Cognitive Foundations of Delay-Discounting*

Ferdinand M. Vieider†[1,2]

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

December 15, 2023

Abstract

I propose and test a model whereby well-known empirical deviations from exponential discounting emerge from the optimal Bayesian combination of noisily perceived time delays with a learned prior summarizing the statistics of the environment. This results in a generative account of delay-discounting which reconciles behavioural patterns that have thus far been modelled mostly separately, and which cannot be captured by any *one* existing model. The model predictions obtain based on a set of assumptions that are plausible from a neuro-biological point of view, and which have been shown to organize empirical findings across different decision domains. I test the model in two pre-registered experiments by manipulating the key mechanisms underlying its behavioural predictions—perceptual noise, and the statistics of the environment. The results show a distinctive empirical fingerprint that cannot be organized by existing models of delay-discounting.

1 Motivation

Humans as well as animals have long been known to substantially discount future rewards. To use an example from a classroom experiment, an outcome received 6 weeks from now was considered only as good as 86% of that same outcome received immediately (point A in figure 1). This implies an annualized discount rate in

*I acknowledge financial support from the Research Foundation—Flanders (FWO) under the project "Causal Determinants of Preferences" (G008021N). The experiments have received ethics approval from the ethics committee at the Faculty of Economics and Business, Ghent University, with number *FEB 2021 M*. The Experiments have been pre-registered at OSF with reference *osf.io/6a2kx*. I am grateful to Ranoua Bouchouicha and Jilong Wu for help in executing the experiments. I am further indebted to Mohammed Abdellaoui, Miguel Ballester, Ranoua Bouchouicha, Soo Hong Chew, Benjamin Enke, Thomas Epper, Thomas Graeber, Levent Gumus, Olivier L'Haridon, Yuchi Li, Nick Netzer, Ryan Oprea, Pietro Ortoleva, Rafael Polania, Christian Ruff, Peter Wakker, and Leeat Yariv for helpful discussions. All errors remain my own.

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

excess of 300% assuming continuous compounding. This finding appears even more puzzling when one considers that it has been obtained using monetary payoffs. Substantial discounting for money is paradoxical from an economic point of view, since one would need to assume immediate consumption of any monetary payouts to explain it (Cohen, Ericson, Laibson and White, 2020). Cubitt and Read (2007) show that for discount rates exceeding the market borrowing rate—as is clearly the case here—no inferences can be drawn on time preferences at all if one assumes agents to abide by standard principles of economic rationality.

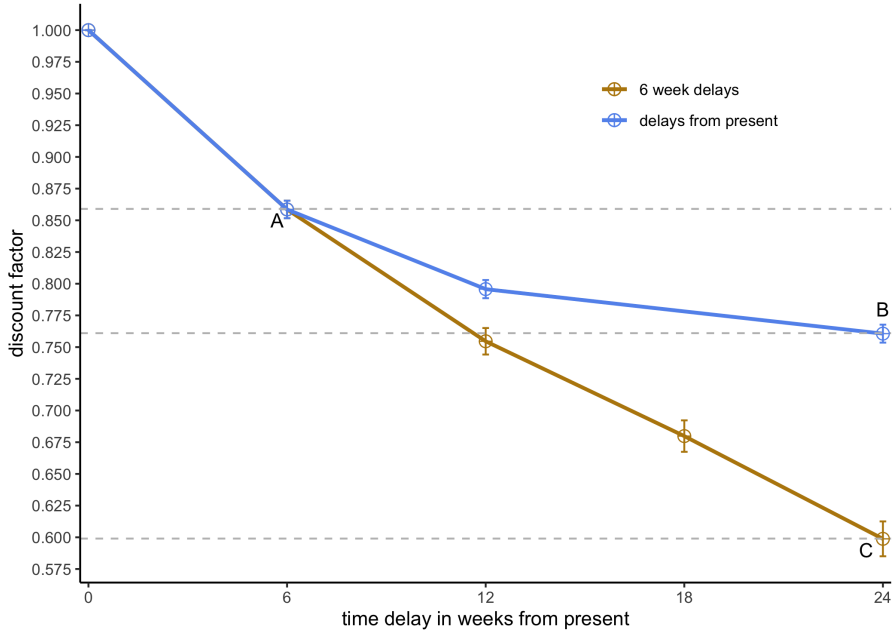

Figure 1: Nonparametric discount functions based on different time delays

Patterns shown are based on an incentivized classroom experiment with 175 students, using tradeoffs between smaller-sooner and larger-later rewards (see Online Appendix F). The figure shows discount functions obtained based on different types of choice stimuli. The function labelled 'delays from present' uses delays of 6, 12 and 24 weeks from the present. The '6 week delays' function is obtained by using delays of 6 weeks from the present, 6 weeks from 6 weeks, 6 weeks from 12 weeks, and 6 weeks from 18 weeks, so that the discount factor for a 12-week delay, δ_{12}, is calculated as $\delta_{12} = \delta_{0,6} \times \delta_{6,12}$, where the subscripts indicate the sooner and later time delays attached to the smaller and larger payoffs, respectively. Following the same procedure, we obtain $\delta_{18} = \delta_{0,6} \times \delta_{6,12} \times \delta_{12,18}$, and $\delta_{24} = \delta_{0,6} \times \delta_{6,12} \times \delta_{12,18} \times \delta_{18,24}$.

Choice patterns in inter-temporal tradeoffs further deviate systematically from the predictions of the standard model of exponential discounting (Samuelson, 1937). The curve leading up to point B in figure 1, obtained from delays of increasing length from the present, exhibits the iconic hyperbolic shape which has received much attention in the discounting literature. Compare this to the curve leading up to point C, which has been obtained by multiplying the discount factors of subsequent 6-

week delays from the present, from 6 weeks, from 12 weeks, and from 18 weeks. The coexistence of these two patterns challenges the majority of descriptive discounting models—including the whole class of "hyperbolic" models—which predicts the two curves to be *the same*. A further puzzle arises from the kink of the function leading up to point C at point A, which captures *present bias*—a disproportionate preference for the smaller-sooner option when it obtains immediately relative to when a common delay is added to both options.[1] This phenomenon is empirically strong (Rohde, 2019; Imai, Rutter and Camerer, 2021; Cheung, Tymula and Wang, 2023),[2] and difficult to subsume under generalized hyperbolic functions.[3]

The descriptive regularities discussed above have all been documented empirically and modelled formally in the literature. Yet few—if any—of the models and functions devised to describe discounting can account for all of these facts *jointly* (see review at the end of this section). Here, I ask a different albeit closely related set of questions: where might such distinctive regularities in inter-temporal choice originate? Could a model of the generative process reconcile the different types of behaviour that have been observed? Can this be done based on principles that are plausible from a neuro-biological perspective? And should we expect the same generative processes to underpin behaviour across different decision domains beyond delay-discounting? The last point is paramount for a neuro-biologically inspired approach, as we would expect the same principles of efficient information processing to leave a distinctive fingerprint across decision domains.

[1] Present bias naturally emerges from a wide class of weighted discount functions. For instance, Sozou (1998) showed how present-biased behaviour may emerge from uncertain hazard rates. Jackson and Yariv (2014) documented the emergence of present bias from the social aggregation of heterogenous preferences. Ebert, Wei and Zhou (2020) show that present bias can emerge based on a broad class of weighted functions, resulting e.g. from decision-makers' uncertainty about the correct discount rate to apply in a given situation.

[2] Online Appendix F presents additional results from the classroom experiment showing that the differences in choice proportions resulting in this kink are statistically significant and economically sizeable. See also experiment I below for further results.

[3] Impatience measured using delays of increasing length from the present, using delays of constant length pushed from the present into the future, and using delays of constant length pushed farther into the future from an initial upfront delay typically displays quantitatively and even qualitatively distinct patterns—see review of the literature at the end of this section, and the evidence from the experiments. Fitting these patterns with one and the same "hyperbolic" function will produce a poor empirical fit. Even more importantly, the *predictive* value of any fitted function for future behaviour will be severely limited by the fact that behaviour differs systematically across the different delay types.

Here, I present a model of delay-discounting based on the premise that people take decisions based on noisy mental representations of choice objects, and in particular, of time delays.[4] Noisy representations of time delays may arise from quick and approximate assessments of inter-temporal tradeoffs, or because time is inherently complicated to quantify objectively for a biological organism. I show that—if delays are perceived with some noise—then it will be optimal to rein in this noise by combination with a Bayesian prior enshrining the statistics of the environment, which needs to be learned based on the very same noisy perceptions. The combination of signal and prior will in turn result in systematic deviations from the exponential benchmark. In particular, it will produce both the apparently hyperbolic shape observed for delays from the present leading up to point B, and the steep discounting observed for shorter time delays leading up to point C. The setup furthermore predicts a disproportionate preference for the present as a separate behavioural phenomenon, thus providing cognitive microfoundations for quasi-hyperbolic discounting (Phelps and Pollak, 1968; Laibson, 1997).[5]

The model predictions arise from principles that are plausible from a neurobiological point of view. In particular, I show that the stylized discounting patterns described above obtain from a choice rule trading off exponentially discounted rewards.[6] This choice rule can be mapped into a log-odds representation, which generates the model's key predictions. Such log-odds representations have been used to explain small-stake risk aversion and probability distortions in decisions under risk (Zhang and Maloney, 2012; Zhang, Ren and Maloney, 2020; Khaw, Li and Woodford, 2021; 2023). It has further been argued that log-odds representations may underly neural representations of choice objects in general, based both on prin-

[4]While I specifically focus on the noisy perception of time delays, augmenting this pespective with the noisy perception of rewards does not impact the conclusions I draw below—see Online Appendix B for a generalization.

[5]It may seem odd that systematic biases could arise from neural mechanisms, which would have evolved to optimally guide behaviour across a wide variety of contexts. Such biases, however, commonly arise from resource constraints, and can result from processes that are highly adaptive from an evolutionary point of view—see e.g. Robson (2001), Robson and Samuelson (2011), Netzer (2009) and Herold and Netzer (2023) for examples.

[6]This does not imply that a decision-maker consciously tries to implement such a choice rule. The framework I present is best thought of as arising from working mechanisms internal to the mind, which may have arissen from evolutionary pressures, and which are typically beyond the reach of the consciousness of the decison-maker.

ciples of efficient representation (Gold and Shadlen, 2001; 2002), and on distinctive empirical footprints in observable behaviour (Glanzer, Hilford, Kim and Maloney, 2019). Bayesian combination of evidence and prior constitutes the key working hypothesis underlying brain functioning in neuroscience (Ting, Yu, Maloney and Wu, 2015; Wei and Stocker, 2015; Ma, Kording and Goldreich, 2023).

The model I propose contributes to a unified theoretical representation of decision-generating processes. The formal principles underlying the model, down to the very log-odds choice rule, are identical to those used for the modelling of small-stakes risk aversion and probability distortions (Khaw et al., 2021; Vieider, 2023). L'Haridon, Oprea, Polania and Vieider (2023) use an identical formal setup to illustrate the cognitive origins of the Ellsberg paradox. In Oprea and Vieider (2023), we test a key prediction of the noisy coding setup, according to which learning should take place even for fully described quantities such as objective risk, and find this hypothesis to be strongly supported by the data. While some of the specific assumptions underlying the formalism may well prove contentious, the reach of the explanations provided across different contexts suggests that the model does capture some key aspects of the decision process.

I test the model in two pre-registered experiments devised to probe the most distinctive implications of the noisy perception account of delay-discounting. Experiment I provides a visual aid for the length of the time delay under form of a ruler of length proportional to the delay. This intervention is designed to focus attention on time delays and to make them more comparable, thereby reducing coding noise—a key driver of behaviour according to the model. Another key implication of the dynamic Bayesian model I present is that the parameters of the prior predicting the most likely stimuli in the environment need to be inferred from the noisy mental signals jointly with the true choice stimuli. Experiment II systematically manipulates the statistics of the experimental environment to test this prediction. The two experiments yield evidence in support of the key mechanisms underlying the theoretical predictions I present. This yields a distinctive empirical fingerprint that cannot be organized under existing discounting models.

The model shares a common formal underpinning with recent papers modelling

the effect of noisy, but otherwise optimal, cognitive processes in complex choice environments (Natenzon, 2019; Khaw et al., 2021; 2023). It also presents affinities with contributions that have showcased the importance of 'cognitive uncertainty' in decision making (Enke and Graeber, 2023). The model's intuition is consistent with experiments that have documented how patterns traditionally ascribed to preferences can be triggered by the complexity of the choice situation (Oprea, 2022). Barretto Garcia, de Hollander, Grueschow, Polania, Woodford and Ruff (2023) show how neural activation during a numerical discrimination task predicts risk attitudes in a behavioural experiment, thus highlighting the neural processes from which such behaviour arises. Imitating the complexities of inter-temporal tradeoffs in an atemporal setting, Enke, Graeber and Oprea (2023) document patterns that closely resemble what one may consider 'typical discounting patterns'.

