Noisy coding of time and reward discounting*

Ferdinand M. Vieider†[1,2]

[1]*RISLαβ, Department of Economics, Ghent University*
[2]*RISLαβ Africa, AIRESS, University Mohammed VI Polytechnic*

December 8, 2022

Abstract

I present a model generating delay-discounting from noisy mental representations of time delays. The optimal combination of noisy signals about time delays with prior information predicts discounting to exhibit present-bias, but to be stationary and delay-dependent once up-front delays are introduced. The inherent stochasticity of the model does away with the need of independent assumptions about the error-generating process. Data collected in an experiment support the need for separate but interacting parameters to capture present-bias and delay-dependence. The noisy coding account of delay-discounting explains why non-trivial discounting is routinely observed in experiments using monetary rewards instead of consumption.

1 Motivation

Humans have been found to deviate systematically from the canonical model of constant reward discounting, according to which the weight attributed to future rewards ought to decrease exponentially with time (Samuelson, 1937). The prevalent finding is instead one of present-bias (Ainslie, 1975; Mazur, 1987; Imai, Rutter and Camerer, 2021). In addition, it has been shown that discounting is *delay-dependent*, in the sense that the measured per-period discount rate decreases in

*I acknowledge financial support from BOF at Ghent University under the project "The role of noise in the determination of preferences". The experiment has received ethics approval from the ethics committee at the Faculty of Economics and Business, Ghent University, with number *FEB 2021 M*. I am grateful to Ranoua Bouchouicha for help in executing the experiment, and to Mohammed Abdellaoui, Benjamin Enke, Thomas Epper, Xavier Gabaix, Emmanuel Kemel, Olivier L'Haridon, Paulo Natenzon, Ryan Oprea, Pietro Ortoleva, Peter Wakker, and Leeat Yariv for helpful comments and discussions. All errors remain my own.

†Sint-Pietersplein 6, 9000 Ghent, Belgium; Email: ferdinand.vieider@ugent.be

the length of the delay used to measure it (Read, 2001; Dohmen, Falk, Huffman and Sunde, 2017). Such patterns have generally been documented in experiments using monetary rewards, even though economic models predict discounting over *consumption*, resulting in difficulties in interpreting the results (Cubitt and Read, 2007; Cohen, Ericson, Laibson and White, 2020). Psychologists and economists have proposed a variety of functional forms to model such discounting patterns (Mazur, 1987; Loewenstein and Prelec, 1992; Laibson, 1997; Read, 2001; Ebert and Prelec, 2007; Bleichrodt, Rohde and Wakker, 2009). The different functional forms, however, are often specialized for the description of specific behavioural patterns, and their relative empirical merit is still debated.

I present a model of time discounting based on the premise that people form probabilistic mental representations of time delays, which are then optimally combined with prior information to arrive at actionable choice parameters. Building on basic neural processes, the model is generative in nature. Starting from an optimal choice rule, the functional forms predicting behaviour emerge from the optimal combination of the noisy signal about the choice stimuli with prior information about the choice environment. When considering delays from the present, this setup yields micro-foundations for the constant-sensitivity discount function proposed by Ebert and Prelec (2007). The need for parsimonious mental representations of the numerical magnitudes involved furthermore results in a discontinuity when the time delay goes to zero, which predicts present-biased behaviour. Once up-front delays are introduced, however, the model predicts behaviour that is delay-dependent but stationary. That is, up-front delays do not matter for discounting except when the earlier payoff obtains immediately.

Traditional functional forms often confound several aspects of discounting. Generalized hyperbolic functions as proposed e.g. by Mazur (1987), Loewenstein and Prelec (1992) and Ebert and Prelec (2007) require one and the same parameter to describe strongly decreasing impatience, present-bias and delay-dependence, depending on the stimuli used to identify the parameters.[1] Other functional forms

[1]Throughout the paper, I define strongly decreasing impatience as $(y, \tau_s) \sim (x, \tau_\ell) \rightarrow (y, \kappa + \tau_s) \prec (x, \kappa + \tau_\ell) \ \forall \ \tau_s, \tau_\ell, \kappa$, where $x > y$ are monetary outcomes and $\tau_s, \tau_\ell, \kappa$ are time delays,

are geared towards capturing specific aspects of behaviour while neglecting others. For instance, the subadditive function proposed by Read (2001) is optimized for delay-dependence, at the expense of present-bias and decreasing impatience; the function proposed by Phelps and Pollak (1968) and popularized by Laibson (1997) captures present bias at the expense of strongly decreasing impatience and delay-dependence. Confounding different behavioural phenomena, or focusing on one specific phenomenon, may impact the predictive performance of the models when extrapolating across contexts, and hence their reliability for policy design.

The noisy coding of time (henceforth: *NCT*) model, on the other hand, generates predictions of present-bias and delay-dependence based on separate mechanisms, which are however closely intertwined. Its predictions rest on a parsimonious parameter space, and the theoretically-derived interactions between the parameters put tight bounds on their values. I test the predictions emerging from the model in an experiment geared at detecting the predicted patterns, as well as patterns that are *not* predicted by the model, such as the strongly decreasing impatience entailed by models of generalized hyperbolic discounting. The data provide clear evidence for present bias when the sooner delay is truly immediate, as has been documented in other recent contributions (Balakrishnan, Haushofer and Jakiela, 2020; Imai et al., 2021). There is also strong evidence for delay-dependence, with measured impatience decreasing systematically in the length of the delay used to measure it. I do, however, find no evidence for strongly decreasing impatience, i.e. once an initial upfront delay has been introduced, impatience does not decrease further as this upfront delay is increased. Testing the predictive performance of the model against standard models augmented by an additive error term, I find the noisy coding model to significantly outperform those models. I further document a strong correlation of coding noise with mathematical abilities. This suggests that the observed patterns do indeed arise from the noisy coding of numerical quantities, since the latter has been shown to underly mathematical ability (Halberda, Mazzocco and Feigenson, 2008; Park and Brannon, 2013).

and $\tau_s < \tau_\ell$. I define present bias where the equivalent condition only holds when when the sooner outcome obtains immediately, i.e. $\tau_s = 0$.

The origin of discounting patterns in noisy perception of time serves to explain why non-trivial discounting is routinely observed in experiments using monetary rewards, instead of only occurring for consumption as predicted by economic theory. This finding has long been considered puzzling, seen how it would need to assume immediate consumption of any monetary payouts (Frederick, Loewenstein and O'Donoghue, 2002; Cubitt and Read, 2007; Cohen et al., 2020). Since in the NCT model delays are subject to noisy perception, it would indeed predict non-standard discounting to occur for money as well as for consumption.[2] The patterns I document also have policy implications. For one, the model suggests that a welfare maximizing policy maker can safely disregard deviations from exponential discounting, since such deviations are due to the noisy perception of choice parameters rather than to true underlying preferences. The results suggest furthermore that strong discounting for short time delays can be counteracted by measures that reduce noisiness in time perception by either facilitating the objective perception of time delays, or increasing attention towards a decision problem.

The NCT model I present is based on an influential theoretical paradigm from neuroscience, according to which the human mind functions like a Bayesian inference machine (Knill and Pouget, 2004; Doya, Ishii, Pouget and Rao, 2007; Vilares and Kording, 2011). Applied to human decision-making, the paradigm holds the promise of unifying many well-known stylized facts, which have been modelled using a variety of disparate ad-hoc models. For instance, noisy coding of rewards may result in small-stake risk aversion (Khaw, Li and Woodford, 2021; Frydman and Jin, 2022). Noisy coding can generate functionals such as described by prospect theory, while at the same time explaining violations of prospect theory (Vieider, 2021). Such models are consistent with recent empirical findings that document how systematic biases in decision making may stem from 'cognitive uncertainty' (Enke and Graeber, 2019; 2021), or from preferences for simplicity (Oprea, 2022).

The Bayesian brain paradigm on which the NCT is built is furthermore part

[2]This does not imply that discounting for money and consumption necessarily need to be *the same*. Indeed, if outcomes are also subject to noisy cognition, we should expect such patterns to differ in systematic ways.

of a larger paradigm of *predictive coding* or *active inference* (Rao and Ballard, 1999; Friston, 2005). According to this paradigm, the mind builds models of the outside world by learning probabilistic prior distributions and actively testing their validity. These priors then serve to decode the incoming signals about the outside world, and thereby allow to make inferences about the stimuli triggering these signals. This system is fully self-organizing, so that models of the world can be built from the signals alone. These models are then checked continuously against the outside reality. This happens by the generation of prediction errors, which lead to rapid changes in the statistics of the prior. In this sense, the resulting Bayesian inference mechanism emerges from an inter-temporal optimization problem, since its parameters are endogenously adjusted to minimize the average probability of surprise over time (aka minimization of free energy; Friston, 2010).

The model I present is closely related to the model of Gabaix and Laibson (2017). The latter present a model in which a perfectly patient Bayesian decision maker receives noisy signal about future *utilities*. If the noisiness of the signal increases linearly with the time delay to the consumption outcome, the combination with a zero-mean prior then results in a hyperbolic discount function such as proposed by Mazur (1987). Seen how Gabaix and Laibson (2017) model noisy perception of future *utilities*, whereas I model noisy perception of *time delays*, the two approaches can be seen as complementary. Augmenting this model by an account whereby the attention paid to simulating future outcomes rationally increases in the outcomes themselves, Gershman and Bhui (2020) showed how the model can account for the absolute magnitude effect.

The NCT model is less directly related to accounts deriving predictions of time discounting based on randomness in choices and preferences (Lu and Saito, 2018; He, Golman and Bhatia, 2019). Such models typically postulate a specific structure for noise, which combined with exponential discounting results in apparent deviations from stationarity. These models, however, cannot separate present-bias from strongly decreasing impatience (i.e., preferences deviate from stationarity even once upfront delays are introduced), and predict no delay-dependence. Re-

cent empirical findings, however, suggest that delay-dependence is related to cognitive uncertainty while present bias is not, which is at odds with the predictions emerging from the random preference literature (Enke and Graeber, 2021). The empirical results I report suggest that such separate mechanisms are indeed needed to properly account for the observed decision patterns, thus considerably increasing the predictive performance of the model. The model furthermore presents interesting interactions with accounts of reward discounting based on the fundamental uncertainty of the future (Halevy, 2008; Epper and Fehr-Duda, 2018), to which I will return in the discussion.

2 The noisy coding of time model

2.1 Noisy logarithmic perception of time delays

I start by briefly reviewing the evidence for some of the key assumptions underlying the model. Humans (as well as animals) rely extensively on an approximate number sense. Numerosities are thought to be represented on a continuous, analog scale, and are thus perceived approximately and with noise (Dehaene and Changeux, 1993; Gallistel and Gelman, 2000; Dehaene, 2003). Such approximate number perception applies even when numbers are presented as abstract Arabic numerals (Moyer and Landauer, 1967; Piazza, Pinel, Le Bihan and Dehaene, 2007). Indeed, it is thought that even formal mathematical knowledge is built on this approximate number sense (Halberda et al., 2008; Park and Brannon, 2013). Noisiness in number perception will be especially relevant when judgments and comparisons are made quickly and intuitively, rather than being based on precise calculations (Whalen, Gallistel and Gelman, 1999).

Neuroscientific evidence further suggests that numbers are mentally represented on a compressed logarithmic scale. Such a mapping delivers parsimonious number representations, avoiding an unbounded increase in the cognitive resources necessary for the representation of larger numerosities (Dehaene and Changeux, 1993; Dayan and Abbott, 2001; Dehaene, 2003). Behavioural studies have consis-

tently shown that discriminability between numerical quantities increases in the difference between the two numbers, and decreases in the magnitude of the baseline number (Moyer and Landauer, 1967; Gallistel and Gelman, 2000; Halberda et al., 2008; Dehaene, 2011). While this may suggest fitting the evidence with a logarithmic function (Fechner, 1860), the logarithm is not unique in being able to account for such evidence. For instance, one could fit the evidence with a distribution with a mean and a standard deviation of the noise term that both increase linearly in the stimulus. Both implement Weber's law, so that for a quantity m we have $\frac{\Delta m}{m} = k$, with k the celebrated just noticeable difference (see e.g. Robson, 2001, and Netzer, 2009, for models using this notion in economics).

One argument in favour of the logarithmic implementation is the parsimony in mental resources this occupies. Direct measurement of the activity of single neurones in number discrimination tasks has further been shown to follow a log-normal distribution with increasing variance as numbers increase on a linear scale (Nieder and Miller, 2003). When numbers are represented on a logarithmic scale, however, a Gausssian density function with identical variance can be shown to fit the data best. Based on statistical parsimony, Dehaene (2003) argues that this constitutes direct evidence for logarithmic compression of numbers. It has furthermore been shown that logarithmic compression is efficient from an information-theoretic point of view, thus facilitating adaptation to diverse environments (Howard and Shankar, 2018). Notice that this conclusion is general, so that logarithmic coding will be optimal for adaptation in the presence of resource constraints *independently from the statistical distribution of stimuli in the environment*.

