., 1984. Tests of an equivalence rule for fixed and variable reinforcer delays. Journal of Experimental Psychology: Animal Behavior Processes 10, 426–436.

Myerson, J., Green, L., 1995. Discounting of delayed rewards: Models of individual choice. Journal of the Experimental Analysis of Behavior 64, 263–276.

O'Donoghue, T., Rabin, M., 1999. Doing it now or later. American Economic Review 89, 103–124.

Ok, E.A., Masatlioglu, Y., 2007. A theory of (relative) discounting. Journal of Economic Theory 137, 214–245.

Read, D., 2001. Is time-discounting hyperbolic or subadditive? Journal of Risk and Uncertainty 23, 5–32.

Samuelson, P.A., 1937. A note on the measurement of utility. Review of Economic Studies 4, 155–161.

Scholten, M., Read, D., 2006. Discounting by intervals: A generalized model of intertemporal choice. Management Science 52, 1424–1436.

Stango, V., Zinman, J., 2009. Exponential growth bias and household finance. Journal of Finance 64, 2807–2849.

Steinberg, L., Sharp, C., Stanford, M.S., Tharp, A.T., 2013. New tricks for an old measure: The development of the Barratt Impulsiveness Scale - Brief (BIS-Brief). Psychological Assessment 25, 216–226.

Strotz, R.H., 1955. Myopia and inconsistency in dynamic utility maximization. Review of Economic Studies 23, 165–180.

Takahashi, T., 2005. Loss of self-control in intertemporal choice may be attributable to logarithmic time-perception. Medical Hypotheses 65, 691–693.

Vieider, F.M., 2021. Noisy coding of time and reward discounting. Working Paper 2021/1036, Department of Economics, Ghent University.

Zauberman, G., Kim, B.K., Malkoc, S.A., Bettman, J.R., 2009. Discounting time and time discounting: Subjective time perception and intertemporal preferences. Journal of Marketing Research 46, 543–556.

Quasi-exponential discounting: Online appendices

Stephen L. Cheung
Kieran MacGibbon
Arquette Milin-Byrne
Agnieszka Tymula

The University of Sydney

A Instructions for the delay discounting task

In this Section, we will ask fifteen questions. Each question involves two dates; these dates will vary from one question to the next.

In each question you will tell us **what amount of money, if paid on the sooner date, would be equally desirable to you as receiving £30 on the later date**.

In answering these questions, you can think about an amount that would make it very hard for you to choose between receiving the amount you tell us sooner or waiting to receive £30 later. Because the sooner and later dates vary, your answers will probably also vary from one question to the next.

Since everyone has different preferences, there are no wrong answers. However, you should **think carefully about your answers and report them truthfully**. Otherwise, if one of these questions is chosen for payment, you might miss out on a payment you would prefer more.

Remember you have one chance in ten to receive a bonus payment determined by one randomly selected question from Sections I and II. Depending on which question is selected, and your response to that question, **you may receive a bonus payment as soon as today, or as late as in 187 days time**. If you are selected for a bonus payment, we will message you with the outcome later today.

If a question from this Section is selected for payment, we will draw a random value £X between £0.10 and £30 (in steps of £0.10) and compare it to your answer for that question. (All values of £X between £0.10 and £30 are equally likely to be chosen.)

- If £X is greater than (or equal to) your answer, you would receive £X on the sooner date.
- If £X is smaller than your answer, you would have to wait to receive £30 on the later date.

While this may sound confusing, what you need to know is that **it is always best to report truthfully** the sooner amount you feel is equally desirable as receiving £30 later. **Otherwise, you may end up getting a payment you prefer less**.

Example:

Suppose we asked you what amount, in 10 days, would be equally desirable to you as receiving £30 in 50 days. Let's say you feel the true answer is £22.50.

- Suppose you report truthfully. Then if the random value of £X is chosen as £23.70 (which is better than £22.50), you would receive £23.70 in 10 days.
- Suppose instead you misreport by stating an amount larger than your true answer, let's say £25. Then if £X is chosen as £23.70 (which is worse than £25), you would have to wait to receive £30 in 50 days instead. Since you actually feel that £30 in 50 days is only as good as £22.50 in 10 days, you would be better off reporting truthfully and getting £23.70 in 10 days.

This example shows how *it is not in your interest to overreport* the sooner amount you feel is equally desirable as £30 later. It can be shown that *it is not in your interest to underreport as well*.

In short, you have the best chance of receiving a bonus payment you prefer the most if you think carefully about the amount of money on the sooner date you feel is equally desirable as receiving £30 on the later date, and reporting it truthfully.

Also, since you won't know until later if one of your answers is chosen for payment, you should think carefully about all of them, treating each one as though it counted for payment.

To enter your answers, click on the line that appears below each question, then drag the slider until you are happy with the value shown. After you click Next, you will not be able to revise your answer.