There exists an extensive theoretical and empirical literature describing the stylized discounting patterns discussed above. Yet few if any of these models can aspire to capturing *all* of the stylized patterns jointly. Psychological explanations of pronounced discounting include the 'miseries of delayed gratification', visceral forces, and difficulties imagining future utilities, among others (see Frederick, Loewenstein and O'Donoghue, 2002, section 2, for a review). These explanations were, however, devised to account for difficulties in delaying *consumption*, and are less suited to explain discounting for money. Discounting for money can be organized by models that depict discounting as arising from the inherent uncertainty of the future (Halevy, 2008; Chakraborty, Halevy and Saito, 2020; Epper and Fehr-Duda, 2023), and which aim at capturing present bias and decreasing impatience. They are, however, less suited for reconciling such patterns with the much steeper discounting emerging from a series of shorter delays.[7]

The discrepancy between the two curves leading up to points B and C in figure 1 indeed creates challenges for the whole class of 'generalized hyperbolic' models (Mazur, 1987; Loewenstein and Prelec, 1992; Laibson, 1997; Ebert and Prelec, 2007;

[7]An exception to this rule is provided by the model of Epper and Fehr-Duda (2023), which decomposes survival probabilities for different 'time intervals', which are subject to separate non-linear transformations. Their conceptualization of time is thus similar to the one used by Scholten and Read (2006), which I will discuss shortly.

Bleichrodt, Rohde and Wakker, 2009), which specifically explain changes in impatience observed in increasing delays from the present, as showcased by the curve leading up to point B. Within this class, the model of Gabaix and Laibson (2017) deserves special attention inasmuch as it uses a formally similar approach to the one I employ here to produce cognitive microfoundations for the proportional discount function of Mazur (1987). I will present a detailed discussion of the relation between the two models after presenting my model.

Read (2001) first documented the discrepancy between the patterns leading up to points B and C in figure 1 under the name of *subadditivity*. Subadditivity—whereby revealed discounting for a given time delay is more pronounced when measured over two or more sub-delays than when measured for the entire delay at once—can account both for the steep discounting leading up to point C (which is due to the many short delays), and for the much flatter function leading up to point B (which is attributed to delays increasing in length). Kable and Glimcher (2010) further showed that delays of different length produce the iconic hyperbolically decreasing patterns *indepedently of the up-front delay*, thus showing that *delay length* rather than distance from the present is the relevant dimension. Scholten and Read (2006) develop a 4 parameter function that has dedicated parameters for a variety of empirical patterns, and Scholten, Read and Sanborn (2014) present a model that trades off reward differences against delay differences. Present bias—the distinctive phenomenon whereby people become more patient when the starting point of a *fixed time delay* is moved from the present to the future—can however not be organized by these models separately from generalized hyperbolic patterns.[8]

A line of investigation originating in addiction research and marketing has emphasized the subjective perception of time, as I do here. In two theoretical notes, Takahashi (2005) hypothesizes that a logarithmic perception of time could underly apparently hyperbolic discounting, and Takahashi (2006) shows that a concave power transformation of time can account for subadditivity. Hyperbolicity and sub-

[8]The models proposed by Scholten and Read (2006) and Scholten et al. (2014) also present a mechanism that allows them to capture the absolute magnitude effect. The outcome derivation in Online Appendix B shows how to integreate noisy perception of rewards into the model I present here. A fully-fledged examiniation of the magnitude effect is, however, beyond the scope of this paper.

additivity thereby both emerge based on subjective time perception, but by means of formally different models. Zauberman, Kim, Malkoc and Bettman (2009) and Kim and Zauberman (2009) document how "hyperbolic discounting" can emerge from logarithmic perception of time, and empirically show that subjective time ratings bear a logarithmic relationship to objective time delays from the present (see also Bradford, Dolan and Galizzi, 2019). Cooper, Kable, Kim and Zauberman (2013) provide evidence that neural activations during subjective time rating tasks are predictive of subsequent discounting behaviour. These studies thus provide valuable insights supporting a central aspect of the account I propose here. The focus of these papers, however, remains on descriptively capturing apparent hyperbolicity in discounting measured from delays of increasing length from the present, and is thus very different from the generative account of delay-discounting I present here.

2 Model

2.1 The choice rule and its mental representation

I model standard tradeoffs between a larger reward x, paid at a later time τ_ℓ, and a smaller reward y, paid at a sooner time τ_s. While I use such simple tradeoffs throughout, the model straightforwardly generalizes to more complex tradeoffs—see Online Appendix C. I start from a choice rule devoid of subjective parameters[9], under which the larger-later reward is chosen whenever

$$e^{-\tau_\ell} x > e^{-\tau_s} y. \tag{1}$$

The choice rule above is optimal inasmuch as discounting is stationary, resulting in consistent choice patterns over time. The substantial discounting of 100% per time unit reflects the fundamental intuition underlying the model that patience—and perhaps *the very meaning of time*—needs to be learned from experience. This take is consistent with extremely pronounced discounting observed in children (Mischel

[9]Note that this choice rule is used without loss of generality. 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.

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).

I now rearrange the terms in (1) to emphasize the comparative nature of the tradeoffs involved. I further take the natural logarithm of both sides, which will be convenient for the modelling of the neural processes, but does not affect the conclusions in any way (see Online Appendix A). We thus obtain:

$$ln\left(\frac{p}{q}\right) > ln\left(\frac{y}{x}\right), \tag{2}$$

where $p \triangleq e^{-\tau_\ell}$ and $q \triangleq e^{-\tau_s}$. This formulation emphasizes the parallels with decisions under risk by conceiving of the exponentials as probabilities. The left-hand side then takes the form of log-odds, which have been described as *ubiquitous* in decision-making (Zhang and Maloney, 2012). Several papers have argued that log-odds may underly mental representations of stimuli in general (Gold and Shadlen, 2001; 2002; Glanzer et al., 2019). Noisy perception of log-odds can produce probability distortions under risk as documented in prospect theory (Zhang et al., 2020; Khaw et al., 2023). An identical choice rule and noisy perception mechanism as the one I use here has been shown to generate probability distortions (Vieider, 2023), account for the Ellsberg paradox (L'Haridon et al., 2023), and explain a gap in decisions between experienced versus described options (Oprea and Vieider, 2023). This shows the unifying power of the modelling approach I apply here.

The central idea underlying the model is that choice quantities are not perceived objectively, but are encoded by a signal suitable for neural representation and calculations. An important substantive implication of (2) is that any such signal will encode time delays occurring between two payment options (see Scholten and Read, 2006, and Scholten et al., 2014, for a similar modelling assumption), since it will apply to $ln\left(\frac{p}{q}\right) = ln\left(\frac{e^{-\tau_\ell}}{e^{-\tau_s}}\right) = -(\tau_\ell - \tau_s)$. This captures the idea that errors in time perception arise from quick comparative tradeoffs between choice options, for which the relative delays to the two payments need to be weighed against the relative rewards. I thus postulate that the choice rule implemeted by the mind will be based

on the perception of the delay $t \triangleq \tau_\ell - \tau_s$. I further assume that the choice rule is mentally implemented on (2) multiplied by -1, and logged once again. The logarithm is again taken purely for computational convenience (see Online Appendix A, for an alternative derivation). This results in the following mental choice rule conditional on the signal r:

$$\mathbb{E}\left[\, ln(t) \,|\, r\right] < ln\left(-ln\left(\frac{y}{x}\right)\right), \tag{3}$$

where I treat rewards as being perceived objectively and without distortion for simplicity. Including noisy inferences on outcome ratios does not affect the conclusions drawn from the model (see Online Appendix B).

The mental signal and the likelihood

Given limits on neuronal resources, the mental signal r will necessarily be affected by noise.[10] Noise in mental representations is indeed a hallmark of brain functioning (Dehaene and Changeux, 1993; Dehaene, 2003; Knill and Pouget, 2004; Vilares and Kording, 2011; Ma et al., 2023). I thus model the mental signal as a single draw from the following likelihood function:

$$r \sim \mathcal{N}(\, ln(t)\,,\, \nu^2\,), \tag{4}$$

where r is a shorthand for $r \,|\, t$, emphasizing that the signal is conditional on a specific time delay t, and where the parameter ν quantifies average coding noise, which I assume to be locally independent of t.[11] It is important to note that coding

[10] Oprea and Vieider (2023) show how such noise may result if the underlying log-odds are encoded by a Beta distribution, the parameters of which could sum up the firing rates of neurons and anti-neurons such as discussed by Gold and Shadlen (2001). This follows from the observation that absolute precision can only be obtained from a Beta distribution in the limit as its parameters converge to infinity, where it converges to a Dirac-delta distribution having all its probability mass attributed to a single point. Any finite number of neurons dedicated to the coding will thus necessarily produce noisy representations, with noise decreasing in the allocation of neural resources at a decreasing rate.

[11] Local independence as defined here implies that the coding noise does not change with the time delays used. This assumption is only expected to hold *locally*, in the sense that one should not expect it to hold when transitioning between different orders of magnitude, or indeed for different numerical representations of identical time delays.

noise will make observed behaviour stochastic, given that the signal r from two subsequent draws may be different even conditional on an identical time delay t being presented to a decision-maker. The accuracy with which time delays are perceived will be proportional to $\lambda \triangleq \nu^{-2}$—the *precision* of the likelihood function.

Physiological and topographic evidence of neural activation directly suggests the use of a normal distribution (Nieder and Miller, 2003; Dehaene, 2003; Harvey, Klein, Petridou and Dumoulin, 2013). The normality assumption is further supported by the log-odds representation, since log-odds are usually normally distributed (a fact that is often exploited in statistics; see e.g. Gelman, Carlin, Stern, Dunson, Vehtari and Rubin, 2014, section 5.3). Expressing the signal r on the logarithmic scale serves to avoid an unbounded increase in the cognitive resources needed to represent larger numbers (Dehaene and Changeux, 1993; Dayan and Abbott, 2001; Dehaene, 2003; Izard and Dehaene, 2008). Such a resource-saving implementation is crucial in terms of neural realism, since the mental resources available for any given task at any one time are limited (Shenhav, Musslick, Lieder, Kool, Griffiths, Cohen and Botvinick, 2017; Heng, Woodford and Polania, 2020; Lieder and Griffiths, 2020; Bhui, Lai and Gershman, 2021).

Logarithmic representation of choice stimuli in the brain has received support from behavioural studies (Glanzer et al., 2019), has been proposed on grounds of computational efficiency (Gold and Shadlen, 2001; 2002) and optimality for adaptation (Howard and Shankar, 2018), and has received direct support from single-neuron measurements of the activation functions of number neurons (Nieder and Miller, 2003; Dehaene, 2003; Nieder, 2016). Zauberman et al. (2009) and Kim and Zauberman (2009) have documented a logarithmic relationship between objective time delays from the present and subjective ratings of those delays. Cooper et al. (2013) have shown that the neural activation signatures measured during a task in which subjects are asked to rate *time delays*, without any outcomes involved, predicted later choices between delayed rewards.

Logarithmic representation of time delays incurs into issues when delays become very small. Since $\lim_{t\to 0} ln(t) = -\infty$, the resources needed to represent small delays would increase unboundedly, thus contradicting the very resource-saving ra-

tionale underlying logarithmic coding. Howard and Shankar (2018) argue that this resource-saving rationale dictates that an arbitrary number be added to the objective quantity (as routinely done in the modelling of sensorimotor tasks; see e.g. Petzschner and Glasauer, 2011). I will thus substitute $s \triangleq \psi + \tau_s$ if $\tau_s > 0$, else $s = 0$, and $\ell \triangleq \psi + \tau_\ell$ to model the mental representations of the objective time delays in order to prevent the numerical representation from becoming boundless as $t \to 0$. While this transformation is inconsequential for $\tau_s > 0$ since $\ell - s = \tau_\ell - \tau_s$, it will have substantive implications for $\tau_s \equiv 0$. This can be thought of as an arbitrary fix of a numerical overflow problem that is itself subject to adjustment whenever the importance of the situation warrants it. In this sense, ψ may arise from inattention when decisions are taken quickly and choice quantities are gauged approximately. Its magnitude thus ought to be subject to attentional manipulation, e.g. by increasing attention through visual salience (see Chew, Wang and Zhong, 2023, for a model of attention).

2.2 Bayesian inference on time delays

Once one accepts the presence of noise, the key question becomes: how can one—*given* the noise in the signals—optimize the quality of the inferences drawn from those signals? It turns out that this can be achieved by combining the noisy signal with prior information enshrining the statistical properties of the environment. To this end, the mind builds generative models of the world, enshrined in probabilistic prior distributions learned over time, which enable the mind to draw causal inferences about the outside quantities triggering the signals.