There is also increasing evidence that time delays are neurally encoded in the same way as described above. Importantly for the purposes of this study, the perception of time (as well as space) has been directly linked to approximate numerosity discrimination (Walsh, 2003; Cantlon, 2012; Howard, MacDonald, Tiganj, Shankar, Du, Hasselmo and Eichenbaum, 2014). It has furthermore been shown that training on numerical discrimination does not only improve numerical discrimination, but also accuracy in time perception (Cappelletti, Ges-

saroli, Hithersay, Mitolo, Didino, Kanai, Kadosh and Walsh, 2013). When combined with stimulation of the areas of the brain where the approximate numerical sense resides, such training indeed showed virtually undiminished effects until at least 16 weeks after the training ended. Importantly, this constitutes *causal* evidence that both numerical perception in general and time perception more specifically relie on the approximate number sense. Behaviourally, Zauberman, Kim, Malkoc and Bettman (2009) have shown that the dependence of subjectively perceived delays on objective delays may be best fit by a logarithmic function.

The logarithmic representation of time delays described above may incur into issues when a time delay τ becomes small. Since $\lim_{\tau \to 0} ln(\tau) = -\infty$, such a case would result in an unbounded increase in resources necessary to represent small time delays. This case would thus contradict the very resource-saving rationale underpinning the logarithmic representation (Howard and Shankar, 2018). To circumvent this issue, neuroscientific studies of continuous variables typically assume that a constant is added to the variable to be represented. At times, this constant has been simply put to 1 to avoid numerical issues (Petzschner and Glasauer, 2011). Howard and Shankar (2018), however, argue for a generic constant κ. I will thus substitute $t \triangleq \kappa + \tau$ to model mental representations of the objective delay τ in order to prevent the numerical representation to become boundless as $\tau \to 0$. While this transformation is inconsequential for large τ, it will have interesting implications as $\tau \to 0$. In empirical estimations, I will use both a parsimonious version, putting $\kappa = 1$, and a version in which κ is endogenously estimated.

2.2 Bayesian inference on time delays

I start by modelling the probability with which a decision-maker chooses a larger-later reward x, paid at time τ, over a smaller-sooner amount y, paid immediately. The central idea underlying the model is that time delays are subject to imprecise mental representations. For simplicity, I treat rewards as being perceived objectively and without distortion (my conclusions are robust to relaxing this assumption—see online appendix E). I start from a choice rule devoid of subjective

parameters[3], under which the larger-later amount will be chosen whenever

$$exp(-\tau)x > y. \tag{1}$$

The choice rule is optimal inasmuch as discounting is exponential, thus resulting in consistent choice patterns over time. The substantial discounting predicted by this choice rule reflects the fundamental intuition underlying the model that an appreciation of future rewards is inherently a learned characteristic. This take is consistent with high discount rates observed in children (Mischel and Ebbesen, 1970; Mischel, Shoda and Rodriguez, 1989; Bettinger and Slonim, 2007), and with the emphasis put on education in economic models endogenizing discounting (Doepke and Zilibotti, 2014; 2017).

The key assumption underlying the model is that the time delay, τ, is not perceived objectively, but rather encoded by a mental signal, r, which can be thought of as a neuronal firing rate. Given limits on neuronal resources, as well as limits imposed by the precision with which such a signal can be optimally decoded (Dayan and Abbott, 2001, chapter 3), this signal will typically be noisy. I further assume that the choice rule is mentally implemented on a log-log scale, which is convenient for computational reasons (Gold and Shadlen, 2001), but does not impact the qualitative conclusions from the model (see online appendix B for a discussion and a robustness analysis). The mental choice rule under which the larger-later reward will be chosen over a smaller-sooner reward paying immediately will then take the following form:

$$E[ln(t)|r] < ln\left[-ln\left(\frac{y}{x}\right)\right], \tag{2}$$

where $E[ln(t)|r]$ is the expectation of the posterior distribution of the time delays conditional on the noisy mental signal r, and where I substituted $t \triangleq \kappa + \tau$ for the objective delay τ in order to prevent the numerical representation to become

[3]Note that augmenting the choice rule by a normatively low discount rate or by a concave utility function does not affect any of the conclusions derived below. I thus focus on a minimalistic setup deprived of any further motives.

boundless as $\tau \to 0$ (see discussion above).

I model the mental signal as a single draw from the distribution:

$$r \sim \mathcal{N}(ln(t), \nu^2), \tag{3}$$

where the parameter ν quantifies coding noise, and is assumed to be independent of τ. The normal distribution derives its justification from the observation that the signal r derives from the activation of thousands of neurones. Whatever the activation function of a single neuron, the composite signal will thus follow a normal law due to the central limit theorem. Furthermore, neural evidence directly points to the normal distribution, since the best fit to the activation of individual neurones is obtained by a normal density with equal variance when numbers are represented on a log scale (Nieder and Miller, 2003; Dehaene, 2003).

Given the signal r, the mind needs to decode the information contained in the signal to infer the outside stimulus that has generated the signal. To do this, the mind builds causal models of the world, enshrined in probabilistic prior distributions learned over time, which enable causal inference about the underlying triggers of the signals. This is achieved by combining the noisy signal r with a prior indicating the probability of different time delays in the environment. Such a process is indeed optimal given the noisiness in the signal, thus reducing error. I will assume a conjugate prior distribution from the normal family:

$$ln(t) \sim \mathcal{N}(\mu, \sigma^2). \tag{4}$$

The log-normal form for the prior seems reasonable in a context where most delays one faces are short, while some are very long, while naturally reflecting the non-negative nature of time delays. Depending on its parameters, it can accommodate a wide variety of patterns. More to the point, the conjugate form to the normal likelihood has important analytical advantages. This not only implies the possibility to show an analytical solution in this paper. Much more importantly, it implies computational convenience for the mind. Obtaining the posterior mean

only requires storing in memory and manipulating 3 quantities. More complex, non-conjugate distributions, on the other hand, may require simulations with thousands of data points to achieve accurate approximations of the posterior, making them an unlikely candidate for neural transformations.[4] Much like the previous assumptions, the distributional assumptions are thus chosen based on their realism in terms of parsimonious neural computations.

Combining the likelihood in equation 3 with the prior in equation 4 by Bayesian updating (see online appendix A for details), we obtain the posterior expectation:

$$E[ln(t)|r] = \beta r + ln(\alpha), \tag{5}$$

where $\beta \triangleq \frac{\sigma^2}{\sigma^2+\nu^2}$ and $\alpha \triangleq e^{(1-\beta)\mu}$. The parameter β governs the degree to which we may expect regression of the posterior mean to the mean of the mental prior, μ. Unexpected, noisy signals falling far from the mean of the prior will thus be discounted more heavily than expected signals characterized by high precision.

Substituting the posterior expectation in equation 5 into the mental choice rule in equation 2 yields a threshold equation capturing the conditions under which the larger-later reward will be chosen over the smaller-sooner reward:

$$r < \beta^{-1}\left[ln\left(-ln\left(\frac{y}{x}\right)\right) - ln(\alpha)\right]. \tag{6}$$

This threshold equation fully describes the condition under which the larger-later amount will be chosen. To result in testable predictions, however, we need to obtain a choice rule devoid of the mental signal r, which is unknown to external observers. Obtaining the z-score for r from the likelihood in equation 3 and comparing it to the z-score of the threshold equation (see online appendix A for

[4] An alternative is to abandon analytical tractability all together, and instead to reintroduce computational tractability by the brain using some form of algorithmic approximation to the posterior. This can be done in an active inference setup by choosing e.g. the posterior mode based on a gradient ascent algorithm, which allows approximating the posterior by some convenient distribution, but adding the divergence between the true and approximate distribution as a penalty term (see e.g. Friston, 2005; 2010; Friston, FitzGerald, Rigoli, Schwartenbeck and Pezzulo, 2017). Notice that—as long as we keep the assumption of noisy logarithmic coding of time—the central results I derive are unlikely to change fundamentally in such a setup.

details), we obtain the following probabilistic choice rule:

$$Pr[(x,\tau) \succ (y,0)] = \Phi\left(\frac{ln\left[-ln\left(\frac{y}{x}\right)\right] - [ln(\alpha) + \beta \times ln(t)]}{\beta \times \nu}\right), \tag{7}$$

where $Pr[(x,\tau) \succ (y,0)]$ represents the probability of the larger-later reward being chosen, and Φ represents the standard normal distribution. In contrast to traditional discounting models, which typically combine a deterministic preference model with an independently and arbitrarily chosen stochastic choice model (He et al., 2019), the probabilistic setup used to derive the noisy coding model produces both the choice parameters *and* the stochastic choice setup, thus resulting in an inherently stochastic model of inter-temporal choice.

2.3 Noisy coding of time delays generates time insensitivity

We can zoom in on the indifference condition implied by the nominator of equation 7 to try and map the stochastic predictions of the model into a more easily interpretable format. Notice that this exercise serves purely the purpose of interpreting the result, with the empirically applicable choice rule remaining in the form of (7). We obtain:

$$ln(\alpha) + \beta \times ln(t) = ln\left[-ln\left(\frac{y}{x}\right)\right]. \tag{8}$$

Taking the expression back to the original scale by exponentiating twice, and exploiting the fact that the sooner amount divided by the later amount will give us a measure of the discount factor under the linear outcome assumption maintained, i.e. $\delta_\tau \triangleq \frac{y}{x}$, we obtain:

$$\delta_\tau = exp(-\alpha t^\beta). \tag{9}$$

The right-hand side of this expression corresponds to the constant-sensitivity discount function proposed by Ebert and Prelec (2007). The time-discriminability parameter β is closely connected to the noisiness of time perception in the present setup. The parameter α is mostly related to expected delays in a given environ-

ment. It is, however, indirectly influenced by noise in perception as well. Ebert and Prelec (2007) show that α constitutes a demarcation between the near future and the far future. Discounting for delays less than $\frac{1}{\alpha}$ will increase in time insensitivity, whereas discounting for delays lager than $\frac{1}{\alpha}$ will decrease. This insight resonates with the interpretation of the NCT model, whereby α captures regression of time perception to the mean of the prior, with opposite biases resulting to either side of it.

There is, however, one sense in which the function is rather different from the original constant sensitivity function. Given the definition of the encoded time delay as $t \triangleq \kappa + \tau$, the function will show a discontinuous drop in the vicinity of $\tau = 0$. That is, $\lim_{\tau \to 0} D(t) = exp(-\alpha\kappa^{\beta})$, where $D(t)$ indicates the mental discount function, and where I assume as usual that $D(0) = 1$, so that immediate payments are not discounted (in the absence of any delay being indicate, there is no time delay to be encoded). For $\kappa = 1$, this expression further simplifies to $\lim_{\tau \to 0} D(t) = exp(-\alpha)$. This introduces a drop in the vicinity of $\tau = 0$ which predicts present-bias—a preference for immediately received rewards over even slightly delayed rewards—which interacts with the other model parameters.

Given the reliance of the model on the noisy discounting of *time delays*, these discount patterns are predicted to occur no matter the outcomes. That is, the model departs rather markedly from traditional accounts in economics, whereby discounting is defined over consumption. The NCT model thus accounts for the fact that non-trivial discounting is routinely observed in experiments with money, which is difficult to rationalize based on standard models in economics (cfr. Cohen et al., 2020, section 4.2). This may also explain why discounting may be impacted by the format of stimulus presentation. Ebert and Prelec (2007) showed that time sensitivity can be improved by presenting visual proxies for time duration, while it will decrease under time pressure, which is consistent with the noisy perception account presented here. Read, Frederick, Orsel and Rahman (2005) showed that when future payouts are described in terms of the payout date instead of the delay, discount rates are considerably lower, and discount functions deviate less from the

exponential benchmark. This is highly consistent with the account presented here, where time *delays* are subject to noisy coding.

2.4 Discounting with up-front delays

The choice rule being implemented will depend on the type of stimuli faced by the decision-maker. The question of how behaviour may change depending on the choice stimuli is thus central in the present setup. Take a tradeoff between delayed outcomes (y, τ_s) and (x, τ_ℓ), where the subscripts s and ℓ stand for *sooner* and *later*, respectively, and as usual $x > y$. The optimal choice rule will be:

$$e^{-\tau_\ell} x > e^{-\tau_s} y. \tag{10}$$

Let us further define mentally coded delays $s \triangleq \kappa + \tau_s$ and $\ell \triangleq \kappa + \tau_\ell$. Rearranging and taking the logarithm of both sides twice as above, we obtain the following mental choice rule based on a signal r_d, encoding the *delay* between the payoffs:

$$E[ln(\ell - s)|r_d] < ln\left[-ln\left(\frac{y}{x}\right)\right]. \tag{11}$$

Importantly, the time *delay* between two payments is now transformed, rather than the individual delays from the present. This rests on the observation that no analytic solution is possible for the transformation of the individual time delays within the log (this is not unique to the use of the log-log rule; see online appendix C for details). The requirement of an analytic solution in the present context does not just serve the convenience of the modeller—far from it, it is a fundamental feature required for the mental implementation of the choice rule. In keep with the general modelling setup here proposed, it is indeed essential that the mind be able to execute the necessary computations and derive a posterior expectation of the quantities involved. Notice that if delays $\tau_\ell - \tau_s$ are gauged approximately rather than calculated precisely, this further adds to the idea that their perception will be noisy. Notice also that the mental transformation of the time variables no longer matters here, since the additive κ drops out when representing the time

delay $d \triangleq \ell - s = \tau_\ell - \tau_s$.