In an experiment, time delays $\boldsymbol{t} = \{t_1, ..., t_n\}$ will be presented exogenously by an experimenter. Delays are not directly observable for the decision-maker, and are mentally represented by a vector of noisy signals, $\boldsymbol{r} = \{r_1, ..., r_n\}$, which are the data used by the mind. The mind uses $\boldsymbol{r}$ to make inferences about the true delays generating the signals, $\boldsymbol{\theta} = \{\theta_1, ..., \theta_n\}$, where $\theta_i \triangleq \mathbb{E}[ln(t_i)|r_i]$ is the posterior expectation of the time delay t_i generating signal r_i. At the same time the mind learns the parameters of the prior $p(\mu, \xi)$, with mean μ and precision $\xi \triangleq \sigma^{-2}$. The mind thus solves a joint inference problem, whereby the joint probability

distribution $p(\boldsymbol{\theta}, \mu, \xi \,|\, \boldsymbol{r})$ is inferred from the noisy signals in $\boldsymbol{r}$. To make the problem more tractable, one can factorize the joint probability distribution:

$$p(\boldsymbol{\theta}, \mu, \xi \,|\, \boldsymbol{r}) = p(\mu, \xi \,|\, \boldsymbol{r}) \times p(\boldsymbol{\theta} \,|\, \boldsymbol{r}, \mu, \xi), \tag{5}$$

where $p(\mu, \xi \,|\, \boldsymbol{\theta})$ is the learned prior distribution parameterized my mean and precision, and $p(\boldsymbol{\theta} \,|\, \boldsymbol{r}, \mu, \xi)$ is the posterior distribution conditional on the parameters of the prior and on the data. I will now derive these elements step by step.

Inference as optimal signal-decoding

In principle, a direct comparison of the delay signal in (4) with (a signal for) the reward ratio would be sufficient to reach a decision (see Thurstone, 1927). However, the average quality of decisions can be improved by combining the noisy signal with prior information enshrining the statistics of the environment. The following prior, for now assumed to be given, captures the statistics of the environment:

$$ln(t) \sim \mathcal{N}(\,\mu\,,\, \xi^{-1}\,). \tag{6}$$

The log-normal form of the prior conforms to the observation that most delays one faces are short while some are very long, and naturally reflects the non-negative nature of time delays. Limpert, Stahel and Abbt (2001) examine a large variety of naturally occurring data, and conclude that there is no single case in which a normal distribution would fit those data better than a log-normal. The conjugate form to the normal likelihood further has important analytical advantages, and allows to quickly access the posterior. Combining the likelihood in (4) with the prior in (6) by Bayesian updating, we obtain the following posterior:

$$p[\,ln(t) \,|\, r\,] = \mathcal{N}\left(\frac{\lambda}{\lambda+\xi}\, r + \frac{\xi}{\lambda+\xi}\,\mu\,,\, \frac{1}{\lambda+\xi}\right), \tag{7}$$

where $\frac{1}{\lambda+\xi} = \frac{\nu^2\sigma^2}{\nu^2+\sigma^2}$ is the variance of the posterior. The posterior mean obtains from a simple linear combination of signal and prior mean, whereby the signal's contribution to the posterior mean is determined by the Bayesian evidence weight

$\beta \triangleq \frac{\lambda}{\lambda+\xi} = \frac{\sigma^2}{\sigma^2+\nu^2}$. I will thus refer to β as *time-discriminability*, or where no ambiguity arises, simply as *discriminability*.

To make the expression in (7) accessible to the experimenter, we will need to replace the mental signal r with an observable quantity. This can be done by means of the *response distribution*, which—given stochasticity in $r \,|\, t$—will capture the distribution of the perceived stimulus, given the *true* stimulus. Let $\theta \triangleq \mathbb{E}[\, ln(t) \,|\, r\,]$ be the expected delay inferred from the noisy signal. We then obtain:

$$p\,(\,\theta \,|\, t\,) = \mathcal{N}\left(\beta\, ln(t) + (1-\beta)\,\mu\,,\; \frac{\lambda}{(\lambda+\xi)^2}\right), \tag{8}$$

where $\frac{\lambda}{(\lambda+\xi)^2} = \frac{\nu^2\sigma^4}{(\nu^2+\sigma^2)^2}$ is the variance of the response distribution (proofs in Online Appendix D). The latter captures the average perceptual noise of a time delay given the true delay being presented to the decision-maker.[12]

The substantive implication of the equation above is that—in the presence of coding noise—unexpected delays falling far from the mean of the prior will be shrunk towards the prior mean more heavily than expected delays. Following (8), we can write the expectation of the response distribution as follows:

$$\begin{aligned} \mathbb{E}[\,\theta \,|\, t\,] &= \mu + \beta\,[\, ln(t) - \mu\,] \\ &= ln(t) + (1-\beta)\,[\,\mu - ln(t)\,]\,. \end{aligned} \tag{9}$$

The first line of the equation conveys the point that we can think of the prior mean as *predicting* the following time delay, so that deviations from that prediction give rise to *prediction errors*. This intuition will be particularly helpful below, when detailing how the prior is learned. The second line emphasizes how regression to the mean will produce systematic deviations from the exponential discounting entailed by the choice rule. To see this, define the bias as the difference between the average inference and the true time delay, $\mathbb{E}[\,\theta \,|\, t\,] - ln(t) = (1-\beta)[\,\mu - ln(t)\,]$. The average inferred time delay will thus be made up of the true time delay plus bias. The

[12] This expression encapsulates the fact that the experimenter can only observe the *average* inference of the decision-maker. Stochasticity in the unobserved mental signal r conditonal on one and the same delay being presented repeatedly then takes the form of variability in the response as captured by the variance of the response distribution.

strength of the bias will be a function of coding noise, converging to an unbiased estimate only as coding noise tends to 0, given that $\lim_{\nu \to 0}(\beta) = 1$. It will further be a function of the distance to the prior mean, capturing how unexpected an observed delay is from the perspective of the learned prior.

Notwithstanding this systematic bias, basing the decision on the posterior mean instead of the maximum likelihood as captured by r is optimal in the sense of minimizing the mean squared error across all choices. To see this, we can use the bias-variance decomposition of the mean squared error:

$$\mathbb{E}\left[\left(\widehat{\theta} - ln(t)\right)^2\right] = \left(\frac{\xi}{\xi + \lambda}\left[\mu - ln(t)\right]\right)^2 + \frac{\lambda}{(\lambda + \xi)^2}, \tag{10}$$

where $\widehat{\theta} \triangleq \mathbb{E}[\,\theta \,|\, t\,]$ is the expectation of the response distribution in (8), the first part on the right-hand side is the squared average bias, $(1 - \beta)^2[\,\mu - ln(t)\,]^2$, and the second part is the variance of the response distribution from (8). The variance of the response distribution entering the definition of the mean squared error is smaller than the variance of the signal r.[13] Furthermore, most deviations $\mu - ln(t)$ will be small due to the nature of the normal distribution. This means that the sum of the squared bias and the variance of the response distribution will stay below the variance of the maximum likelihood estimator in (4) when aggregating across many stimuli (see Ma et al., 2023, section 4.5, for an illustration and detailed discussion). It is in this sense that the biased Bayesian estimator is optimal.[14]

Learning the statistics of the environment

I have so far treated the prior as fixed and known. In reality, however, the parameters of the prior, μ and ξ, are unknown and will need to be learned from the data, conditional on a hyperprior.[15] The expression in (6) above might suggest

[13] This obtains simply from $\frac{\lambda}{(\lambda+\xi)^2} < \lambda^{-1}$, which will be the case for any $\xi > 0$, i.e. for any proper prior variance that is smaller than infinity, i.e. for any $\sigma^2 < \infty$.

[14] This result holds in particular because the quantity of data on which the inference is based will necessarily be small, given that choice quantities have to be assessed quickly. For a general discussion of the optimality of Bayesian estimators in such contexts form a machine learning perspective, see Bishop (2006), chapter 3.

[15] The hyperprior constitutes the hierarchical level above the prior. Hierarchical structures with multiple hierarchical levels are hypothesized to constitute the fundamental building blocks of neu-

that learning of the prior should be based on the true time delays. This quantity, however, is not directly accessible to the mind. The mind thus necessarily needs to learn based on the inferences it draws about the time delays triggering the signals. This will indeed be optimal for the same reasons discussed above—given noise in the signals, the mind cannot do better than relying on its posterior inferences about the true choice stimuli. This gives rise to a recursive learning equation, whereby new inferences are disciplined by accumulated past inferences.[16]

Conditional on the data, $\boldsymbol{r}$, the posterior estimates of μ and ξ will be correlated. To simplify the problem, we can factorize the joint prior into the product of a marginal prior of the precision and the prior of the mean, conditional on the precision, $p(\mu, \xi \,|\, \boldsymbol{r}) = p(\mu \,|\, \boldsymbol{r}, \xi) \times p(\xi \,|\, \boldsymbol{r})$. The conjugate prior to the Gaussian used for the inference will take the form of a Normal-Gamma distribution, $\mathcal{NG}(\mu_0, \kappa_0, \zeta_0, \rho_0)$, which can be written as follows:

$$p(\mu, \xi \,|\, \boldsymbol{r}, \mu_0, \kappa_0, \zeta_0, \rho_0) = \mathcal{N}\left(\mu \,|\, \boldsymbol{r}, \mu_0, \kappa_0 \xi\right) \times \mathcal{G}(\xi \,|\, \boldsymbol{r}, \zeta_0, \rho_0), \tag{11}$$

where $\mathcal{G}$ indicates the Gamma distribution, and where the precision of the normal prior of μ will depend on $\kappa_0 \xi$. The parameter κ_0 thereby serves to create a dependency between the conditional prior for the mean and the marginal prior of the precision, with κ_0 capturing the weight attributed to the prior relative to the data. The parameters ζ_0 and ρ_0 govern the shape and rate of the precision prior, respectively, with $\mathbb{E}[\xi_0] = \zeta_0/\rho_0$ and $\mathbb{V}[\xi_0] = \zeta_0/\rho_0^2$.

I will start from the updating equations of the prior mean, based on the inferred

ral architectures (Friston, 2005), and can present advantages in the highly structured environments we typically encounter in the world. I will return to the interpretation of such hyperpriors in the discussion.

[16] Since the mind cannot directly access $ln(t_i)$, the only alternative would be to base learning directly on the noisy signals r_i. This incurs into the same issues as discussed above for the inference stage—while the signal is unbiased, it is also highly volatile, so that prior learning based on $\boldsymbol{r}$ would be unstable and potentially dependent on chance realizations in early draws.

time delay in a given round i as derived above (proofs in Online Appendix D):

$$\begin{aligned} p(\,\mu_i\,|\,\boldsymbol{r},\mu_0,\kappa_0\xi\,) &= \mathcal{N}\left(\widehat{\mu}_{i-1}+\frac{1}{\kappa_{i-1}+1}\,(\theta_i-\widehat{\mu}_{i-1})\;,\;\left[\,(\kappa_{i-1}+1)\,\widehat{\xi}_{i-1}\,\right]^{-1}\right) \\ &= \mathcal{N}\left(\widehat{\mu}_{i-1}+\frac{\beta_{i-1}}{\kappa_{i-1}+1}\,[r_i-\widehat{\mu}_{i-1}]\;,\;\left[\,(\kappa_{i-1}+1)\,\widehat{\xi}_{i-1}\,\right]^{-1}\right), \end{aligned} \tag{12}$$

where $(\kappa_{i-1}+1)\,\widehat{\xi}_{i-1}$ is the precision of the posterior mean, and where I use the 'hat' on μ and ξ to indicate the expectation of the parameters defining the prior, given that these parameters are now themselves uncertain quantities. The expected posterior mean, $\widehat{\mu}_i$, will thus be a convex combination of the prior mean, $\widehat{\mu}_{i-1}$, and the prediction error, $(\,\theta_i-\widehat{\mu}_{i-1}\,)$, weighed by the strength of the data relative to the weight of the prior, $\frac{1}{1+\kappa_{i-1}}$. The parameter κ_{i-1} captures the confidence in the prior relative to the data, and is updated as $\kappa_i=\kappa_{i-1}+1$.

The last line of the equation above uses the definition of θ to show how learning is based on the prediction error of the *signal*, r_i, weighed by discriminability β_{i-1}, which will typically be smaller than 1. This implies that we should expect the updating of the prior mean to be 'conservative'—the signal will only be taken into account in proportion to its informational content. This will result in incomplete learning, given that for repeated stimuli we can substitute $ln(t_i)$ for r_i in the last line of the equation to obtain a learning equation observable to the experimenter (the variance will however be different in this case—see Online Appendix D for details).[17] Updating will further decline over time in a given environment due to the presence of $\kappa_{i-1}+1$ in the denominator of the learning rate.

To complete the description of the learning process, we can take a look at the updating equations for the precision of the prior:

$$\begin{aligned} p(\xi_i\,|\,\boldsymbol{r},\zeta_0,\rho_0) &= \mathcal{G}\left(\zeta_{i-1}+\frac{1}{2}\;,\;\rho_{i-1}+\frac{1}{2}\frac{\kappa_{i-1}}{\kappa_{i-1}+1}\,(\theta_i-\widehat{\mu}_{i-1})^2\right) \\ &= \mathcal{G}\left(\zeta_{i-1}+\frac{1}{2}\;,\;\rho_{i-1}+\frac{1}{2}\frac{\kappa_{i-1}\,\beta_{i-1}^2}{\kappa_{i-1}+1}\,(r_i-\widehat{\mu}_{i-1})^2\right). \end{aligned} \tag{13}$$

[17] The observation that learning will be incomplete raises the question of whether we can still consider the decision process optimal. That question, however, does not have a simple answer, and will thus be considered in the discussion.

The posterior sum of squares, ρ_i, thus combines the prior sum of squares, ρ_{i-1} with a squared *prediction error*, arising from the discrepancy between the inferred delay and the prior mean, $(\theta_i - \widehat{\mu}_{i-1})^2$. The last line of the equation shows how the precision will be updated based on the squared difference between the prior mean and the *signal*. This, in turn, means that the variance will systematically depend on the coding noise as well as on the variance of the true stimuli.[18] Between subjects, we should thus expect a positive correlation between coding noise ν and the prior standard deviation σ, which constitutes a testable implication of (13).[19] The learning of the precision once again depends on discriminability β_{i-1}, but other than seen for the mean, does not slow down over time. Changes in the environment resulting in larger prediction errors will thus yield temporary increases in the learning rate.

2.3 The stochastic choice rule

We can now use the response distribution in (8) jointly with the dynamic equations from the previous section to arrive at the following stochastic choice rule:

$$Pr[(x_i, \tau_{\ell,i}) \succ (y_i, \tau_{s,i})] = \Phi\left(\frac{ln\left(-ln\left(\frac{y_i}{x_i}\right)\right) - \beta_{i-1}\, ln\left(\alpha_{i-1} t_i\right)}{\beta_{i-1}\, \nu}\right), \tag{14}$$

where $Pr[(x_i, \tau_{\ell,i}) \succ (y_i, \tau_{s,i})]$ is the probability of the larger-later reward being chosen in round i, Φ represents the standard normal cumulative distribution function, and $\alpha_{i-1} \triangleq exp\left(\frac{(1-\beta_{i-1})}{\beta_{i-1}}\, \widehat{\mu}_{i-1}\right)$ captures *impatience*. Dividing the weight of the prior, $1 - \beta_{i-1}$, by the weight attached to the logged time delay, β_{i-1}, thereby

[18] This is most easily seen when reformulating the equation above from the point of view of the experimenter for n time delays, in which case two separate squared deviations for the prediction error of the stimulus and for coding noise will be added to the prior sum of squares—see Online Appendix D for the derivation of such an equation.

[19] The equations I present here predict a causal relation going from coding noise to the variance of the prior. In general, however, we should also expect the coding noise to efficiently adapt to the variation of stimuli in the environment (Laughlin, 1981; Polania, Woodford and Ruff, 2019; Heng et al., 2020; Prat-Carrabin and Woodford, 2022; Heng, Woodford and Polania, 2023), thus resulting in a feedback cycle that further strengthens the correlation. Other than existing models of efficient coding, which typically assume the statistics of the environment to be objectively know to the decision-maker, the learning setup I present here suggests that efficient adaptation will in turn depend on the learned quantities of the prior. Note, however, that what counts for learning is coding noise *relative to the variance of the prior*, ν/σ, which we may think of as a more or less stable individual characteristic.

serves to put the parameter on the scale of subjective time. This allows for easier interpretation of the behavioural effects of the parameter.

According to (14), the probability of choosing the larger-later option will decrease in the ratio of the sooner to the later reward y_i/x_i and in the time delay t_i. For values of $t_i < (\alpha_{i-1})^{-1}$, time discriminability $\beta_{i-1} < 1$ will increase perceived time, making it appear longer than it truly is. For values $t_i > (\alpha_{i-1})^{-1}$, the opposite will be the case, with perceived time delays being compressed relative to their true duration. Conditional on a given value of μ_{i-1}, α_{i-1} furthermore depends on discriminability β_{i-1}, with $\lim_{\beta \to 1}(\alpha) = 1$. As noise vanishes we should thus witness convergence towards the exponential coding rule in (2).

In contrast to traditional discounting models, which combine a deterministic preference model with an independently chosen stochastic choice model (He, Golman and Bhatia, 2019), the probabilistic setup used here produces an inherently stochastic model of inter-temporal choice. The standard deviation of the decision noise, $\beta_{i-1}\,\nu = \frac{\nu\,\sigma_{i-1}^2}{\nu^2+\sigma_{i-1}^2}$, corresponds to the standard deviation of the response distribution in (8). This gives us a new sense of the optimality of the Bayesian combination of signal and prior. Whereas variation of the signal r increases unboundedly in coding noise ν, response noise $\beta\,\nu$ is non-monotonic in ν, increasing up to $\nu = \sigma$ and decreasing again thereafter. Intuitively, for $\nu > \sigma$ the pull towards the prior overpowers the signal, reining in the effect of further increases in coding noise. The interdependency between numerator and denominator further addresses issues arising from scale arbitrariness when both are defined independently as documented by Apesteguia and Ballester (2018).[20]

2.4 Behavioural implications

Present bias and delay-dependent discounting

When modelling delays from the present, i.e. for $\tau_s \equiv 0$ and $t = \ell$, equation (14) can be seen as providing stochastic microfoundations for the constant sensitivity discount function of Ebert and Prelec (2007). This is most easily seen by zooming in

[20] I am indepted to Miguel Ballester for pointing this out.

on the stochastic equality condition in the numerator of (14).[21] Dropping subscripts for simplicity, let $\delta_\ell \triangleq \frac{y}{x}$. Exponentiating the numerator twice yields the following expression at the point of stochastic indifference:

$$\delta_\ell = exp\left(-(\alpha\ell)^\beta\right). \tag{15}$$

Ebert and Prelec (2007) showed that α creates a demarcation between the near and the far future. Discounting for delays shorter than $1/\alpha$ will increase in time insensitivity relative to the exponential benchmark, whereas discounting for delays longer than $1/\alpha$ will decrease (see also Kim and Zauberman, 2009, for the use of the same function to fit subjective ratings of time delays from the present). This insight resonates with the interpretation of the model presented above, whereby α captures regression of time perception to the mean of the prior.

The model predictions, however, deviate from those of the constant sensitivity model in several respects. The definition of the time delay as $\ell \triangleq \psi + \tau_\ell$ implies that the function will show a discontinuous drop in the vicinity of $\tau_\ell = 0$. That is, $\lim_{\tau_\ell \to 0} D(t) = exp(-(\alpha\psi)^\beta)$, where $D(t)$ indicates the discount function. Assuming as usual that $D(0) = 1$, so that immediate payments are not discounted—in the absence of any delay being indicated, there is no time delay to be encoded—this introduces a discontinuity between $\tau_s \equiv 0$ and $\tau_s > 0$ which predicts present-bias—a preference for immediate rewards over even slightly delayed rewards. Present-bias is thus captured by ψ, introduced to fix a numerical overflow problem. This number may itself be subject to attentional modulation.

Once up-front delays are introduced, the noisy perception model predicts that *delays between the two options* are nonlinearly distorted, whereas the constant sensitivity model predicts transformations of the individual time delays from the present, so that $\delta_{\tau_s,\tau_\ell} = exp(-\alpha^\beta(\tau_\ell^\beta - \tau_s^\beta))$. Much like other functions from the non-exponential family, the constant sensitivity model can thus not account for sub-additive discounting. The noisy time perception model, on the other hand, predicts that impatience systematically depends on the time delays used to measure it. For

[21] I define 'stochastic equality' as the larger-later reward being chosen 50% of the time.

a given delay length, however, the function is stationary, i.e. it is independent of the value taken by the up-front delay τ_s except in the special case where $\tau_s = 0$.[22] The fact that this function obtains naturally from an intuitive optimal choice rule may explain the pervasiveness of subadditivity (Read, 2001; Read and Roelofsma, 2003; Scholten and Read, 2006; Dohmen, Falk, Huffman and Sunde, 2017; Enke et al., 2023).

Other than in the constant sensitivity model, time discriminability β is closely linked to decision errors, and the model is inherently stochastic. In the limit as $\nu \to 0$ subadditivity will disappear, and revealed impatience will be independent of the length of the time delays used. The model parameters are furthermore subject to endogenous adjustment dependent on the statistics of the environment—a key feature which distinguishes the model not only from the constant sensitivity model of Ebert and Prelec (2007), but from all existing discounting models.

Relation to Gabaix and Laibson (2017)

Gabaix and Laibson (2017) propose a model that is formally related to the setup presented here. They assume Bayesian agents who are perfectly patient, but who perceive future *utilities* with some noise, so that $s_\tau \sim \mathcal{N}(\, u(x_\tau)\,,\, \omega_\tau^2\,)$, where $u(x_\tau)$ is the utility of a reward x received at time τ, and s_τ is the noisy simulation of that utility. This simulation is combined with a prior $u(x_\tau) \sim \mathcal{N}(\widehat{x}, \eta^2)$. This yields the posterior expectation $E[u(x_\tau) \,|\, s_\tau] = \widehat{x} + D(\tau)(s_\tau - \widehat{x})$. Assuming that the noisiness of the signal increases linearly in the time delay τ, i.e. $\omega_\tau^2 = \omega^2 \times \tau$, they obtain $D(\tau) \triangleq \frac{\eta^2}{\eta^2+\omega^2\tau} = \frac{1}{1+\frac{\omega^2}{\eta^2}\tau}$, which takes the form of the proportional discount function proposed by Mazur (1987), $D(\tau) = \frac{1}{1+\chi\tau}$, with $\chi \triangleq \frac{\omega^2}{\eta^2}$. Note that while future utilities are subject to noisy simulations, time delays are perceived

[22] While decreasing impatience is much-discussed in the economics literature, its empirical status is unclear. Most of historical discussion of decreasing impatience has used delays of increasing lengths from the present (e.g., Thaler, 1981; Ebert and Prelec, 2007; Zauberman et al., 2009). Read (2001) found no evidence for strongly decreasing impatience across 3 experiments. 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). See, however, Bleichrodt, Gao and Rohde (2016) for evidence of strongly decreasing impatience for both health and money.

objectively. The model thus predicts discounting to be proportional to the time delay from the present, but does not predict subadditivity or present-bias. The model of Gabaix and Laibson (2017) provides micro-foundations for psychological accounts according to which future utilities are uncertain. The model thus presents a view that is highly complementary to the one presented here.[23]

3 Experiment I: Manipulation of perceptual noise

One central prediction of the model is that many of the distinctive discounting patterns known from the literature arise from the noisy perception of *time*. Manipulating how delays are displayed thus ought to affect noise in time perception, and thereby shape observed discounting behaviour. In this experiment I thus randomize whether subjects see time delays described by means of textual statements, or whether they are given a visual aid representing time delays by rulers with length proportional to the delays being presented. The visual aid manipulation I use is meant to specifically focus attention onto the time dimension.

The visual aid is shown in figure 2. In the control or *textual* condition, time delays are described in the standard, verbal format. In the *visual* condition, time delays are depicted by means of time lines which are proportional in length to the time delay between now and the moment the reward is paid out. In an additional *outcome* condition the rewards in the textual condition are emphasized by increasing the font in which outcomes are displayed and highlighting them in bold. This serves as a placebo treatment, to ensure that any effects in the visual condition cannot be ascribed to outcomes appearing more prominent as well. I ran the experiment with 451 subjects on Prolific UK. Subjects were paid a fixed fee for their time, and made hypothetical choices between time-delayed payment options. There is no evidence that using hypothetical instead of real choices impacts

[23]For instance, the model presented here does not predict strongly decreasing impatience, in the sense of impatience decreasing for identical time delays as they are pushed farther into the future (Prelec, 2004). While the experiments in this paper do not support strongly decreasing impatience, this may change for consumption goods and when delays become very long, being measured e.g. on a time scale of years instead of weeks. In such cases, a combination of the two models may be desirable, which is straightforward given the models' technical similaries.

observed discounting behaviour (Cohen et al., 2020). Online Appendix F further reports an incentivized classroom experiment I ran as a stability check, and which shows identical qualitative results as the textual treatment.

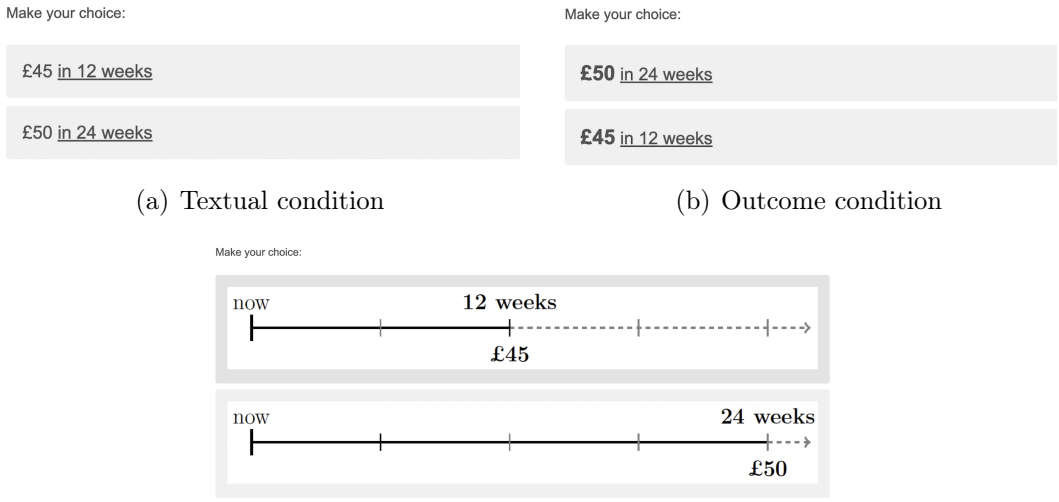

(a) Textual condition

(b) Outcome condition

(c) Visual condition

Figure 2: Choise situation by treatment

Panel 2(a) shows a choice situation in the *textual condition*. Panel 2(c) shows the equivalent choice situation in the *visual* condition. Panel 2(b) shows the choice situation in the outcome condition. Subjects could indicate their choice by clicking on their preferred option, and subsequently moving on the the following screen.