We can now proceed as above and combine the likelihood and prior by Bayesian updating to obtain the posterior expectation. Substituting the expectation into the mental choice rule and comparing the z-scores (see online appendix C for details) yields a stochastic choice rule which maps into the following discount factor:

$$\delta_{s,\ell} = exp(-\alpha(\ell - s)^\beta). \tag{12}$$

This discount factor captures delay-dependence, since the delay between payments is compressed by $\beta < 1$. Delay-dependence aside, however, the function is stationary. Indeed, the function is identical for any delay d independent of the value taken by the up-front delay τ_s. In other words, only the *inter-reward delay* matters, while the up-front delay τ_s is inconsequential. The constant sensitivity function of Ebert and Prelec (2007), on the other hand, would take the form $exp(-\alpha(\tau_\ell^\beta - \tau_s^\beta))$, thus implying a transformation of the single delays from the present.

The fact that this function obtains naturally from an intuitive optimal choice rule may explain the pervasiveness of subadditivity (Dohmen et al., 2017; Enke and Graeber, 2021), given the model's *predictive* interpretation. Notice that while Ebert and Prelec (2007) also ascribed a causal and hence predictive interpretation to their constant-sensitivity function (p. 1436), the absence of differentiation by stimulus type in their setup confounds delay-dependence with strongly decreasing impatience, since both are captured by one and the same parameter.

2.5 Fit with the literature

In this section, I briefly review some of the most closely related discount functions used in the literature. While an exhaustive review of discounting models and functions is beyond the scope of this article, this will allow me to compare the key features of the NCT model with those of some formally related functions.

Samuelson (1937) famously introduced a model in which utilities of future outcomes are discounted by an exponentially decreasing discount function:

$$D(\tau) = exp(-\rho\tau), \tag{13}$$

where D is the discount function and ρ is the subjective discount rate. The model thus entails constant impatience and no delay-dependence. It is also deterministic—a feature that it shares will all discount functions discussed here in sequence.

While the model above remains the normative benchmark, many descriptive deviations from that benchmark have been observed over the years (Thaler, 1981; Loewenstein and Prelec, 1992; Read, 2001). Many types of discount functions have been proposed to account for such 'anomalies'. A popular form in economics is the quasi-hyperbolic discount function (Phelps and Pollak, 1968; Laibson, 1997):

$$D(\tau) = \beta \times exp(-\rho\tau), \tag{14}$$

where $\beta \triangleq 1$ for immediate payments, and $\beta < 1$ thus captures present-biased behaviour when one of the payouts is immediate. Once an up-front delay is included, β drops out and discounting is exponential. Recent meta-analyses provide clear evidence for present-bias in experiments using monetary rewards, although present-bias may be even stronger when time discounting is measured for consumption (Imai et al., 2021). Balakrishnan et al. (2020) have shown that present-bias crucially depends on the present payout being truly *immediate*.

A parallel literature has discussed hyperbolic discounting. Hyperbolic functions are often motivated by the common difference effect, an anomaly whereby the discount rate for a given time delay decreases as a common delay is added to both outcomes (Loewenstein and Prelec, 1992). Empirically, however, the function has most often been fit using delays of different length from the present (Thaler, 1981; Kirby, 1997; Ebert and Prelec, 2007; Zauberman et al., 2009), thus potentially confounding strongly decreasing impatience and delay-dependence (Read, 2001). In addition to the constant sensitivity function proposed and axiomatized

by Ebert and Prelec (2007) already discussed above, Mazur (1987) introduced an early one-parameter function capturing 'proportional discounting':

$$D(\tau) = \frac{1}{1+\zeta\tau}, \tag{15}$$

where ζ is the subjective discounting parameter. Discounting is by definition proportional to the time delay, and the function can thus not separately capture delay-dependence and strongly decreasing impatience.[5]

Another popular function is the hyperbolic function proposed in the seminal paper of Loewenstein and Prelec (1992). The function takes the following form:

$$D(\tau) = \left(\frac{1}{1+\zeta\tau}\right)^{\frac{\xi}{\zeta}}, \tag{16}$$

where $\zeta \geq 0$ measures the degree of decreasing impatience. The proportional discount function seen above emerges as a special case when $\xi = \zeta$. Again, decreasing impatience, present bias, and delay dependence cannot be separately quantified.

Bleichrodt et al. (2009) proposed a more flexible double-exponential family of functional forms, that allow to overcome limitations of the above function in capturing very strong patterns of decreasing impatience. Their constant relative decreasing impatience function converges to the function of Ebert and Prelec (2007) when immediate payouts are included. Their constant absolute decreasing impatience function takes the following form:

$$D(\tau) = \begin{cases} e^{r\times e^{-c\tau}-1} & \text{if } c > 0 \\ e^{-r\tau} & \text{if } c = 0 \\ e^{-r\times e^{-c\tau}} & \text{if } c < 0, \end{cases} \tag{17}$$

where the parameter c indicates the convexity of the discount function $D(.)$. For $c = 0$, the function converges to exponential discounting. Once again, the function is geared towards capturing decreasing impatience, but foresees no separate

[5]Technically, the function cannot account for delay-dependence at all. Devide τ into 2 equal subintervals τ_1 and τ_2. Then $\frac{1}{1+\zeta\tau_1} \times \frac{1+\zeta\tau_1}{1+\zeta\tau_2} = \frac{1}{1+\zeta\tau}$, whereas delay-dependence entails $\delta_{\tau_1} \times \delta_{\tau_2} < \delta_\tau$. Any delay-dependence will thus create difficulties for the model.

mechanisms for present-bias and delay-dependence.

Notwithstanding the large literature on generalized hyperbolic discount functions, the evidence on impatience decreasing for constant delays as the up-front delay increases is mixed at best. Read (2001) found no evidence for strongly decreasing impatience across 3 experiments. Developing and testing a novel index for decreasing impatience, Rohde (2019) found evidence for present-bias, but not for further decreases in impatience with larger up-front delays. He et al. (2019) concluded from two experiments that "decreasing impatience is not as robust as is widely held" (p. 63). Similar conclusions were reached by a number of other recent papers carefully controlling for delay-dependence (Attema, Bleichrodt, Rohde and Wakker, 2010; Cavagnaro, Aranovich, McClure, Pitt and Myung, 2016). On the other hand, Bleichrodt, Gao and Rohde (2016) report evidence for strongly decreasing impatience for both health and money. Ultimately, however, the global evidence remains difficult to assess due to the variety of measurement methods and stimulus types used in the literature. It is worth noting that the predictions of the NCT model apply to a binary choice setup, which may be difficult to map into other methodologies used in the literature. For instance, Read and Roelofsma (2003) explicitly tested direct choice against matching techniques, and found some evidence for strongly decreasing impatience in matching, but none in choice tasks.

While an in-depth review of the literature on decreasing impatience is beyond the scope of this article, there is no doubt that at least *some* papers have documented such patterns. A possible explanation is that these findings may be driven by an additional motive I have not modelled. Halevy (2008) proposed a model whereby choice patterns reminiscent of hyperbolic discounting may arise in trade-offs between timed, nominally sure amounts because of the inherent uncertainty of the future. If agents have some uncertainty about future payoffs, and if they exhibit nonlinear probability distortions, then this will result in present-biased behaviour. Epper and Fehr-Duda (2018) further showed how the uncertainty of the future could result in strongly decreasing impatience, as well as accounting for a number of other stylized facts that have been documented in the literature.

The uncertainty in these models is captured by the subjective probability that a future payment will really take place and can be enjoyed. The uncertainty in the model I have presented concerns the mental representation of the time delay itself. That being said, the two types of uncertainty may well co-exist, and the model here presented could easily be generalized to include the inherent uncertainty of the future—a point to which I will return in the discussion.

Finally, Read (2001) proposed the following functional form to account for delay-dependent preferences between a sooner delay τ_s and a later delay τ_ℓ:

$$D(\tau_s, \tau_\ell) = \frac{1}{1 + \zeta(\tau_\ell - \tau_s)^\gamma}, \tag{18}$$

where ζ again captures discounting proportional to the time *delay*, and $0 < \gamma < 1$ captures delay-dependence. Kable and Glimcher (2010) propose a variant of this function imposing $\gamma \triangleq 1$, which captures what they refer to as 'as soon as possible (ASAP) discounting', so called because it captures seemingly hyperbolic patterns for increasing delays no matter what the up-front delay. The function is founded on neural data which failed to detect separate coding of an 'immediacy signal' in the brain—an observation that is consistent with the NCT model implementation of present-bias. The main differences with the NCT model are that these models cannot account for present-bias, and that the two parameters are nominally independent. However, Enke and Graeber (2021) report evidence for both present-bias and delay-dependence using monetary payoffs, and suggest that present-bias and delay-dependence may arise from different mechanisms.

Several recent models have proposed setups in which hyperbolic discounting emerges due to random in decisions and in tastes. Lu and Saito (2018) proposed a random preference model in which the discount function of a decision maker is stochastic. He et al. (2019) combined an exponential discounting model with both stochastic choice under the form of a logit model, and stochastic tastes under the form of a random preference model. A central feature of these models is that they arrive at predictions of decreasing impatience even when the underlying discounting is exponential. These models are, however, geared towards decreasing

impatience, and none of them predicts delay-dependence, making them rather distinct from the NCT model proposed here.

Gabaix and Laibson (2017) recently proposed a setup that is formally related to the model presented here. They assume Bayesian agents who are perfectly patient, but who perceive future *utilities* with some noise, so that $s_t \sim \mathcal{N}(u(x_t), \xi_t^2)$, where $u(x_t)$ is the utility of a reward x at time t, ad s_t is the noisy simulation of that utility. This simulation is combined with a prior $u(x_t) \sim \mathcal{N}(\widehat{\mu}, \zeta^2)$. This yields the posterior expectation $E[u(x_t)|s_t] = \widehat{\mu} + D(t)(s_t - \widehat{\mu})$, where $D(t) \triangleq \frac{\zeta^2}{\zeta^2+\xi^2} = \frac{1}{1+\frac{\xi^2}{\zeta^2}}$. Assuming that the error variance of the simulation increases linearly in the time delay t, i.e. $\xi_t^2 = \eta^2 t$, this yields $D(t) = \frac{1}{1+\gamma t}$, where $\gamma \triangleq \frac{\xi^2}{\zeta^2}$. For $\widehat{\mu} = 0$, this yields a stadard discounting equation akin to the one proposed by Mazur (1987). Whatever the value of μ, discounting will exhibit regression to the mean of the prior which increases with the time horizon t. Note that while the predictions are again quite distinct from those of the NCT model inasmuch as it does not predict present-bias or delay-dependence, the formal setup is similar and the insights emerging from the model are complementary to the ones presented here.

3 Experiment I: Patterns of reward discounting

3.1 Experiment

Bachelor students attending an introductory class in behavioural economics at Ghent University were invited to take part in a classroom experiment. The students had been exposed to the basics of expected utility theory, but had not covered time preferences yet, nor had they been exposed to descriptive theories of decision making under risk and uncertainty. Students were told to bring a laptop or tablet to class to participate in an experiment. They were told that the anonymized aggregate data would be used to illustrate typical aggregate choice patterns for teaching purposes. They were also told that 10 students would be randomly extracted to play one of their choices for real money immediately after the experiment. Overall, 175 students participated in the experiment and provided

a complete set of responses during the allocated time.

The time delays used in the experiment are depicted in figure 1. They were chosen using simulations to allow for optimal identification of all model components. Importantly, I chose the stimuli in such a way as to allow for the identification of patterns predicted by the model—such as present-bias and delay-dependence—as well as to allow for the identification of patterns *not* predicted by the model—such as strongly decreasing impatience. In particular, comparison of AB to BC and of AC to CE allows for the identification of present-bias. Comparison of AB and BC versus AC, of CD and DE versus CE, and of all the 6 week delays (AB, BC , CD , DE) and 12 week delays (AC , CE) with the full delay over 24 weeks (AE) allow for the identification of delay-dependence. Strongly decreasing impatience can be identified from the comparison of BC, CD, and DE.