The time delays used in the experiment are depicted in figure 9. They were chosen using simulations to allow for optimal identification of all model components. In particular, comparison of AB to BC and of AC to CE allows for the identification of present-bias. Comparison of AB and BC with AC, of CD and DE with 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 subadditivity. Strongly decreasing impatience can be identified from the comparison of BC, CD, and DE.[24]

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 B for a discussion of robustness to this assumption). The earlier amounts ranged between £33 and £49 inclusive in steps of £1. Each screen presented one single choice, and

[24]Following Prelec (2004), I define strongly decreasing impatience as $(y, \tau_s) \sim (x, \tau_\ell) \to (y, d + \tau_s) \prec (x, d + \tau_\ell)$, where d, τ_s, τ_ℓ are nonnegative time delays, $\tau_s < \tau_\ell$, and $x > y$ are monetary outcomes.

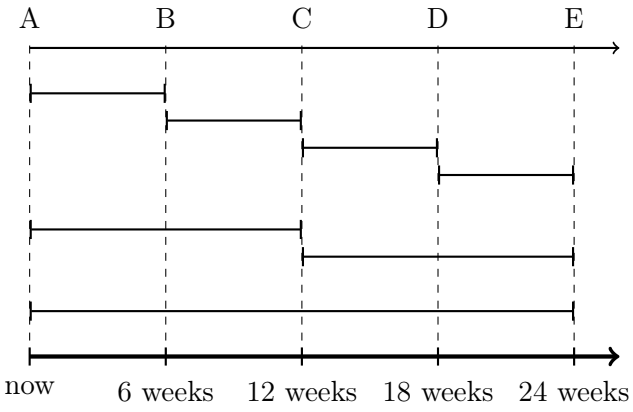

Figure 3: Time horizons used in experiment

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, given 4 delays of 6 weeks, 2 of 12 weeks, and 1 of 24 weeks, the stimuli are well-fitted by a log-normal distribution.

the individual choice pairs were presented completely at random. The position of the smaller-sooner versus larger-later option was also randomized. This design was implemented to fit the discrete binary choice setup modelled. Subjects made 158 choices, which took about 15 minutes on average. All stimuli were presented at least once. In addition, 40 randomly selected stimuli were repeated with replacement. Identification of response noise, which plays a central role in the model, is thus assured by i) repeated observations of the same stimuli; and ii) monotonicity violations between similar stimuli.

A key prediction is that the visual aid for time will decrease coding noise. Importantly, however, by focusing attention on the time dimension, the visual aid manipulation is predicted to also decrease present-bias. This prediction arises from the observation that the visual aid does not only simplify the appraisal and comparison of time delays, but that it also impacts attention paid to the time delays themselves. If the attentional interpretation of ψ is warranted, we should thus expect to see a decrease in present-bias. The outcome manipulation serves as a placebo. It may either not make any difference, or it may further increase coding noise of time and hence decrease delay-discriminability β by focusing attention on the outcome dimension and away from time.

Note that the predictions above are distinctive features of the noisy coding setup I use. Traditional preference-based accounts would not predict any difference in behaviour based on how time delays are displayed, nor would standard psychological explanations based on difficulties in postponing consumption.[25] We also would not expect any differences based on the inherent uncertainty of the future, which remains identical across treatments. Error models do not make any predictions about systematic effects either. It is of course conceivable that the magnitude of the error could vary with the presentation format. Note, however, that models such as those proposed by Lu and Saito (2018) and He et al. (2019) would then predict differences in *decreasing impatience*—a prediction that is quite distinct from those presented here. The same holds true for the model of Gabaix and Laibson (2017), given that in the latter time is perceived *objectively* and noise affects the time-proportionality parameter. The predictions that come closest are those based on accounts meant to capture the fragility or subjective nature of time. For instance, a visual manipulation has been proposed previously by Ebert and Prelec (2007), and Zauberman et al. (2009) have used priming interventions to affect time perception. Both these papers have, however, focused exclusively on delays from the present to quantify "hyperbolic discounting". These manipulations can thus not disentangle strongly decreasing impatience from subadditivity and present-bias—a distinction which is crucial here.[26]

3.1 Results

Figure 4 shows nonparametric discount functions by condition. Panel A shows the discount functions for delays of increasing length from the present. The function in the visual treatment is the highest of the three for the 6-week delay, and the lowest of the three for the 24-week delay, putting it closest to exponential discounting. This shows that the visual treatment has the effect of increasing time discriminability, as

[25] Of course, there could still be a difference affecting the noise attached to such models. Such noise, however, is typically attached to the discounted utility, and modelled as "while noise" independent from the preference parameters.

[26] In particular, a model estimating a hyperbolic-type function will subsume any subadditivity in decreasing impatience. It will also not distinguish between a drop near 0 from general patterns of declining impatience. See Enke et al. (2023) for a discussion of the problems arising from this.

hypothesized. Panel B shows discount functions obtained by multiplying discount factors for subsequent 6-week delays from the present, from 6 weeks, from 12 weeks, and from 18 weeks. Once again, the function in the visual treatment is the closest to an exponential function amongst the three, showing increased time discriminability.

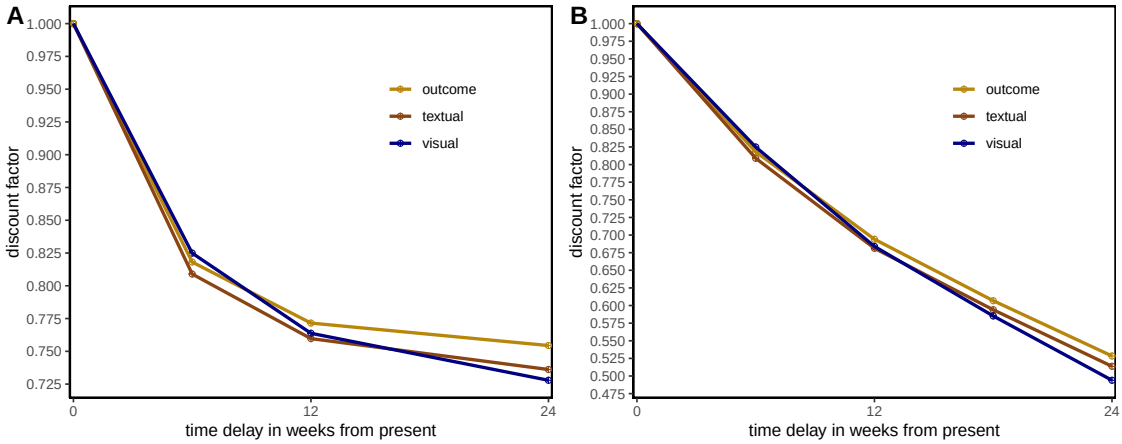

Figure 4: Nonparametric discount functions by treatment

The figure shows nonparametric discount functions by treatment, with linear interpolation between the data points. Panel A shows the functions obtaining when using delays of increasing length from the present. Panel B shows the function obtaining when multiplying subsequent 6-week delays, i.e. by using delays of 6 weeks from the present, 6 weeks from 6 weeks, 6 weeks from 12 weeks, and 6 weeks from 18 weeks, so that the discount factor for a 12-week delay, δ_{12}, is calculated as $\delta_{12} = \delta_{0,6} \times \delta_{6,12}$, where the subscripts indicate the sooner and later time delays attached to the smaller and larger payoffs, respectively. Following the same procedure, we obtain $\delta_{18} = \delta_{0,6} \times \delta_{6,12} \times \delta_{12,18}$, and $\delta_{24} = \delta_{0,6} \times \delta_{6,12} \times \delta_{12,18} \times \delta_{18,24}$. The figure for subsequent 12-week delays looks similar, and is not shown.

Next let us have a look at the most distinctive prediction arising from the increased attention to the time delays hypothesized in the visual condition—the decrease in present-bias. Figure 5 shows decumulative choice proportions of the larger-later reward as the smaller-sooner reward increases. Panel A shows the choice proportions in the textual condition, comparing a 6-week delay from the present to a 6-week delay from 6 weeks. Panel D shows the choice proportions in the textual treatment for a 12-week delay from the present compared to choice proportions for a 12-week delay from 12 weeks. In both cases, there is clear evidence for present-bias ($p < 0.001$ in both cases; Wilcoxon signed-rank test).

Panels B and E show the equivalent figures for the visual condition. The choice proportions for the delays from the present and for the delays from the upfront delay are now virtually indistinguishable ($p > 0.47$ in both cases). One can furthermore see that the difference is closed from the side of the immediate delays, i.e. we clearly witness *a reduction in present-bias*, rather than an overall increase in impatience.

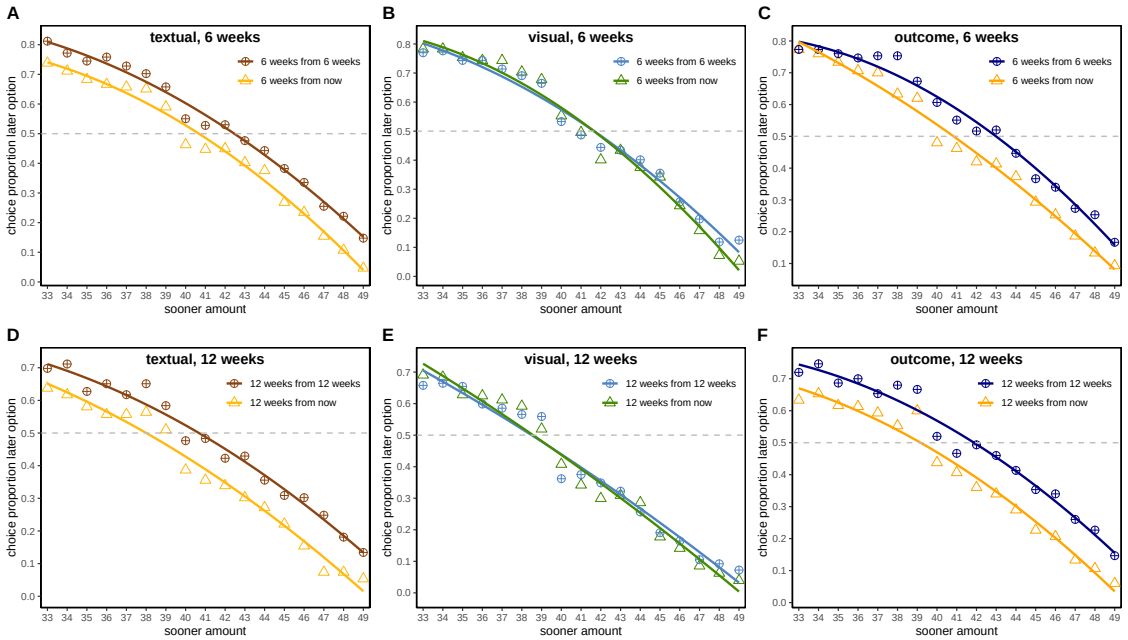

Figure 5: Decumulative choice porportions for larger-later option by treatment

Decumulative choice proportions of the larger-later reward as the smaller-sooner reward increases from £33 through £49. The later amount was kept constant at £50 throughout. Curves fit to the choice proportions are based on second degree polynomials. Panel A compares the choice proportions for a 6-week delay from the present, to those for a 6-week delay from 6 weeks in the textual condition. Panel B shows the equivalent comparison for the visual condition. Panel C shows once more the same comparison for the outcome condition. Panels D to F show the equivalent comparisons using a 12 week delay from now versus a 12 weeks delay from 12 weeks, for the textual condition (panel D), the visual condition (panel F), and the outcome condition (panel F).

Finally, panels C and F show decumulative choice proportions of the larger-later option in the outcome condition for delays of 6 weeks and 12 weeks, respectively. Present-bias is again strong in both cases ($p < 0.001$). This shows how the effects of the visual treatment cannot be explained by a confound deriving from the increased visibility of the *outcome* dimension.

Figure 6 shows a comparison of the structurally estimated parameters across treatment conditions. Panel A shows the empirical cumulative distribution function of the individual-level coding noise means. The curves are clearly distinct. Coding noise in the visual condition is lower than in the textual condition ($p = 0.005$, Wilcoxon ranksum test on individual means), as well as compared to the outcome condition ($p < 0.001$). Coding noise is furthermore significantly larger in the outcome condition compared to the textual condition ($p = 0.001$). Panel B shows time-discriminability parameters β. Discriminability shows systematic differences across the conditions. In particular, discriminability in the visual condition is significantly larger than in both the textual condition ($p = 0.016$) and the outcome

condition ($p < 0.001$). Discriminability in the textual condition is in turn larger than in the outcome condition ($p = 0.007$). This shows the causal effect of the visual aid on noise in delay perception.