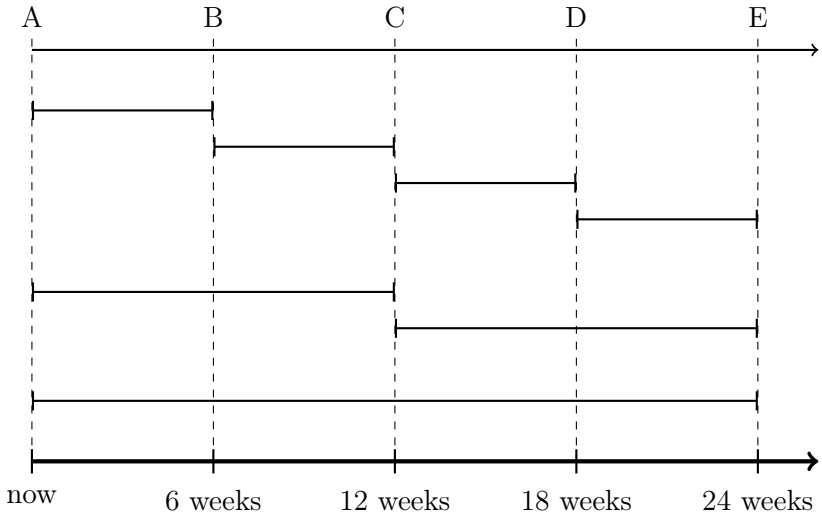

Figure 1: Time horizons used in experiment

Illustration of time delays used in the experiment. The maximum delay, indicated by AE, was 24 weeks. This delay was divided into 4 different sub-periods of 6 weeks, AB, BC, CD, and DE; and into 2 different sub-delays of 12 weeks, AC and CE. Comparison of AB to BC and of AC to CE allows for the identification of present-bias. Comparison of BC to CD and DE, and of the latter two, allow for the identification of strongly decreasing impatience. Comparison of long delays with their constituent parts allow for the identification of delay-dependence. Notice that, give 4 delays of 6 weeks, 2 of 12 weeks, and 1 of 24 weeks, the stimuli are indeed well-fiitted by a log-normal distribution.

The future outcome was fixed at €50. The choice of such a round, invariant amount was meant to ensure that outcomes are perceived objectively, rather than being subject to noisy perceptions themselves (see online appendix E for a discussion of robustness to this assumption). The earlier amounts ranged between

€33 and €49 inclusive, allowing for discount factors between 0.66 and 1 (discount rates between 0 and 52%) per period. Each screen presented one single choice, and the individual choice pairs were presented completely at random. This design was implemented to fit the discrete binary choice setup predicted by the model. Subjects made a total of 158 choices, which took about 20 minutes on average. All stimuli were presented at least once. In addition, 40 randomly selected stimuli were repeated. The repeated extractions were executed with replacement, so that the same stimulus may recur a number of times. Identification of decision noise, which plays a central role in the model, is thus assured by i) repeat observations of the same stimuli; and ii) monotonicity violations between similar stimuli.

Before the start of the experiment, the lecturer presented the instructions (online appendix G), to make sure that everybody had an understanding of the tasks. Since the tasks consisted of a binary choice paradigm, the actual explanations of the tasks were very simple. The lecturer, however, emphasized the procedural aspects of the payout mechanism. Both immediate and future payouts were made by bank transfer. Bank transfers between all major Belgian banks are immediate. This was emphasized in the instructions, and subjects were told that in case of an immediate payout being extracted to count for real pay, the lecturer would execute the payment directly and wait for the money to arrive on the student's account. The student would then be asked to verify if the money had arrived and sign a receipt. In case of a future payment, the lecturer signed a certificate on university letterhead. The certificate contained the amount to be paid and the date on which it would be paid, and it was signed by the lecturer. The certificate also contained the address and telephone number of the lecturer, and students were encouraged to contact the lecturer in case they changed bank accounts or they had any doubts about the payment. All time delays were chosen in such a way as to fall within the same academic year, to keep the costs of approaching the lecturer low, and to further reassure subjects of the future payment guarantee.

3.2 Econometric approach

In addition to the nonparametric analysis, I will present structural estimations of the model parameters. The model presented above is inherently stochastic, so that it can be directly implemented without any further need for separate assumptions about the error structure. I augment traditional models, which are deterministic in nature, by a normally distributed, additive error term, which constitutes the most common error model used in the literature. While some error models have been proposed that may result in non-stationary discounting patterns even when starting from an optimal setup of exponential discounting (Lu and Saito, 2018; He et al., 2019), these error models result in predictions that are quite distinct from those emerging from the model here presented, and which are rather closer to traditional setups—a point to which I will return in the discussion.

I use a Bayesian random-parameter setup to obtain individual-level parameter estimates jointly with aggregate estimates, which serve as endogenously-estimated priors for the individual estimates. Compared to purely aggregate estimates, such a setup has the advantage of producing individual-level estimates of the parameters of interest; compared to individual-level estimates, such a model discounts noisy outliers, thus resulting in increased predictive performance (Conte, Hey and Moffatt, 2011; von Gaudecker, van Soest and Wengström, 2011; Baillon, Bleichrodt and Spinu, 2020). I thus obtain the posterior estimate $\mathcal{L}_i(\boldsymbol{\theta}_n|z)$ given choice data z over the individual parameter vector $\boldsymbol{\theta}_n$ from

$$\mathcal{L}_i(\boldsymbol{\theta}_n|z) \propto p(z|\boldsymbol{\theta}_n) \times p(\boldsymbol{\theta}_n), \tag{19}$$

where z takes the value 1 if the larger-later reward is chosen and 0 otherwise, and where the likleihood $p(z|\boldsymbol{\theta}_n)$ is defined as follows:

$$p(z|\boldsymbol{\theta}_n) = (Pr[(x, \tau_\ell) \succ (y, \tau_s)])^z \times (1 - Pr[(x, \tau_\ell) \succ (y, \tau_s)])^{1-z}, \tag{20}$$

with $Pr[(x, \tau_\ell) \succ (y, \tau_s)]$ taking the form of the choice probability in equation 7 for the noisy coding model. For the standard models estimated below, on the other

hand, the choice probability will be defined by the functional form underlying the model, plus a normally distributed additive error term, assumed to take the form of 'white noise' (Hey and Orme, 1994; Bruhin, Fehr-Duda and Epper, 2010).

Finally, the prior distribution for the individual-level parameters, $p(\boldsymbol{\theta}_n)$, takes the following form:

$$p(\boldsymbol{\theta}_n) = \mathcal{N}(\overline{\boldsymbol{\theta}}, \boldsymbol{\Sigma}), \tag{21}$$

where $\overline{\boldsymbol{\theta}}$ is a vector containing the aggregate parameter means, and $\boldsymbol{\Sigma}$ is a variance-covariance matrix of the individual-level model parameters. Both $\overline{\boldsymbol{\theta}}$ and $\boldsymbol{\Sigma}$ are endogenously estimated from the aggregate data, and serve as priors for the individual-level estimates. The hyperpriors for the parameters in $\overline{\boldsymbol{\theta}}$ and $\boldsymbol{\Sigma}$ are chosen in a way as to be mildly regularizing, thus helping the algorithm to converge, but being wide enough to accommodate any plausible parameter values that may emerge from the data. This follows best practices in Bayesian econometrics (McElreath, 2016), and the estimates reported below are not sensitive to changes in the hyperpriors used, given that the amount of data can easily overpower any prior at the aggregate level.

I maximize the logged sum over the choice-level observations i of the likelihood function described above using Bayesian simulations in Stan (Carpenter, Gelman, Hoffman, Lee, Goodrich, Betancourt, Brubaker, Guo, Li and Riddell, 2017), launched from RStan (Stan Development Team, 2017). I conduct comparisons between different models using leave-one-out cross-validation (Gelman, Hwang and Vehtari, 2014; Vehtari, Gelman and Gabry, 2017). The choice of cross-validation methods for model selection serves to avoid the overfitting of existing data, instead focusing on the *predictive* performance of the models. This ensures coherence with the random-parameter setup used, and constitutes a more adequate test of model fit in the present setting than alternative methods geared towards optimizing the fit to existing data.

3.3 Results

I start the analysis with an examination of the nonparametric evidence for present-bias, shown in figure 2. The figure plots decumulative choice proportions for the larger-later reward as the sooner-smaller reward increases from €33 through €49. Panel 2(a) compares the choice patterns for a 6 week delay from the present to those for a 6 week delay from an upfront delay of 6 weeks (i.e., for the delays AB vs BC in figure 1). The choice proportions for the 6-week delay from 6 weeks are shifted to the north-east of the choice patterns for the 6 week delay from the present, thus indicating present-bias ($p = 0.011$, two-sided Wilcoxon signed-rank test on individual-level choice proportions for the later option). Results for the 12 week delay from the present versus a 12-week delay from 12 weeks, shown in panel 2(b), are very similar, and again indicate an increase in patience following the introduction of the up-front delay ($p \ll 0.001$).

Figure 3 examines the evidence for strongly decreasing impatience by showing discounting patterns for 6-week delays from up-front delays of 6, 12, and 18 weeks (i.e., comparing time delays BC , CD, and DE). No clear differences in discounting for the different upfront delays is apparent from the graph. Statistically, there is no difference in discounting between a 6-week delay from 6 weeks and a 6-week delay from 12 weeks ($p = 0.393$). There is also no difference when comparing a delay of 6 weeks from 12 weeks to a 6-week delay from 18 weeks ($p = 0.182$), or when comparing a 6-week delay from 6 weeks to a 6-week delay from 18 weeks ($p = 0.076$). Although this last comparison is marginally significant, it goes in the opposite direction of the previous comparison, resulting in an overall null result.

To test delay-dependence, I identify a probabilistic non-parametric discount factor from the sooner amount at which a subject starts choosing the smaller-sooner reward 50% of the time, divided by the later amount. Figure 4 shows two examples of delay-dependence. Panel 4(a) compares a 12 week delay from the present (AC) to its two underlying 6 week delays (AB and BC). Delay-dependence predicts $\delta_{0,12} > \delta_{0,6} \times \delta_{6,12}$, where δ is the discount factor and subscripts indicate time delays in weeks. The great majority of data points falls above the 45° line,

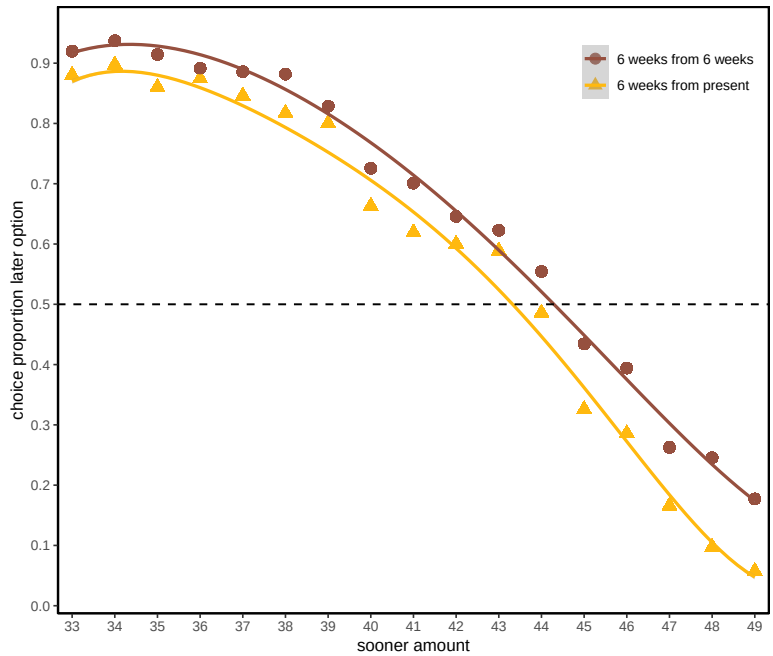

(a) present-bias, 6 weeks

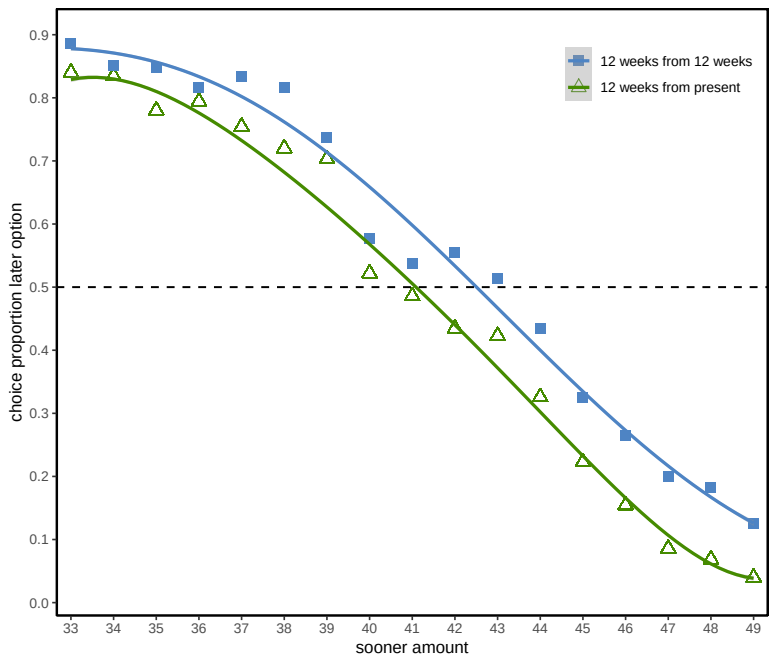

(b) present-bias, 12 weeks

Figure 2: Non-parametric illustration of present bias

Panel 2(a) shows the comparisons of a 6 week delay from the present versus a 6 week delay from an up-front delay of 6 weeks, using decumulative choice proportions for the larger-later reward as the sooner reward increases. Panel 2(b) shows the decumulative choice proportions for the larger-later reward comparing a 12 week delay from the present to a 12-week delay from an up-front delay of 12 weeks. Monetary amounts are in Euros The choice proportions are calculated as the choices for the later option relative to the total number of choices in a given choice problem. The nonparametric choice proportions are fit with a 5th degree polynomial, and include a 95% confidence interval.

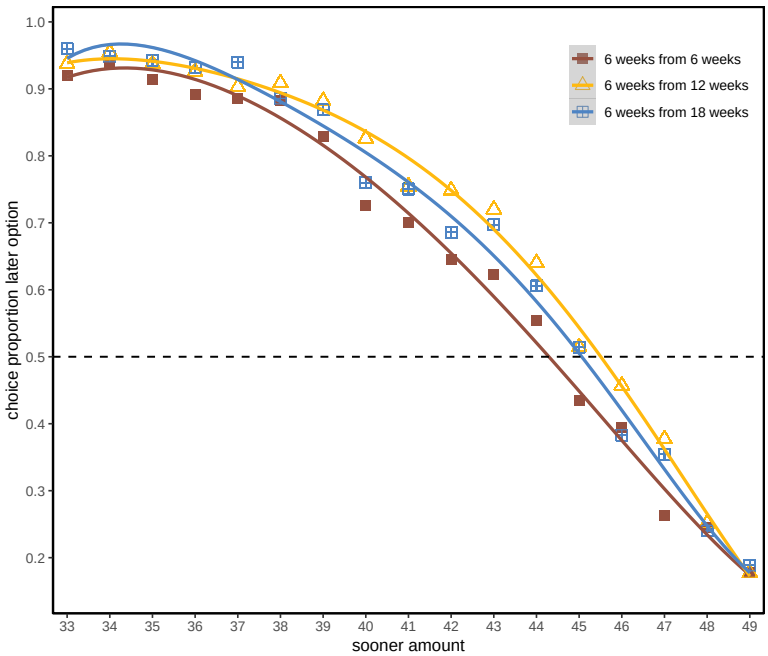

Figure 3: Discounting for 6 week delay from different up-front delays

The figure shows decumulative choice proportions for the larger-later reward as the sooner reward increases for 6-eek delays from 6 weeks, from 12 weeks, and from 18 weeks. No clear differences emerge indicating no support for generalized decreasing impatience. The nonparametric choice proportions are fit with a 5th degree polynomial, and include a 95% confidence interval.

indicating delay-dependence. Importantly, this also holds true when we exclude calculated discount factors $\delta_{0,6} \times \delta_{6,12}$ smaller than 0.66, which could otherwise bias the findings due to censoring effects (p $\ll$ 0.001).

Panel 4(b) shows the patterns for the longest 24 week delay (AE) against the product of the discount factors for the 4 underlying 6-week periods (AB, BC, CD , DE). Delay-dependence is now revealed by behavioural patterns indicating that the discount factor over the whole period is larger than the product of the discount factors for the different underlying 6-week periods, $\delta_{0,24} > \delta_{0,6} \times \delta_{6,12} \times \delta_{12,18} \times \delta_{18,24}$. This effect can be seen to be very pronounced, with all points falling to the north-west of the 45 degree line. Once again, this effect is highly significant even after accounting for censoring and excluding individuals for whom the calculated discount factor from the shorter delays is lower than 0.66 ($p \ll 0.001$). Delay-dependence in the remaining comparisons CE versus CD times DE, and AE versus AC times CE are similar, and are shown in online appendix F.

I next estimate the noisy coding model, and test its predictive performance

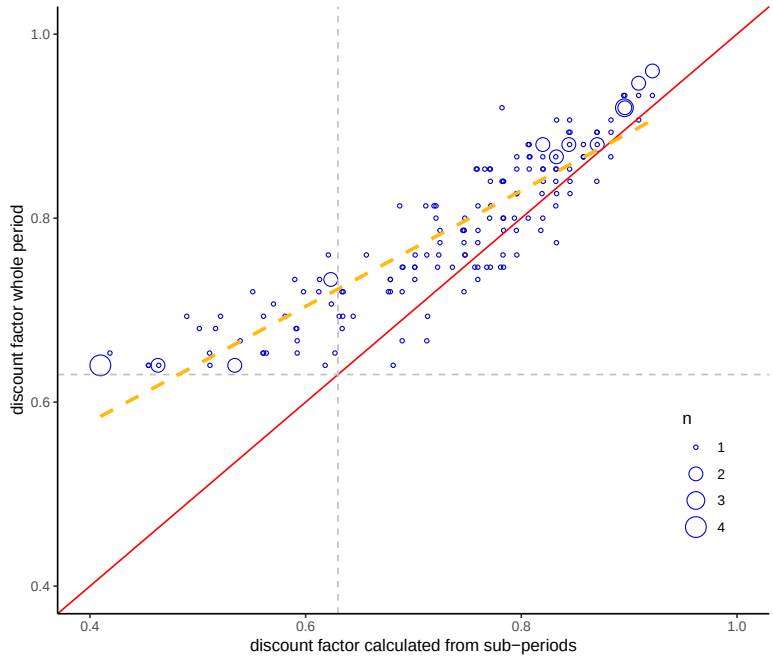

(a) delay-dependence, 12 weeks from present

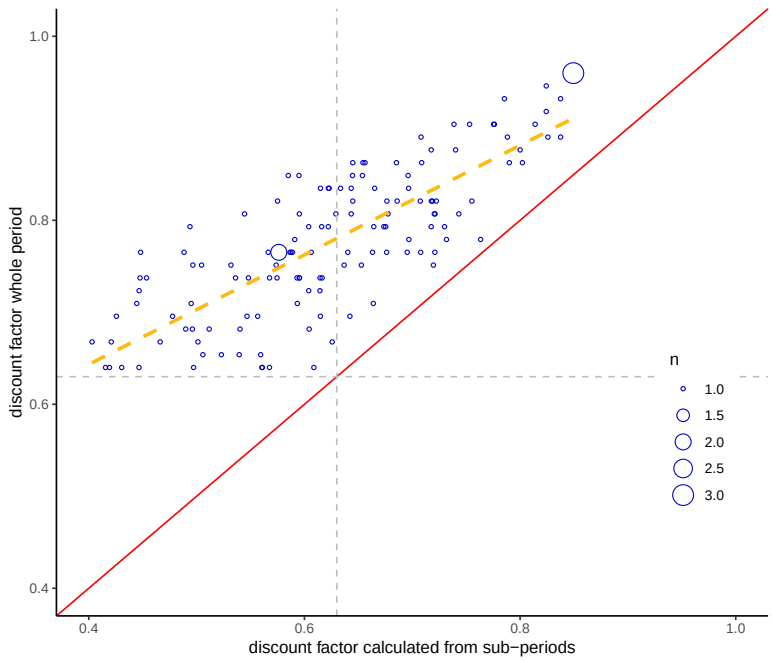

(b) delay-dependence, 24 weeks vs 6-week delays

Figure 4: Non-parametric illustration of subadditivitys

The subadditivity comparisons are obtained by comparing a discount factor over the whole period with the product of the discount factors of the subperiods. The pattern in panel 4(a) thus obtains from a comparison of the discount factor $\delta_{0,12}$ with the product of the two discount factors $\delta_{0,6} \times \delta_{6,12}$, where the subscripted numbers indicate the extremes of the time delays. The pattern in panel 4(b) obtains from a comparison of the discount factor $\delta_{0,24}$ to the product of all underlying 6-week discount factors, $\delta_{0,6} \times \delta_{6,12} \times \delta_{12,18} \times \delta_{18,24}$. Other comparisons for subadditivity are similar. Dashed lines indicate correlations.

against different standard models augmented by a normally distributed, additive error term, as most commonly used in the literature. The point here is not to see whether the NCT model outperforms every other model, combined with any possible stzochastic choice model, of which there is a staggering variety of combinations (see Regenwetter, Cavagnaro, Popova, Guo, Zwilling, Lim and Stevens, 2018, for a review). The point is rather to see whether the specific predictions of the model and the interaction between the parameters predicted may give it an edge in terms of predictive performance when compared to other, 'typical' implementations.

Table 1: Mean parameter estimates and predictive performance test

model	discount parameter	distortion parameter	noise	ELPD difference
noisy coding of time	0.028	0.656	1.	0
	(0.001)	(0.017)	()	–
noisy coding of time ($\kappa \equiv 1$)	0.029	0.683	1.006	−65.2
	(0.001)	(0.017)	(0.012)	(10.4)
subadditive discounting	0.046	0.596	0.124	−72.7
	(0.002)	(0.016)	(0.001)	(31.8)
constant sensitivity discounting	0.030	0.763	0.152	−288.6
	(0.001)	(0.014)	(0.002)	(36.9)
quasi-hyperbolic discounting	0.012	0.950	0.152	−288.6
	(0.000)	(0.003)	(0.002)	(36.6)
hyperbolic discounting	0.021	0.049	0.152	−298.6
	0.001	(0.005)	(0.002)	(37.4)
as-soon-as-possible discounting	0.016	–	0.146	−301.3
	(0.000)	–	(0.002)	(36.3)
proportional discounting	–	0.017	0.152	−329.5
	–	(0.000)	(0.002)	(37.4)
constant absolute decreasing impatience	0.014	0.938	0.154	−403.5
	(0.000)	(0.449)	(0.002)	(24.7)
exponential discounting	0.014	–	0.154	−403.5
	(0.000)	–	(0.002)	(24.7)

The parameters and standard deviations (in parentheses) shown in the table have been obtained in aggregate-level Bayesian estimations. The ELPD difference refers to the expected log pointwise predictive density difference, and measures the predictive performance of a model relative to the best-performing model measured by leave-one-out cross-validation (Vehtari et al., 2017). The ELPD difference constitutes a truly Bayesian generalization of the deviance information criterion. Results are virtually identical if I base the model tests on the Watanabe-Aikake Information criterion (WAIC) instead. Negative values indicate lower performance.

he test results, as well as the average parameter estimates, are shown in table 1. The noisy coding model easily outperforms the exponential model. It also outperforms the quasi-hyperbolic model. This is hardly surprising given the strongly delay-dependent patterns documented above, which cannot be captured by these models. However, the NCT model also clearly outperforms the hyperbolic class models (including proportional and hyperbolic discounting), as well

as the constant absolute decreasing impatience model, and ASAP discounting. The noisy coding model, furthermore, also outperforms the standard version of the constant-sensitivity model, where $\delta_{s,\ell} = exp(-\widehat{\alpha}(\ell^{\widehat{\beta}} - s^{\widehat{\beta}}))$, with the 'hat' distinguishing the parameters from the equivalent parameters in the noisy coding model. While the constant-sensitivity model can account for delay-dependence, it predicts the co-existence of delay-dependence and strongly decreasing impatience, which are described by one and the same parameter. We have, however, seen that delay-dependence is strong in the data, while strongly decreasing impatience is not supported. This explains the poor predictive performance of the constant-sensitivity model in the present data. Perhaps unsurprisingly, the model coming closest in performance is Read's subadditive model. The version of the model including a a parameter κ furthermore performs better than the version fixing that parameter to 1, with $\kappa = 4.237$, (standard error: 0.361).

This issue is further illustrated in panel 5(a) of figure 5, which plots the time-sensitivity parameter $\widehat{\beta}$ from the constant sensitivity model against the time-discriminability parameter β from the noisy coding model. The figure shows that the two parameters are widely dispersed and largely independent from each other, further highlighting that these parameters capture partly different behavioural features. Panel 5(b) further shows the correlation between the time-discriminability parameter β and an equivalent subadditivity parameter γ proposed by Read (2001) to capture delay dependence, where $\delta_{\tau_s,\tau_\ell} = exp(-\widehat{\alpha}(\tau_\ell - \tau_s)^\gamma)$. The two parameters are now highly correlated, although this correlation gets less tight for larger parameter values. Nevertheless, the noisy coding model also outperforms a purely subadditive model (ELPD diff: -48.9; SE diff: 3.5). The reason for this is that the subadditive model has a dedicated parameter that *only* accounts for subadditivity, and thus risks overfitting noisy data patterns. The time-discriminability parameter β emerging from the noisy coding model, on the other hand, accounts for delay-dependence as well as being related to decision noise ν and contributing to the explanation of present bias through its effect on α. These interdependencies create tight constraints on estimated parameters, which in turn provide a closer