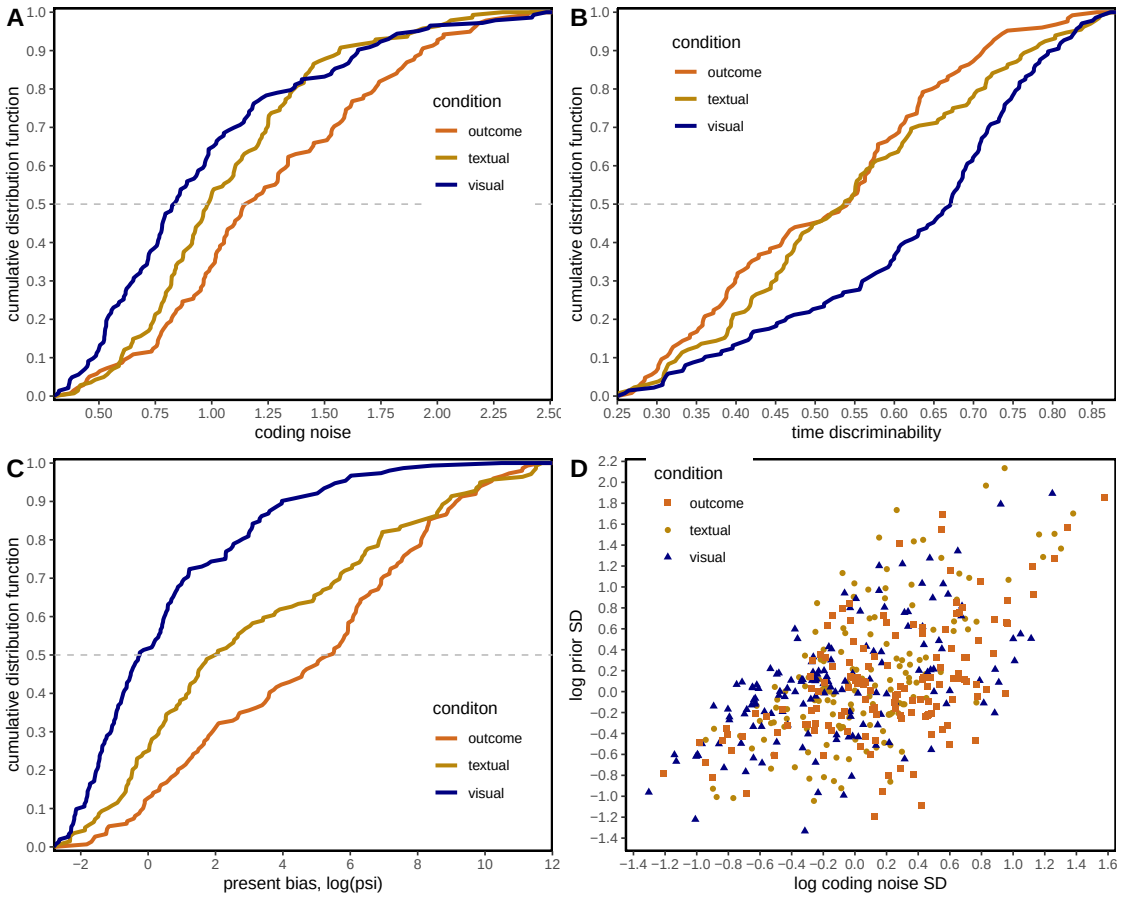

Figure 6: Comparison of structural model parameters

The figure shows individual level parameters, defined as the means of the posterior distributions. Panel A shows the patterns for coding noise, ν. Panel B shows time-discriminability, β. Panel C shows the logged present-bias parameter, $log(\psi)$. Panel D shows a scatter plot of logged coding noise SD $ln(\nu)$ against the logged prior SD, $ln(\sigma)$. Some outliers may be cut from the graphs to improve the visualization.

Panel C shows the cumulative distributions for the natural logarithm of the present bias parameter, $ln(\psi)$. The parameters in both the textual condition and the outcome condition are much larger than they are in the visual condition ($p < 0.001$ in both comparisons).[27] The present-bias parameter is furthermore significantly smaller in the textual than in the outcome condition ($p = 0.008$). The pronounced difference we see between the visual treatment and the textual and outcome treat-

[27]While ψ is the main parameter driving present bias in the model, the overall incidence of present bias will be governed by $(\alpha\psi)^{\beta}$. All the differences discussed in the main text remain statistically significant if I use this measure instead.

ments gives credence to the hypothesis that present-bias is subject to attentional modulation. Finally, panel D plots the logged standard deviation of the coding noise, $ln(\nu)$, against the logged standard deviation of the prior, $ln(\sigma)$. In all three treatment conditions, the two parameters are strongly positively correlated between subjects ($\rho > 0.47$, $p < 0.001$ in all 3 conditions according to a Spearman rank correlation). This corresponds to the prediction of the joint inference model, thus providing suggestive evidence in favour of the learning mechanism modelled.

4 Experiment II: Moving the prior

The visual display manipulation just discussed showed one of the central mechanisms underlying the model at work. A second distinctive mechanism underlying the model concerns the effect of the prior mean. The regression to the mean of the prior underlying time distortion yields a distinctive prediction: one ought to be able to systematically change patience simply by manipulating the choice context. In particular, we would expect that choice behaviour measured on a common set of test tasks ought to exhibit more patience when it is preceded by an adaptation choice set containing relatively short time delays, than when it is preceded by a set with relatively long time delays.

Note that this is a distinctive prediction emerging from the Bayesian inference model I have presented above. Neither exponential, nor hyperbolic or quasi-hyperbolic functions contain a mechanism predicting that choices will be influenced by context effects. Standard interpretations of these models hold that the observed patterns are due to preferences, which are implicitly treated as exogenous. The prediction is also at odds with accounts ascribing non-standard discounting patterns to the inherent uncertainty of the future, since the latter cannot possibly be impacted by the nature of the choice set. Noisy simulations of future rewards will remain the same, since I keep the rewards constant, and time is perceived objectively in that model. Even subadditive models and accounts that have discussed the subjective nature of time do not contain any formal mechanism predicting this sort of effect. The prediction whereby patience should systematically change according to

the initial stimuli presented in an experiment is thus unique to the joint Bayesian inference model I present here.

The experiment consists of binary choices between smaller-sooner and larger-later rewards shown using a standard textual display. There are two parts: part 1 contains the adaptation stimuli, and part 2 the test stimuli. The test stimuli in part 2 are like the delays AB, AC, CE, and AE used in experiment I, except that all delays are doubled (i.e., AC now corresponds to 24 weeks instead of 12). The reason for this is that doubling the test delays allows me to both lengthen and shorten the delays in the adaptation phase, while still using round delays in weeks. The later reward was kept fixed at £50, while the sooner reward varied between £25 and £49 in steps of £1. This yields 100 test choices, to which I add 20 randomly repeated choices. Choices were presented one-by-one in random order.

Part 1 contained the adaptation phase, which constitutes the main treatment manipulation. Subjects in the *long* condition were presented with binary choices just like in the test phase, but ranging over longer delays of up to 92 weeks. Subjects in the *short* condition were presented with binary choices involving delays of at most 12 weeks. The conditions contained the following 4 delays:

$$D_s = \{\,(0,6),(0,8),(6,12),(6,18)\,\}\ ,\ \ D_l = \{\,(0,84),(0,92),(6,84),(12,82)\,\},$$

where D_s stands for short delays, and D_l for long delays, with all delays indicated in weeks. The delays were constructed in such a way as to keep the variance of delays constant across treatment conditions over the whole experiment. As in the test phase, the later monetary reward was kept fixed at £50, while the sooner rewards varied between £25 and £49 in steps of £1. This yields 100 adaptation trials, which were presented to subjects in random order. The change from adaptation phase to test phase was seamless, i.e. subjects were given no sign or warning that there were different parts to the experiment.

The main prediction concerns the prior mean μ, and hence the impatience parameter α. In particular, both should be larger in the *long* condition than in the *short* condition, thus yielding more impatient behaviour in the long condition. This

is a distinctive prediction emerging from the model, given that standard models are static and do not allow for the adaptation of 'preference parameters' to the statistics of the environment. While additional model parameter may also change following the dynamic learning equations presented above, I will here focus on changes in impatience, following the pre-registered analysis plan.

4.1 Results

Panel A in figure 7 shows nonparametric discount functions for the two conditions obtained from delays of increasing length from the present, and panel B shows nonparametric discount functions obtained from subsequent 24 week delays. The discount functions in the *short* condition, where subjects were presented with relatively short delays in the *adaptation* phase, stays clearly above the function observed in the *long* condition, where subjects were presented with relatively long time delays during the adaptation phase. This is the case in both comparisons, and the effects are economically sizeable. This result thus supports a central prediction of the joint inference model, whereby the initial time delays experienced during the experiment should systematically influence subsequent behaviour.

Panel C further shows decumulative choice proportions for the larger-later reward obtaining 24 weeks from the present as the sooner reward increases from £25 through £49, and panel D shows equivalent decumulative choice proportions for the 24 week interval starting 24 weeks from the present (relative choice proportions for other delays look very similar). The choice proportions for the larger-later option are clearly larger in the *short* condition ($p < 0.001$ in both cases). The difference is further sizeable, amounting to about 10 percentage points for the lowest sooner rewards. The nonparametric evidence thus leaves no doubt that being exposed to shorter delays in an initial adaptation phase yields higher levels of patience than being initially exposed to long delays, whereby patience is measured on identical choice alternatives. This constitutes evidence for a key mechanisms underlying the joint inference model—that "revealed preferences" ought to systematically depend on the statistics of the environment—which sets its prediction apart from the predictions of all existing accounts of time discounting.

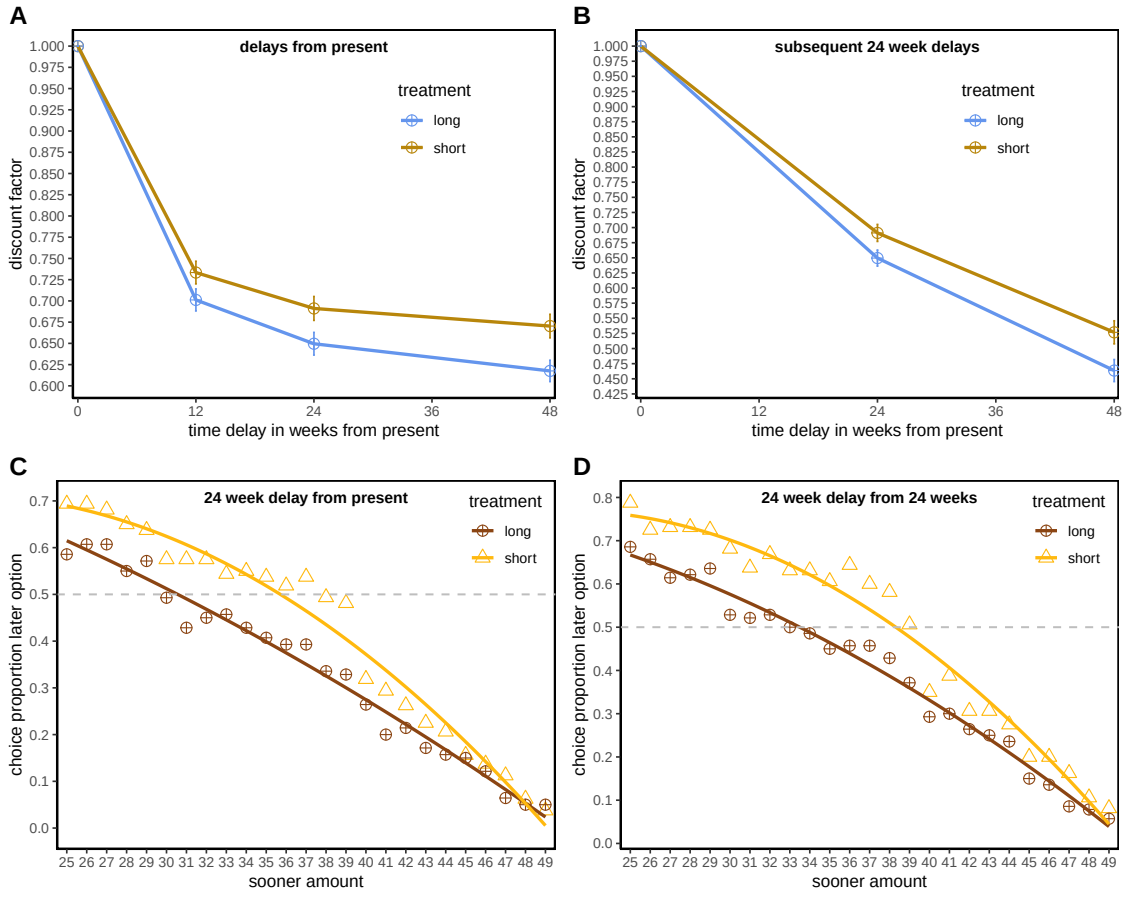

Figure 7: Effects of different adaptation choices, nonparametric

The figure shows the effect of different time delays shown during an adaptation phase preceding the test phase on which the graphs are based. Panel A shows nonparametric discount functions using delays of different length from the present. Panel B shows nonparametric discount functions calculated based on subsequent 24 week delays, so that $\delta_{48} = \delta_{0,24} \times \delta_{24,48}$. Panel C compares decumulative choice proportions for the larger-later option as a function of the earlier payment for a delay of 24 weeks from the present (line fit by a second-order polynomial). Panel D shows decumulative choice proportions for a 24 week delay from an upfront delay of 24 weeks.