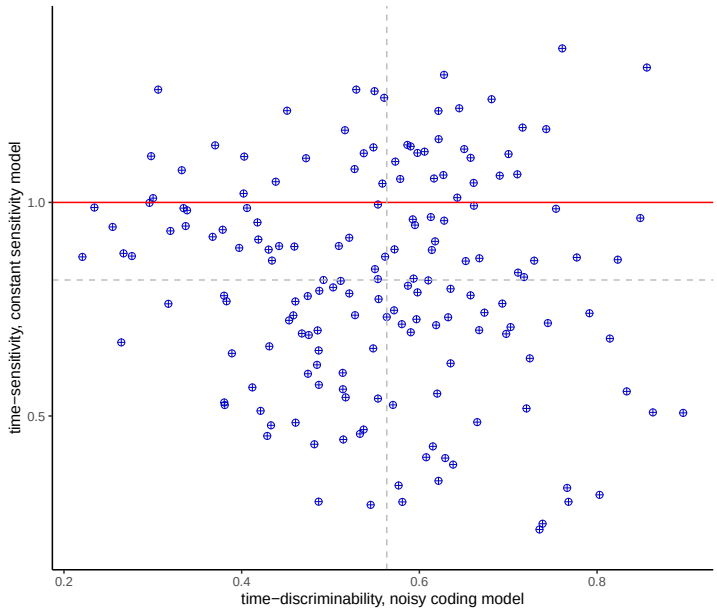

(a) time sensitivity in CS vs NCT models

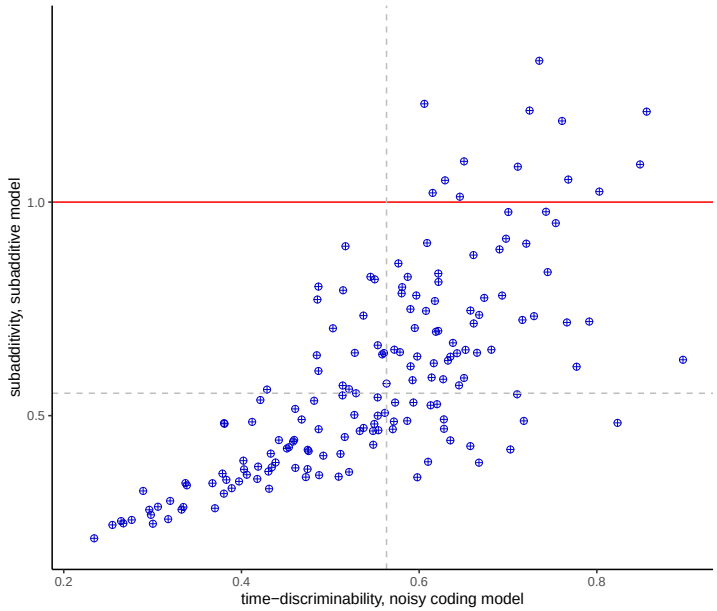

(b) time sensitivity in SA vs NCT model

Figure 5: Scatter plots of time sensitivity parameters from different models

The plotted parameters originate from a Bayesian hierarchical estimation of the individual-level parameters. Panel 5(a) plots the time-discriminability parameter estimated from the noisy coding model against the time-sensitivity parameter $\widehat{\beta}$ estimated from the constant sensitivity model, where $\delta_{s,\ell} = exp(-\widehat{\alpha}(\ell^{\widehat{\beta}} - s^{\widehat{\beta}}))$, with the 'hat' distinguishing the parameters from those in the noisy coding model. The estimates based on the constant sensitivity model fall into a much wider range than those from the noisy coding model, as well as being higher on average, highlighting their double role of fitting delay-dependence *and* strongly decreasing impatience. Panel 5(b) plots the noisy coding time-discriminability parameter against the time-compression parameter γ from the subadditive model, where $\delta_{\tau_s,\tau_\ell} = exp(-\widehat{\alpha}(\tau_\ell - \tau_s)^\gamma)$. The two parameters now show a much closer correlation, which however becomes less tight as γ gets large.

fit to the data patterns and enhanced predictive performance.

4 Discussion

The model I have presented predicts stochastic choice behaviour in tradeoffs between smaller-sooner and larger-later amounts from the noisy coding of time delays. A related recent literature in psychology and economics has studied the effects of randomness in preferences on inter-temporal discounting (Lu and Saito, 2018; He et al., 2019). These papers have highlighted that particular patterns of randomness in responses or preferences may result in hyperbolic as-if discounting, thus resulting in predictions similar to those obtained from the observation that decisions may be subject to 'cognitive uncertainty' (Enke and Graeber, 2021). Procedurally, these models have commonalities with traditional modelling approaches: they start from postulating an arbitrary error structure, and then investigate the consequences of this error structure for behaviour. The noisy coding model, on the other hand, derives predictions on both choice parameters and choice probabilities from the noisy coding of the delay stimuli. While the noisy coding model predicts delay-dependence and present-bias based on different mechanisms, these other models account for decreasing impatience by postulating randomness in preferences, and do not contain a separate mechanism accounting for delay-dependence in discounting, which I have found to be a central feature of the data.

The noisy coding model shares common foundations with recent models applied to decision making under risk and uncertainty. In particular, Khaw, Li and Woodford (2020) showed that a model in which outcomes are subject to noisy mental coding and subsequent Bayesian decoding can account for small stake risk aversion, which may otherwise be hard to reconcile with reasonable levels of risk aversion over larger stakes. In Vieider (2021), I showed that if probabilities are noisily encoded as well as outcomes, then a noisy coding setting will naturally generate prospect-theory like patterns, as well as predicting violations of prospect theory. While still in its infancy, this approach thus appears to hold the promise of a unification of descriptive approaches to individual decision-making.

Using a formally related Bayesian setup, Gabaix and Laibson (2017) derived a prediction of stringly decreasing impatience. The latter propose a model whereby future *utilities* are perceived with noise. If this noise in the perception of utilities is linearly increasing in the time delay attached to the utility, then the combination of these noisy perceptions with a prior results in posterior expectations taking the form of the hyperbolic discounting function proposed by Mazur (1987). Their modelling setup thus formalizes the intuition that hyperbolic discounting patterns may derive from difficulties people have imagining their future selfs. Their prediction of strongly decreasing impatience may well be complementary to the predictions derived here, and given the formal similarity, the two models would be straightforward to combine. In particular, increased difficulty in simulating non-monetary outcomes may further contribute to explaining differences in discounting patterns observed for money and consumption.

An unresolved question concerns why strongly decreasing impatience is sometimes observed in experiments using monetary payoffs. A recent literature has attributed seemingly hyperbolic patterns to the inherent uncertainty of the future (Halevy, 2008; Epper, Fehr-Duda and Bruhin, 2011; Epper and Fehr-Duda, 2018). If a decision maker has some doubts about a future payout obtaining, and if this decision maker exhibits the certainty effect, then they will exihibit seemingly hyperbolic patterns of discounting. I have been particularly careful to eliminate future uncertainty by having the experiment conducted by a trusted person (the lecturer) and in an environment where inter-actions are repeated an ongoing. That being said, future uncertainty could be easily incorporated into the model using the model of Vieider (2021), where I use a formally identical setup to model likelihood-distortions.

Decision making under risk and over time have been depicted as closely related (Prelec and Loewenstein, 1991). A particular relation between the two approaches is between probability distortions (probabilistic insensitivity) and time distortions (time insensitivity). In the noisy coding of time model I presented, time insensitivity can be traced back to noise in the encoding of time delays. In a twin paper

on decision making under uncertainty (Vieider, 2021), I show that likelihood-sensitivity can be traced back to noise in the coding of likelihood ratios in binary wagers using a formally equivalent setup. This is consistent with the results of Epper et al. (2011), who documented a correlation between the two phenomena, albeit based on the idea that the causation runs directly from probability distortions to time distortions. Disentangling these two different accounts of causality may constitute an interesting question for future research.

The predictions and results here reported also matter for policy and welfare considerations. For one, non-stationary preferences become irrelevant for welfare considerations once one incorporates their origin in the noisy perception of delays. Since such choice patterns are purely due to the noisy perception of time, a welfare-maximizing policy maker can safely ignore them. What is more, understanding the underlying mechanisms generating inter-temporal choice behaviour ought to allow one to leverage these mechanisms in policy design. In particular, the malleability of the choice patterns emerging from the model ought to make it possible to directly manipulate choice patterns by means of representational interventions. This opens the door for directly manipulating apparent preferences in policy interventions, which can thus be added to the behavioural policy maker's tool kit.

References

Abdellaoui, Mohammed, Han Bleichrodt, Olivier L'Haridon, and Corina Paraschiv (2013) 'Is There One Unifying Concept of Utility? An Experimental Comparison of Utility under Risk and Utility over Time.' *Management Science* 59(9), 2153–2169

Ainslie, George (1975) 'Specious reward: A behavioral theory of impulsiveness and impulse control.' *Psychological Bulletin* 82(4), 463–496

Attema, Arthur E, Han Bleichrodt, Kirsten IM Rohde, and Peter P Wakker (2010) 'Time-tradeoff sequences for analyzing discounting and time inconsistency.' *Management Science* 56(11), 2015–2030

Baillon, Aurélien, Han Bleichrodt, and Vitalie Spinu (2020) 'Searching for the reference point.' *Management Science* 66(1), 93–112

Balakrishnan, Uttara, Johannes Haushofer, and Pamela Jakiela (2020) 'How soon is now? evidence of present bias from convex time budget experiments.' *Experimental Economics* 23(2), 294–321

Bettinger, Eric, and Robert Slonim (2007) 'Patience among children.' *Journal of Public Economics* 91(1-2), 343–363

Bleichrodt, Han, Kirsten I. M. Rohde, and Peter P. Wakker (2009) 'Non-hyperbolic time inconsistency.' *Games and Economic Behavior* 66(1), 27–38

Bleichrodt, Han, Yu Gao, and Kirsten IM Rohde (2016) 'A measurement of decreasing impatience for health and money.' *Journal of Risk and Uncertainty* 52(3), 213–231

Bruhin, Adrian, Helga Fehr-Duda, and Thomas Epper (2010) 'Risk and Rationality: Uncovering Heterogeneity in Probability Distortion.' *Econometrica* 78(4), 1375–1412

Cantlon, Jessica F (2012) 'Math, monkeys, and the developing brain.' *Proceedings of the National Academy of Sciences* 109(Supplement 1), 10725–10732

Cappelletti, Marinella, Erica Gessaroli, Rosalyn Hithersay, Micaela Mitolo, Daniele Didino, Ryota Kanai, Roi Cohen Kadosh, and Vincent Walsh (2013) 'Transfer of cognitive training across magnitude dimensions achieved with concurrent brain stimulation of the parietal lobe.' *Journal of Neuroscience* 33(37), 14899–14907

Carpenter, Bob, Andrew Gelman, Matthew D Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell (2017) 'Stan: A probabilistic programming language.' *Journal of Statistical Software* 76(1), 1–32

Cavagnaro, Daniel R, Gabriel J Aranovich, Samuel M McClure, Mark A Pitt, and Jay I Myung (2016) 'On the functional form of temporal discounting: An optimized adaptive test.' *Journal of Risk and Uncertainty* 52(3), 233–254

Cohen, Jonathan, Keith Marzilli Ericson, David Laibson, and John Myles White

(2020) 'Measuring time preferences.' *Journal of Economic Literature* 58(2), 299–347

Conte, Anna, John D. Hey, and Peter G. Moffatt (2011) 'Mixture models of choice under risk.' *Journal of Econometrics* 162(1), 79–88

Cubitt, Robin P, and Daniel Read (2007) 'Can intertemporal choice experiments elicit time preferences for consumption?' *Experimental Economics* 10(4), 369–389

Dayan, Peter, and Laurence F Abbott (2001) *Theoretical neuroscience: computational and mathematical modeling of neural systems* (Computational Neuroscience Series)

Dehaene, Stanislas (2003) 'The neural basis of the weber–fechner law: a logarithmic mental number line.' *Trends in cognitive sciences* 7(4), 145–147

_ (2011) *The number sense: How the mind creates mathematics* (OUP USA)

Dehaene, Stanislas, and Jean-Pierre Changeux (1993) 'Development of elementary numerical abilities: A neuronal model.' *Journal of cognitive neuroscience* 5(4), 390–407

Doepke, Matthias, and Fabrizio Zilibotti (2014) 'Culture, Entrepreneurship, and Growth.' In 'Handbook of Economic Growth,' vol. 2

_ (2017) 'Parenting with style: Altruism and paternalism in intergenerational preference transmission.' *Econometrica* 85(5), 1331–1371

Dohmen, Thomas, Armin Falk, David Huffman, and Uwe Sunde (2017) 'The robustness and pervasiveness of sub-additivity in intertemporal choice.' In 'Working Paper'

Doya, Kenji, Shin Ishii, Alexandre Pouget, and Rajesh PN Rao (2007) *Bayesian brain: Probabilistic approaches to neural coding* (MIT press)

Ebert, Jane E. J., and Drazen Prelec (2007) 'The Fragility of Time: Time-Insensitivity and Valuation of the Near and Far Future.' *Management Science* 53(9), 1423–1438

Enke, Benjamin, and Thomas Graeber (2019) 'Cognitive uncertainty.' Technical Report, National Bureau of Economic Research

_ (2021) 'Cognitive uncertainty in intertemporal choice.' Technical Report, Mimeo

Epper, Thomas, and Helga Fehr-Duda (2018) 'Risk in time: The intertwined nature of risk taking and time discounting.' *Working Paper*

Epper, Thomas, Helga Fehr-Duda, and Adrian Bruhin (2011) 'Viewing the future through a warped lens: Why uncertainty generates hyperbolic discounting.' *Journal of Risk and Uncertainty* 43(3), 169–203

Fechner, Gustav Theodor (1860) (Kessinger's Legacy Reprints)

Frederick, Shane, George Loewenstein, and Ted O'Donoghue (2002) 'Time Discounting and Time Preference: A Critical Review.' *Journal of Economic Literature* 40(2), 351–401

Friston, Karl (2005) 'A theory of cortical responses.' *Philosophical transactions of the Royal Society B: Biological sciences* 360(1456), 815–836

_ (2010) 'The free-energy principle: a unified brain theory?' *Nature Reviews Neuroscience* 11(2), 127–138

Friston, Karl, Thomas FitzGerald, Francesco Rigoli, Philipp Schwartenbeck, and Giovanni Pezzulo (2017) 'Active inference: a process theory.' *Neural Computation* 29(1), 1–49

Frydman, Cary, and Lawrence J Jin (2022) 'Efficient coding and risky choice.' *Quarterly Journal of Economics* 136, 161–213

Gabaix, Xavier, and David Laibson (2017) 'Myopia and discounting.' Technical Report, National bureau of economic research

Gallistel, Charles R, and Rochel Gelman (2000) 'Non-verbal numerical cognition: From reals to integers.' *Trends in cognitive sciences* 4(2), 59–65

Gelman, Andrew, Jessica Hwang, and Aki Vehtari (2014) 'Understanding predictive information criteria for bayesian models.' *Statistics and computing* 24(6), 997–1016

Gershman, Samuel J, and Rahul Bhui (2020) 'Rationally inattentive intertemporal choice.' *Nature communications* 11(1), 1–8

Gold, Joshua I, and Michael N Shadlen (2001) 'Neural computations that underlie decisions about sensory stimuli.' *Trends in cognitive sciences* 5(1), 10–16

Halberda, Justin, Michèle MM Mazzocco, and Lisa Feigenson (2008) 'Individual differences in non-verbal number acuity correlate with maths achievement.' *Nature* 455(7213), 665–668

Halevy, Yoram (2008) 'Strotz Meets Allais: Diminishing Impatience and the Certainty Effect.' *The American Economic Review* 98(3), 1145–1162

He, Lisheng, Russell Golman, and Sudeep Bhatia (2019) 'Variable time preference.' *Cognitive Psychology* 111, 53–79

Hey, John D., and Chris Orme (1994) 'Investigating Generalizations of Expected Utility Theory Using Experimental Data.' *Econometrica* 62(6), 1291–1326

Howard, Marc W, and Karthik H Shankar (2018) 'Neural scaling laws for an uncertain world.' *Psychological review* 125(1), 47

Howard, Marc W, Christopher J MacDonald, Zoran Tiganj, Karthik H Shankar, Qian Du, Michael E Hasselmo, and Howard Eichenbaum (2014) 'A unified mathematical framework for coding time, space, and sequences in the hippocampal region.' *Journal of Neuroscience* 34(13), 4692–4707

Imai, Taisuke, Tom A Rutter, and Colin F Camerer (2021) 'Meta-analysis of present-bias estimation using convex time budgets.' *The Economic Journal* 131(636), 1788–1814

Kable, Joseph W, and Paul W Glimcher (2010) 'An ?as soon as possible? effect in human intertemporal decision making: behavioral evidence and neural mechanisms.' *Journal of neurophysiology* 103(5), 2513–2531

Khaw, Mel Win, Ziang Li, and Michael Woodford (2020) 'Cognitive imprecision and small-stakes risk aversion.' *The Review of Economic Studies, forthcoming*

— (2021) 'Cognitive imprecision and small-stakes risk aversion.' *The Review of Economic Studies* 88(4), 1979–2013

Kirby, Kris N (1997) 'Bidding on the future: Evidence against normative discounting of delayed rewards.' *Journal of Experimental Psychology: General* 126(1), 54

Knill, David C, and Alexandre Pouget (2004) 'The bayesian brain: the role of uncertainty in neural coding and computation.' *TRENDS in Neurosciences* 27(12), 712–719

Laibson, David (1997) 'Golden Eggs and Hyperbolic Discounting.' *The Quarterly Journal of Economics* 112(2), 443–478

Loewenstein, George, and Drazen Prelec (1992) 'Anomalies in Intertemporal Choice: Evidence and an Interpretation.' *The Quarterly Journal of Economics* 107(2), 573–597

Lu, Jay, and Kota Saito (2018) 'Random intertemporal choice.' *Journal of Economic Theory* 177, 780–815

Mazur, James E (1987) 'An adjusting procedure for studying delayed reinforcement.' *Quantitative analyses of behavior* 5, 55–73

McElreath, Richard (2016) *Statistical Rethinking: A Bayesian Course with Examples in R and Stan* (Academic Press)

Mischel, Walter, and Ebbe B Ebbesen (1970) 'Attention in delay of gratification.' *Journal of Personality and Social Psychology* 16(2), 329

Mischel, Walter, Yuichi Shoda, and Monica L Rodriguez (1989) 'Delay of gratification in children.' *Science* 244(4907), 933–938

Moyer, Robert S, and Thomas K Landauer (1967) 'Time required for judgements of numerical inequality.' *Nature* 215(5109), 1519–1520

Netzer, Nick (2009) 'Evolution of time preferences and attitudes toward risk.' *American Economic Review* 99(3), 937–55

Nieder, Andreas, and Earl K Miller (2003) 'Coding of cognitive magnitude: Compressed scaling of numerical information in the primate prefrontal cortex.' *Neuron* 37(1), 149–157

Oprea, Ryan (2022) 'Simplicity equivalents.' *Working Paper*

Park, Joonkoo, and Elizabeth M Brannon (2013) 'Training the approximate number system improves math proficiency.' *Psychological science* 24(10), 2013–2019

Petzschner, Frederike H, and Stefan Glasauer (2011) 'Iterative bayesian estimation as an explanation for range and regression effects: a study on human path integration.' *Journal of Neuroscience* 31(47), 17220–17229

Phelps, E. S., and R. A. Pollak (1968) 'On Second-Best National Saving and Game-Equilibrium Growth.' *The Review of Economic Studies* 35(2), 185–199

Piazza, Manuela, Philippe Pinel, Denis Le Bihan, and Stanislas Dehaene (2007) 'A magnitude code common to numerosities and number symbols in human intraparietal cortex.' *Neuron* 53(2), 293–305

Prelec, Drazen, and George Loewenstein (1991) 'Decision Making over Time and Under Uncertainty: A Common Approach.' *Management Science* 37(7), 770–786

Rao, Rajesh PN, and Dana H Ballard (1999) 'Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects.' *Nature neuroscience* 2(1), 79–87

Read, Daniel (2001) 'Is Time-Discounting Hyperbolic or Subadditive?' *Journal of Risk and Uncertainty* 23(1), 5–32

Read, Daniel, and Peter HMP Roelofsma (2003) 'Subadditive versus hyperbolic discounting: A comparison of choice and matching.' *Organizational behavior and human decision processes* 91(2), 140–153

Read, Daniel, Shane Frederick, Burcu Orsel, and Juwaria Rahman (2005) 'Four score and seven years from now: The date/delay effect in temporal discounting.' *Management Science* 51(9), 1326–1335

Regenwetter, Michel, Daniel R Cavagnaro, Anna Popova, Ying Guo, Chris Zwilling, Shiau Hong Lim, and Jeffrey R Stevens (2018) 'Heterogeneity and parsimony in intertemporal choice.' *Decision* 5(2), 63

Robson, Arthur J (2001) 'The biological basis of economic behavior.' *Journal of Economic Literature* 39(1), 11–33

Rohde, Kirsten IM (2019) 'Measuring decreasing and increasing impatience.' *Management Science* 65(4), 1700–1716

Samuelson, Paul A. (1937) 'A Note on Measurement of Utility.' *The Review of Economic Studies* 4(2), 155–161

Stan Development Team (2017) 'RStan: the R interface to Stan.' R package version 2.17.2

Thaler, Richard (1981) 'Some empirical evidence on dynamic inconsistency.' *Economics Letters* 8(3), 201–207

Vehtari, Aki, Andrew Gelman, and Jonah Gabry (2017) 'Practical bayesian model evaluation using leave-one-out cross-validation and waic.' *Statistics and computing* 27(5), 1413–1432

Vieider, Ferdinand M. (2021) 'Decisions under uncertainty as bayesian inference.' Working Paper WP 21/1022 of Faculty of Economics and Business Administration, Ghent University, Belgium 21/1022, Ghent University, Faculty of Economics and Business Administration

Vilares, Iris, and Konrad Kording (2011) 'Bayesian models: the structure of the world, uncertainty, behavior, and the brain.' *Annals of the New York Academy of Sciences* 1224(1), 22

von Gaudecker, Hans-Martin, Arthur van Soest, and Erik Wengström (2011) 'Heterogeneity in Risky Choice Behaviour in a Broad Population.' *American Economic Review* 101(2), 664–694

Walsh, Vincent (2003) 'A theory of magnitude: common cortical metrics of time, space and quantity.' *Trends in cognitive sciences* 7(11), 483–488

Whalen, John, Charles R Gallistel, and Rochel Gelman (1999) 'Nonverbal counting in humans: The psychophysics of number representation.' *Psychological science* 10(2), 130–137

Zauberman, Gal, B Kyu Kim, Selin A Malkoc, and James R Bettman (2009) 'Discounting time and time discounting: Subjective time perception and intertemporal preferences.' *Journal of Marketing Research* 46(4), 543–556

ONLINE APPENDIX

A Model derivation for delays from the present

To obtain actionable choice quantities, we need to combine the likelihood and prior by Bayesian updating. Since r is a scalar, the posterior distribution of the time delay will be:

$$P\left[ln(t)|r\right] \sim \mathcal{N}\left(\frac{\sigma^2}{\sigma^2+\nu^2} \times r + \frac{\nu^2}{\sigma^2+\nu^2} \times \mu \ , \ \frac{\nu^2\sigma^2}{\nu^2+\sigma^2}\right), \tag{22}$$

where $\beta \triangleq \frac{\sigma^2}{\sigma^2+\nu^2}$, so that $(1-\beta) \triangleq \frac{\nu^2}{\sigma^2+\nu^2}$. The parameter β thus gives us the weight attributed to the likelihood relative to the prior mean, μ. It depends on the relative uncertainty associated to the mental signal versus the mental prior. Further defining $\alpha \triangleq e^{(1-\beta)\mu}$ gives us the expectation of the posterior, $E[ln(t)|r]$. The threshold equation obtains by plugging $E[ln(t)|r] = \beta r + ln(\alpha)$, back into the mental choice rule in equation 2 and solving for r. We then exploit the known distributional properties of the mental signal r. Taking the z-score of the likelihood by subtracting the mean and dividing by the standard deviation, we obtain $\mathcal{Z} = \frac{r-ln(t)}{\nu}$. We can the compare this z-score to an equivalent z-score obtained from the threshold equation, $\mathcal{Z}_t = \frac{r-\beta^{-1}\left[ln(-ln(\frac{y}{x}))-ln(\alpha)\right]}{\nu}$. Subtracting one from the other gives us the probabilistic choice rule in equation 7.

B Alternative choice rule in log space

In my preferred specification, I have assumed that the choice rule is transformed from its original scale by taking the log twice before entering the mental transformation process. Here I discuss the implications of using the once-logged choice rule instead (not logging the choice rule at all results in equivalent conclusions). Starting from the optimal choice rule, we take the log once and rearrange, to obtain $\tau < -ln\left(\frac{y}{x}\right)$. The likelihood and prior for the time delay remain the same as in the main text. However, the expectation of the posterior will be different,

since it is obtained from a log-normal distribution: $E[t|r] = e^{\beta r + (1-\beta)(\mu + \frac{1}{2}\sigma^2)}$. Let $\widehat{\alpha} \triangleq e^{(1-\beta)(\mu+\frac{1}{2}\sigma^2)}$. Substituting the posterior expectation back into the optimal choice rule, we obtain $e^{\beta r + ln(\widehat{\alpha})} < -ln\left(\frac{y}{x}\right)$. Taking the logarithm of both sides and rearranging, we obtain the threshold equation $r < \beta^{-1}\left[ln\left(-ln\left(\frac{y}{x}\right)\right) - ln(\widehat{\alpha})\right]$. We can now proceed as above to obtain the actionable choice rule. Notice how the only difference from the specification in the main text is in the definition of $\widehat{\alpha}$. The discounting parameter is thus influenced not only by the mean of the prior, but also by its variance. The *qualitative* predictions of the model remain unaltered.

C Derivation with up-front delays

Take the log-log choice rule $ln(\tau_\ell - \tau_s) < ln\left[-ln\left(\frac{y}{x}\right)\right] - ln(\rho)$. Assume now that the mental signal directly concerns the time delay $\ell - s = \tau_\ell - \tau_s$, giving the likelihood $r_{\ell-s} \sim \mathcal{N}(ln(\ell - s), \nu^2)$ and prior $ln(\ell - s) \sim \mathcal{N}(ln(\xi), \sigma^2)$. Combining likelihood and prior by Bayesian updating, and substituting the expectation of the posterior back into the choice rule yields $ln(\alpha) + \beta \times ln(\ell - s) < ln\left[-ln\left(\frac{y}{x}\right)\right]$. After exponentiating twice, we obtain the mental choice rule $exp(-\alpha(\ell - s)^\beta) > \frac{y}{x}$, which abstracts from decision noise for simplicity. Now take the case where $\ell - s = s$. The function then simplifies to $exp(-\alpha(s)^\beta)$, i.e. it becomes identical to the function one obtains for a delay τ_s from the present, except for the added 1 producing present bias. Increasing the up-front delay to $2s$, $3s$, or indeed ns while keeping the time interval constant has no impact. In other words, the up-front delay is meaningless for this setup except when immediate payouts are involved, and the function is thus stationary.