Figure 8 shows structural estimates of model parameters. The parameters have been estimated purely on the stimuli of the test phase, which are identical across treatments. Panel A displays the empirical cumulative distribution function of the logged impatience parameter, $ln(\alpha)$. The distribution of impatience parameters obtained from the condition with short delays during the adaptation phase lies below the mean estimated from the condition with long delays in the adaptation phase ($p = 0.001$, Wilcoxon ranksum test).[28] Panel B shows a scatter plot of discriminability β against $ln(\alpha)$. The graph illustrates how subjects with large discriminability

[28] A similar difference emerges when looking at the mean of the prior, μ. Here, we also find the parameter to be smaller for subjects who have been exposed to short delays rather than long delays during the adaptation phase, indicating less impatience ($p = 0.002$).

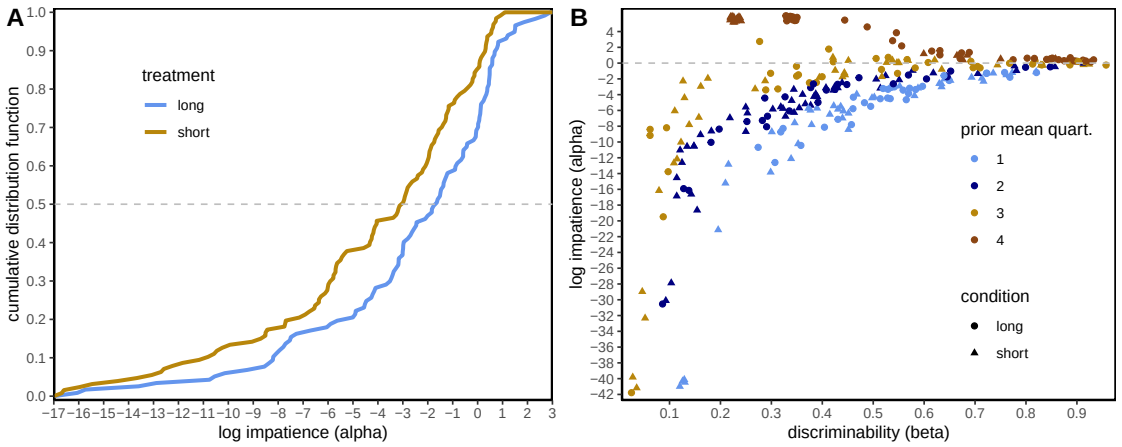

Figure 8: Structural parameters after prior adaptation

The figure shows the effect of different time delays shown during an adaptation phase preceding the test phase on structural model parameters. Panel A shows the empirical cumulative distribution function of the logged impatience parameter, $ln(\alpha)$. Panel B shows a scatter plot between discriminability β and the logged impatience parameter, $ln(\alpha)$, with colours indicating quartiles of the prior mean μ, and shapes showing the treatment condition. Some observations may be cut for an improved visual display.

β also have logged impatience parameters converging towards 0, which here corresponds to a correct assessment of the average delay on the subjective time scale.[29] For values of $\beta < 0.5$, where coding noise exceeds the prior standard deviation, impatience parameters diverge in both directions, indicating higher behavioural dispersion when the prior is learned imperfectly. These results further illustrate the causal effect of the statistical environment on behaviour.

5 Discussion and Conclusion

I have presented a *generative* account of delay-discounting. Other than descriptive models, which aim to mathematically capture observed behaviour as accurately as possible, such a model aims to represent the underlying processes from which observed choice patterns may arise. Modelling this decision process in a way that is plausible from a neuro-biological point of view, and using a choice rule and modelling setup identical to those used to model decision processes under risk, uncertainty,

[29]The estimations use a time scale that has been normalized to yearly delays by division of the delayes in weeks with 52. The impatience parameter α is measured on that same scale, so that convergence towards 0 of $ln(\alpha)$ indicates convergence towards 1 of α, thus indicating the perception of 1 time unit on the subjective scale. Convergence of β towards 1, however, also implies that the subjective time scale converges towards the objective scale, thus yieling convergence towards the optimal choice rule.

and ambiguity, I have shown the model to reconcile several findings that remain challenging for any *one* descriptive model to organize. Two experimental tests showed support for the central working mechanisms of the model—coding noise and the prior mean. The experimental results present a unique empirical fingerprint that cannot be organized by any existing model of discounting.

The model predictions and experimental findings described above are consistent with a number of stylized facts in the delay-discounting literature. For instance, Ebert and Prelec (2007) and Zauberman et al. (2009) have documented how the time dimension is 'fragile', subjective, and subject to manipulations. The model I have presented shares a common intuition with these accounts of discounting, but also differs from them in important ways. One, time perception in the Bayesian inference model applies to delays between different payment options, rather than to the time attached to single rewards, which allows it to reconcile apparent hyperbolicitly in delays of different length from the present with subadditivity (see also Scholten and Read, 2006, and Scholten et al., 2014). Two, it affords a central role to the prior, thus predicting context effects of the sort documented in experiment II, which are not predicted by the other models. The central role played by described time delays in the model also helps to shed light on other findings in the literature. For instance, Read, Frederick, Orsel and Rahman (2005) have documented a 'date/delay effect', whereby discounting weakens substantially when delayed payoffs are associated with calendar dates rather than with time delays (see also LeBoeuf, 2006). This is highly consistent with the model I presented, since in the absence of any delays to be noisily encoded, we would expect the effect of noisy delay perception to subside.

The noisy coding of time model I have presented makes stochastic predictions on choice behaviour in tradeoffs between smaller-sooner and larger-later amounts based on the noisy perception of *time delays*. A related literature has studied the effects of randomness in choices and/or 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. The predictions emerging from these two classes of models are quite distinct. While the noisy coding model predicts delay-dependence and present-bias

based on different mechanisms, these error models predict that decreasing impatience may arise from randomness in preferences. The predictions of these models are thus similar to those of the hyperbolic discounting literature, even though the patterns arise from noise rather than preferences.

Some of the specific mechanisms underlying the model I presented may well prove contentious. It is, however, important to note that many of the central predictions of the model do not hinge on a literal interpretation of the modelling assumptions. For instance, we would not expect the mind to encode the parameters of the normal distributions and to calculate the precise posteriors. The exact neural underpinnings of Bayesian representations remain disputed at this point (Ma and Jazayeri, 2014). While the statistics of the environment could be encoded by populations of neurons (Ma, Beck, Latham and Pouget, 2006), they may alternatively be represented by a limited number of samples that are updated with new samples over time (Sanborn and Chater, 2016; Prat-Carrabin, Wilson, Cohen and Azeredo da Silveira, 2021). Although the precision with which a Bayesian mechanism is implemented will differ between these coding paradigms, the key point is that (approximate) calculations such as the ones used here are well within the reach of the human mind. This makes the model plausible from a neuro-biological point of view.

I have described the neural processes underlying inter-temporal decisions as optimal, subject to resource constraints. Such a description is in line with the convergence towards a "resource-rational" view of decision-making in neuroscience (Lieder and Griffiths, 2020; Bhui et al., 2021), although different conceptions of resource-rationality coexist (Ma and Woodford, 2020). It is thus important to note that the optimality I describe is strictly conditional on the noisy logarithmic comparison of choice stimuli. Noise is an inherent feature of neural activation, and logarithmic coding has been argued for based on efficiency grounds, using direct physiological evidence, and for computational reasons. I thus take these elements as a starting point for the model presented in this paper. Conditional on noisy logarithmic coding, however, the combination of noisy signals with a prior capturing the statistical properties of the environment can in general be considered optimal.

The feature of the dynamic updating equations whereby stimuli are only learned imperfectly may introduce additional bias, thereby questioning the optimality of a decision process that replies on imperfectly learned environmental statistics. Once again, however, a central point of the model is that—in the presence of coding noise—the mind may not be able to do better than using these noisy signals to the best effect possible. Learning based on a discounted signal that is taken into account in proportion to its informational content may then well be the best the mind can do. This is especially true if—as hypothesized here—the learning is based on a hyperprior which summarizes the wider experiences of the individual across a variety of contexts. Given that the world we live in is highly structured (after all, a year has the same duration regardless of where on lives), we should expect most specific decision situations (such as e.g. an experiment) to present similarities to the wider context (e.g., typical delays one faces in the real world). Relying on higher-level hierarchies as a starting point for situation-specific learning may then again be considered optimal, although a detailed investigation of this issue is beyond the scope of this paper.

Multi-level hierarchical structures have been hypothesized to constitute the fundamental building blocks of neural architectures (Friston, 2005). One can think of hyperpriors as having evolved for optimal adaptation, whereby general statistics of the environment will also influence those statistics in most specific situations occurring within that same environment. In this sense, the hierarchical structure of the priors can be seen as reflecting statistical regularities in the world, and thus to be optimal for adaptation. Adaptation to specific environments will, in turn, be optimal whenever perfect perception of all choice stimuli is not possible due to processing constraints—a central feature of humans, given the sheer amount of sensory stimuli we are exposed to during every waking minute, and the computational demands that accurate processing of such stimuli would impose (Bossaerts, Yadav and Murawski, 2019).

Hyperpriors can then be interpreted as the place where something akin to the traditional conception of "preferences" may originate within the model. We would expect the parameters of the hyperprior to be affected by lifetime experiences across

a wide variety of contexts an individual has been exposed to. One may thereby expect particularly strong effects arising from salient experiences—high stakes outcomes, experiences during formative years, persistent regularities in the environment—as well as from educational inputs as modelled in economics (Doepke and Zilibotti, 2014; 2017). Patience that is advantageous in given environments may then well be passed on through the generations, thus yielding traits that persist over time and across contexts (Galor and Özak, 2016). Isolating the parameters of the hyperprior may thus be our best bet for getting at stable components of behaviour that characterize an individual. Whether we want to call such traits 'preferences', or whether we would rather devise another term to refer to them, is then mostly a question of semantics and taste.

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

Apesteguia, Jose, and Miguel A Ballester (2018) 'Monotone stochastic choice models: The case of risk and time preferences.' *Journal of Political Economy* 126(1), 74–106

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

Barretto Garcia, Miguel, Gilles de Hollander, Marcus Grueschow, Rafael Polania, Michael Woodford, and Christian C. Ruff (2023) 'Individual risk attitudes arise from noise in neurocognitive magnitude representations.' *Nature Human Behaviour, forthcoming*

Bettinger, Eric, and Robert Slonim (2007) 'Patience among children.' *Journal of*

Public Economics 91(1-2), 343–363

Bhui, Rahul, Lucy Lai, and Samuel J Gershman (2021) 'Resource-rational decision making.' *Current Opinion in Behavioral Sciences* 41, 15–21

Bishop, Christopher M. (2006) *Pattern recognition and machine learning,* vol. 4 (Springer)

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

Bossaerts, Peter, Nitin Yadav, and Carsten Murawski (2019) 'Uncertainty and computational complexity.' *Philosophical Transactions of the Royal Society B* 374(1766), 20180138

Bradford, W David, Paul Dolan, and Matteo M Galizzi (2019) 'Looking ahead: Subjective time perception and individual discounting.' *Journal of Risk and Uncertainty* 58, 43–69

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

Chakraborty, Anujit, Yoram Halevy, and Kota Saito (2020) 'The relation between behavior under risk and over time.' *American Economic Review: Insights* 2(1), 1–16

Cheung, Stephen L, Agnieszka Tymula, and Xueting Wang (2023) 'Quasi-hyperbolic present bias: A meta-analysis.' *Life Course Centre Working Paper*

Chew, Soo Hong, Wenqian Wang, and Songfa Zhong (2023) 'Attention theory.' *Mimeo*

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

Cooper, Nicole, Joseph W Kable, B Kyu Kim, and Gal Zauberman (2013) 'Brain activity in valuation regions while thinking about the future predicts individual discount rates.' *Journal of Neuroscience* 33(32), 13150–13156

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

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'

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

Ebert, Sebastian, Wei Wei, and Xun Yu Zhou (2020) 'Weighted discountingÑon group diversity, time-inconsistency, and consequences for investment.' *Journal of Economic Theory* 189, 105089