We can compare the setup just derived to the setup we get when comparing two delays from the present. In that framework, the choice between the same two options would be written as $exp(-\alpha\ell^\beta)x > exp(-\alpha s^\beta)y$, or equivalently $exp(-\alpha(\ell^\beta - s^\beta)) > \frac{y}{x}$. The latter more closely corresponds to the standard setup in (Ebert and Prelec, 2007), with hyperbolicity deriving from the separate transformation of the two time delays. Setting $\ell - s = s$ does not make the function identical to the one we obtain for a delay s from the present. At the same

time, however, this feature does not result in separate predictions about strongly decreasing impatience and delay-dependence. That is, it is not a good *predictive* function in the sense that it cannot separately account for these two potentially important phenomena, instead using one and the same parameter to fit both.

Decomposing the log of the time difference. One may think that a prediction of strongly decreasing impatience could be obtained by separately transforming the two time delays within the logarithm in $ln(\tau_\ell - \tau_s) < ln\left[-ln\left(\frac{y}{x}\right)\right] - ln(\rho)$. As it turns out, this approach has no closed-form solution. Let us define a function $h(k) = ln(e^k - 1)$. We can then decompose the difference within the log as $ln(s) + h[ln(\ell) - ln(s)] < ln\left[-ln\left(\frac{y}{x}\right)\right]$. We can now distribute the logs of the single time delays. The issue, however, is that the function h makes it impossible to *jointly* distribute the noisy mental signals for $ln(x)$ and $ln(y)$. This, in turn, means that we can derive no analytic solution for the error term. Note that this issue is not due to use of the twice-logged choice rule, but generalizes to a setup starting from a once-logged choice rule. Indeed, in that case the same probelm will occur during a subsequent stage of the derivation, which exponentiating to obtain the posterior expectation of the unlogged time delay (cfr. above).

D A choice rule for generalized hyperbolicity

The failure to predict strongly decreasing impatience independently of delay length begs the question of whether a different mental representation of the choice stimuli *could* yield a prediction of such behaviour. A candidate choice rule that could produce such behaviour is $exp\left(-\rho \times \frac{\tau_\ell}{\tau_s}\right) > \frac{x-y}{x}$. This choice rule captures the intuition that the relative increase in waiting time is traded off against the relative gain of taking the later amount instead of the sooner, so that we refer to it as *relative delay discounting.* In log-log space, the choice rule takes the form $ln(\tau_\ell) - ln(\tau_s) < ln\left[-ln\left(\frac{x-y}{x}\right)\right]$, which allows us to jointly distribute the mental signals for the two delays τ_s and τ_ℓ. Encoding and decoding the two time delays results in the posterior expectations $E[ln(s)|r_s] = \beta r_s + (1-\beta)\mu_s$ and $E[ln(\ell)|r_\ell] =$

$\beta r_\ell + (1-\beta)\mu_d$ following the same derivations as above. Substituting everything into the choice rule and solving for the mental signals, we obtain the threshold equation $r_d - r_s < \beta^{-1}\left[ln\left[-ln\left(\frac{x-y}{x}\right)\right] - ln(\alpha)\right]$. Further let $\mu_d = ln(\lambda_d)$ and $\mu_s = ln(\lambda_s)$. Then $\alpha \triangleq \left(\frac{\lambda_d}{\lambda_s}\right)^{1-\beta}$. Using the joint distribution of the two mental signals, which will itself be normal, obtaining the z-score, comparing this z-score to the threshold, and rearranging, we obtain the following probabilistic choice rule:

$$P[(x,\ell) \succ (y,s)] = \Phi\left[\frac{\beta^{-1}\left[ln\left[-ln\left(\frac{x-y}{x}\right)\right] - ln(\alpha)\right] - [ln(\ell) - ln(s)]}{\sqrt{2}\nu}\right], \quad (23)$$

where the standard deviation is now $\sqrt{2}\nu$ due to the joint distribution of two independent mental signals. Abstracting from the decision noise to zoom in on the indifference condition underlying the nominator, and exponentiating twice to take the expression back to the original scale, we obtain $exp\left[-\alpha\left(\frac{\ell}{s}\right)^\beta\right] = \frac{x-y}{x}$, where α may be different from the time preference parameter α discussed in the main text, depending on the prior expectation of the later versus the sooner delay. This function now is clearly hyperbolic, as it explicitly depends on the up-front delay τ_s. The choice rule, however, may not be very intuitive. Indeed, the model performs not only worse than the noisy coding mode derived in the text (ELPD diff: 5181.2, SE diff: 87.4), but also worse than the standard constant sensitivity model (ELPD diff: 4916.3, SE diff: 81.2), than the subadditive model (ELPD diff: 5132.3, 87.3), and even than the quasi-hyperbolic (ELPD diff: 4916.2, SE diff: 81.2) and exponential models (ELPD diff: 4801.3, SE diff: 80.9). Given that the latter models fail to predict some of he major patterns present in the data, such as delay-dependence, the verdict on this choice rule is indeed damning.

E Generalization to include outcome distortions

I have assumed outcomes to be perceived without noise. If outcomes are also subject to noisy coding, the two dimensions will need to be derived jointly to accurately represent the decision-making process. Starting from the linear in log-log optimal choice rule, let $z \triangleq -ln(\frac{y}{x})$. Assuming that the log of the ratio of the

smaller to larger payment is encoded in a way similar to time delays, we obtain the following likelihood and prior: $r_z \sim \mathcal{N}(ln(z), \nu_o^2)$, $ln(z) \sim \mathcal{N}(\mu_o, \sigma_o^2)$. Using the usual procedure and after defining $\gamma \triangleq \frac{\sigma_o^2}{\sigma_o^2+\nu_o^2}$, we obtain the posterior expectation $E[ln(z)|r_z] = \gamma r_z + (1-\gamma)\mu_o$. Substituting this expression together with the one for the time delay into the choice rule, we obtain the threshold equation $\beta r - \gamma r_z < -ln(\theta)$, where $\theta \triangleq \frac{\alpha}{\zeta}$ and $\zeta \triangleq e^{(1-\gamma)\mu_o}$. This parameter is a composite of the time prior and outcome prior means. In particular, it captures the intuition that an expectation that outcomes will increase over time, as captured by ζ, will lower the observed discount rate. Distributing the signals on the left-hand side jointly, obtaining the z-score, and comparing it to the threshold z-score, we obtain the following probabilistic choice rule:

$$Pr[(x,\ell) \succ (y,0)] = \Phi\left(\frac{ln(\theta) + \beta \times ln(t) - \gamma \times ln\left[-ln\left(\frac{y}{x}\right)\right]}{\sqrt{\nu^2\beta^2 + \nu_o^2\gamma^2}}\right). \tag{24}$$

Zooming in on the indifference condition, exponentiating twice, and letting $\delta_\tau \triangleq \frac{y}{x}$, we obtain $\delta_\tau = exp(-\theta t^{\frac{\beta}{\gamma}})$. The result of introducing outcome distortions is thus two-fold: 1) it rescales decision noise; and 2) it modifies the time distortion parameter, which is now given by $\frac{\beta}{\gamma}$ instead of β. Notice, however, that—other than in a standard setup (Abdellaoui, Bleichrodt, L'Haridon and Paraschiv, 2013)—the two parameters β and γ cannot be separately identified in empirical data. To see this, take a preference condition $(x, \tau_\ell; v, \tau_s) \succ (y, \tau_s)$. In the present setup we obtain $\left(\frac{y-v}{x}\right)^\gamma$, so that the equation is only identified up to a power. The conclusion that outcome distortions would only affect the value of β, but have no separate meaning, may then be seen as an additional prediction emerging from the model.

F Additional results

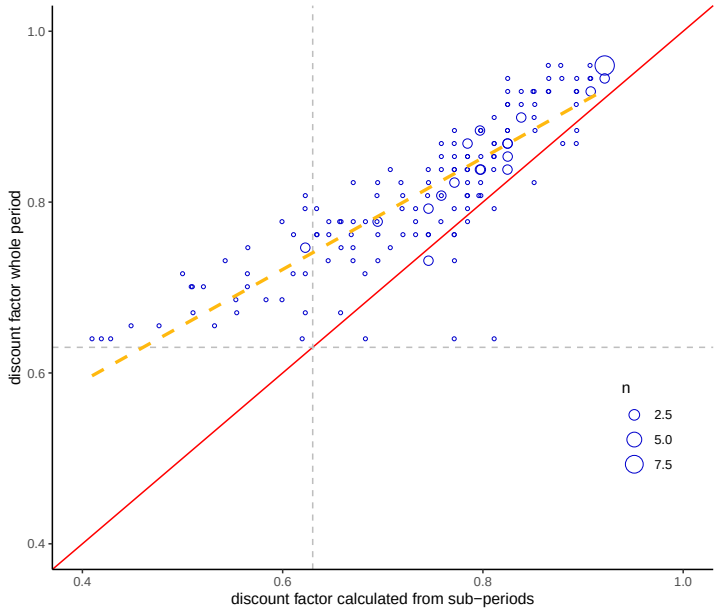

(a) subadditivity, 12 weeks from 12 weeks

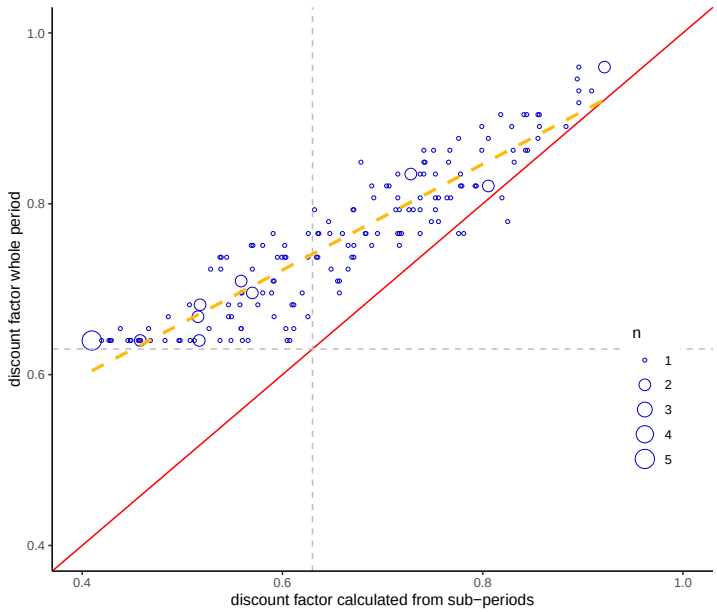

(b) subadditivity, 24 weeks vs 12-week delays

Figure 6: Non-parametric illustration of subadditivity—additonal comparisons

The subadditivity comparisons are obtained by comparing a discount factor over the whole period with the product of the discount factors of the subperiods. The pattern in panel 6(a) thus obtains from a comparison of the discount factor $\delta_{12,24}$ with the product of the two discount factors $\delta_{12,18} \times \delta_{18,24}$, where the subscripted numbers indicate the extremes of the time delays. The pattern in panel 6(b) obtains from a comparison of the discount factor $\delta_{0,24}$ to the product of all underlying 12-week discount factors, $\delta_{0,12} \times \delta_{12,24}$. Dashed lines indicate correlations.

G Experimental Instructions

Thank you for taking part in this experiment. You will be asked to take some decisions involving time delays. On each screen, you will be asked to choose between an amount of money that is paid at **a sooner moment in time**, and an amount that is paid **at a later moment in time**. Time delays are always indicated in weeks from today. Here is an example of a choice task:

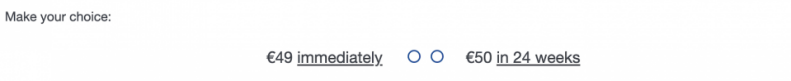

You will be presented repeatedly with such tasks, and you are asked to indicate your choice for each one of those tasks. The experiment will take approximately 20 minutes. At the end of this class, **10 participants will be randomly drawn to play one of their choice for real money**. Notice that both the amounts and the time delays involved may change from screen to screen. Please consider the information carefully and choose your preferred option.

At the end of the experiment, **the lecturer will announce the students selected to play for one of their choices**. If you have chosen the **immediate amount** in the randomly selected choice, we will transfer the corresponding amount to your bank account immediately while you are waiting. Since transfers amongst all major banks are immediate, you will be asked to check your bank account and **confirm the receipt of the money before leaving**.

If you have chosen a **delayed amount** in the randomly selected choice, then that amount will be paid to you on the indicated day. For instance, an amount indicated as obtaining **in 4 weeks** will be paid 4 weeks from today, i.e. on Monday November 15th. The lecturer will make a note of your bank account to organize the transfer. You will receive **a certificate signed by the lecturer to guarantee the transfer**. The certificate will indicate the amount to be paid and the date of transfer. It will also contain the contact details of the lecturer, for the case that you change bank account or have any questions concerning the transfer.