Enke, Benjamin, and Thomas Graeber (2023) 'Cognitive uncertainty.' *Quarterly*

Journal of Economics

Enke, Benjamin, Thomas Graeber, and Ryan Oprea (2023) 'Complexity and time.' Technical Report, National Bureau of Economic Research

Epper, Thomas, and Helga Fehr-Duda (2023) 'Risk in time: The intertwined nature of risk taking and time discounting.' *Journal of the European Economic Association, forthcoming*

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

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

Galor, Oded, and Ömer Özak (2016) 'The agricultural origins of time preference.' *American Economic Review* 106(10), 3064–3103

Gelman, Andrew, John B Carlin, Hal S Stern, David B Dunson, Aki Vehtari, and Donald B Rubin (2014) *Bayesian data analysis,* vol. 2 (CRC press Boca Raton, FL)

Glanzer, Murray, Andrew Hilford, Kisok Kim, and Laurence T Maloney (2019) 'Generality of likelihood ratio decisions.' *Cognition* 191, 103931

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

_ (2002) 'Banburismus and the brain: decoding the relationship between sensory stimuli, decisions, and reward.' *Neuron* 36(2), 299–308

Green, David Marvin, John A Swets et al. (1966) *Signal detection theory and psychophysics,* vol. 1 (Wiley New York)

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

Harvey, Ben M, Barrie P Klein, Natalia Petridou, and Serge O Dumoulin (2013) 'Topographic representation of numerosity in the human parietal cortex.' *Science* 341(6150), 1123–1126

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

Heng, Joseph A, Michael Woodford, and Rafael Polania (2020) 'Efficient sampling and noisy decisions.' *Elife* 9, e54962

_ (2023) 'Efficient numerosity estimation under limited time.' *bioRxiv* pp. 2023–07

Herold, Florian, and Nick Netzer (2023) 'Second-best probability weighting.' *Games and Economic Behavior* 138, 112–125

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

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

Izard, Véronique, and Stanislas Dehaene (2008) 'Calibrating the mental number line.' *Cognition* 106(3), 1221–1247

Jackson, Matthew O, and Leeat Yariv (2014) 'Present bias and collective dynamic choice in the lab.' *American Economic Review* 104(12), 4184–4204

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

_ (2023) 'Cognitive imprecision and stake-dependent risk attitudes.' Technical Report

Kim, B Kyu, and Gal Zauberman (2009) 'Perception of anticipatory time in temporal discounting.' *Journal of Neuroscience, Psychology, and Economics* 2(2), 91

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

Laughlin, Simon (1981) 'A simple coding procedure enhances a neuron's information

capacity.' *Zeitschrift für Naturforschung c* 36(9-10), 910–912

LeBoeuf, Robyn A (2006) 'Discount rates for time versus dates: The sensitivity of discounting to time-interval description.' *Journal of Marketing Research* 43(1), 59–72

L'Haridon, Olivier, Ryan Oprea, Rafael Polania, and Ferdinand M. Vieider (2023) 'Cognitive foundations of ambiguity atttitudes.' *Mimeo*

Lieder, Falk, and Thomas L Griffiths (2020) 'Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources.' *Behavioral and brain sciences* 43, e1

Limpert, Eckhard, Werner A Stahel, and Markus Abbt (2001) 'Log-normal distributions across the sciences: keys and clues.' *BioScience* 51(5), 341–352

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

Ma, Wei Ji, and Mehrdad Jazayeri (2014) 'Neural coding of uncertainty and probability.' *Annual Review of Neuroscience* 37, 205–220

Ma, Wei Ji, and Michael Woodford (2020) 'Multiple conceptions of resource rationality.' *Behavioral and Brain Sciences*

Ma, Wei Ji, Jeffrey M Beck, Peter E Latham, and Alexandre Pouget (2006) 'Bayesian inference with probabilistic population codes.' *Nature neuroscience* 9(11), 1432–1438

Ma, Wei Ji, Konrad Paul Kording, and Daniel Goldreich (2023) *Bayesian Models of Perception and Action: An Introduction* (MIT press)

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

Natenzon, Paulo (2019) 'Random choice and learning.' *Journal of Political Economy* 127(1), 419–457

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

Nieder, Andreas (2016) 'The neuronal code for number.' *Nature Reviews Neuroscience* 17(6), 366–382

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*

Oprea, Ryan, and Ferdinand M. Vieider (2023) 'Risk, noise, and experience.' *Mimeo*

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

Polania, Rafael, Michael Woodford, and Christian C Ruff (2019) 'Efficient coding of subjective value.' *Nature neuroscience* 22(1), 134–142

Prat-Carrabin, Arthur, and Michael Woodford (2022) 'Efficient coding of numbers explains decision bias and noise.' *Nature Human Behaviour* 6(8), 1142–1152

Prat-Carrabin, Arthur, Robert C Wilson, Jonathan D Cohen, and Rava Azeredo da Silveira (2021) 'Human inference in changing environments with temporal structure.' *Psychological Review* 128(5), 879

Prelec, Drazen (2004) 'Decreasing Impatience: A Criterion for Non-stationary Time Preference and "Hyperbolic" Discounting.' *Scandinavian Journal of Economics* 106(3), 511–532

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

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

Robson, Arthur J, and Larry Samuelson (2011) 'The evolutionary foundations of preferences.' In 'Handbook of social economics,' vol. 1 (Elsevier) pp. 221–310

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

Sanborn, Adam N, and Nick Chater (2016) 'Bayesian brains without probabilities.' *Trends in cognitive sciences* 20(12), 883–893

Scholten, Marc, and Daniel Read (2006) 'Discounting by intervals: A generalized model of intertemporal choice.' *Management Science* 52(9), 1424–1436

Scholten, Marc, Daniel Read, and Adam Sanborn (2014) 'Weighing outcomes by time or against time? Evaluation rules in intertemporal choice.' *Cognitive Science* 38(3), 399–438

Shenhav, Amitai, Sebastian Musslick, Falk Lieder, Wouter Kool, Thomas L Griffiths, Jonathan D Cohen, and Matthew M Botvinick (2017) 'Toward a rational and mechanistic account of mental effort.' *Annual review of neuroscience* 40, 99–124

Sozou, Peter D (1998) 'On hyperbolic discounting and uncertain hazard rates.' *Proceedings of the Royal Society of London. Series B: Biological Sciences* 265(1409), 2015–2020

Takahashi, Taiki (2005) 'Loss of self-control in intertemporal choice may be attributable to logarithmic time-perception.' *Medical Hypotheses* 65(4), 691–693

_ (2006) 'Time-estimation error following weber–fechner law may explain subadditive time-discounting.' *Medical hypotheses* 67(6), 1372–1374

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

Thurstone, Louis L (1927) 'A law of comparative judgment.' *Psychological review* 34(4), 273

Ting, Chih-Chung, Chia-Chen Yu, Laurence T Maloney, and Shih-Wei Wu (2015) 'Neural mechanisms for integrating prior knowledge and likelihood in value-based probabilistic inference.' *Journal of Neuroscience* 35(4), 1792–1805

Vieider, Ferdinand M. (2023) 'Decisions under uncertainty as bayesian inference on choice options.' *Management Science, forthcoming*

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

Wei, Xue-Xin, and Alan A Stocker (2015) 'A bayesian observer model constrained by efficient coding can explain'anti-bayesian'percepts.' *Nature neuroscience* 18(10), 1509–1517

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

Zhang, Hang, and Laurence T Maloney (2012) 'Ubiquitous log odds: a common representation of probability and frequency distortion in perception, action, and cognition.' *Frontiers in neuroscience* 6, 1

Zhang, Hang, Xiangjuan Ren, and Laurence T Maloney (2020) 'The bounded rationality of probability distortion.' *Proceedings of the National Academy of Sciences* 117(36), 22024–22034

ONLINE APPENDIX

A Choice rule on orignal scale

In my preferred specification, I have assumed that the choice rule is transformed from its original scale by taking the log twice. This does in no way affect the conclusions. If we work on the original scale, the optimal choice rule is $exp(-t) > \frac{y}{x}$. We now derive the posterior expectation of the mental time delay t on the logarithmic scale just like above. To obtain the posterior expectation of the time delay t itself, we exploit the properties of the log-normal distribution, which has a mean $\mu + \frac{1}{2}\sigma_p^2$, where $\sigma_p^2 \triangleq \frac{\nu^2\sigma^2}{\nu^2+\sigma^2}$ is the posterior variance. We thus obtain

$$\mathbb{E}[\,t\,|\,r\,] = exp\left(\beta r + (1-\beta)\mu + \frac{1}{2}\sigma_p^2\right) = exp\left(\beta r + (1-\beta)(\mu + \frac{1}{2}\sigma^2)\right).$$

Substituting this mental quantity for t into a choice rule $exp(-\mathbb{E}[t\,|\,r]) > \frac{y}{x}$ and rearranging, we get $exp\left(-exp(\beta r + (1-\beta)\widehat{\mu})\right) > \frac{y}{x}$, where $\widehat{\mu} \triangleq \mu + \frac{1}{2}\sigma^2$. Taking the logarithm of both sides, multiplying by -1, and taking the logarithm again, yields $\beta r + (1-\beta)\widehat{\mu} < ln\left(-ln(\frac{y}{x})\right)$. The left-hand side of this equation can be used to derive a response distribution as in (8). The subsequent derivation follows the steps in the main text. The only difference from the choice rule presented in the main text is now in the definition of the impatience parameter, which takes the form $\tilde{\alpha} \triangleq exp\left(\frac{(1-\beta)}{\beta}\widehat{\mu}\right)$. Note, however, that the two models deliver predictions that are empirically indistinguishable. Indeed, the two models fit to the same data show a difference in predictive fit that is exactly equal to 0.

B 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. 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: $r_z \sim \mathcal{N}(ln(z), \nu^2)$, $ln(z) \sim \mathcal{N}(\mu_o, \sigma_o^2)$,

where I assume a common coding noise ν for reasons of empirical identifiability. Defining $\gamma \triangleq \frac{\sigma_o^2}{\sigma_o^2+\nu^2}$, we obtain the posterior expectation $E[ln(z)|r_z] = \gamma r_z + (1-\gamma)\mu_o$. Deriving the common response distribution for the weighted difference in signals and normalizing, like done in the main text, we obtain the following stochastic choice rule:

$$Pr[(x,\tau_\ell) \succ (y,0)] = \Phi\left(\frac{\gamma \times ln\left[-ln\left(\frac{y}{x}\right)\right] - \beta \times ln(\alpha\, t)}{\nu\sqrt{\beta^2+\gamma^2}}\right), \tag{16}$$

where I have dropped the time-subscripts for simplicity's sake.

Notice that the addition of outcome transformations now affects mostly the model parameters α and β. That is, one could rewrite the numerator of the expression in parentheses as $ln\left[-ln\left(\frac{y}{x}\right)\right] - \frac{\beta}{\gamma} \times ln(\alpha t)$ (as long as one rewrites the denominator as $\frac{\nu}{\gamma} \times \sqrt{\beta^2+\gamma^2}$). One can then see that the addition of a parameter $\gamma \leq 1$ would simply serve to reinforce the conclusions already reached in the main text. If γ were to be found to vary in the reward magnitude, however, this additional parameter could account for the absolute magnitude effect—a topic that is beyond the scope of this paper.

C Generalization to other types of time tradeoffs

In the main text, I have discussed simple tradeoffs between a smaller-sooner and larger-later option. Here I show that the model can be generalized to more complex tradeoffs in a straightforward way. To avoid notational overload, I will use a static representation that drops the indices for the round in which a stimulus is observed.

Take for instance choice options having payouts at different time horizons, such as the ones used by Abdellaoui, Bleichrodt, L'Haridon and Paraschiv (2013) to disentangle discounting from inter-temporal utility. A sooner reward y, τ_s is then traded off against a stream of rewards $(x, \tau_\ell; z, \tau_s)$, where $x > y > z$. This can be accommodated by replacing the choice rule in (2) with the following expression:

$$ln\left(\frac{p}{q}\right) > ln\left(\frac{y-z}{x}\right),$$

with all other quantities defined as in the main text. Any additional payments occurring with the same two delays can be handled similarly.

The model can furthermore be generalized to any time tradeoffs following the setup in Appendix A of Vieider (2023) for risk. Take a stream of rewards $\boldsymbol{x} = (x_1, ..., x_n)$, where subscripts index different time delays. Now take a separate stream $\boldsymbol{y} = (y_1, .., y_n)$. Let us assume for argument's sake that the latter has more smaller amounts obtaining at sooner points in time. The generalization of the choice rule in (2) entails that the larger-later stream will be chosen whenever:

$$\sum_{j=2}^{n} \frac{p_j}{p_1} \frac{(x_j - y_j)}{(y_1 - x_1)} > 1,$$

where $p_j \triangleq exp(-\tau_j)$. This equation takes the form of an optimal choice rule used in signal detection theory (Green, Swets et al., 1966). The reference event in the denominator is arbitrary (just like in multinomial logit models), and is here taken to be the first payoff. Note, however, that all pairwise comparisons are enshrined within this equation, as can be seen by taking the ratio of the two terms one wants to compare in the numerator. That is, subjects do not necessarily need to compare everything to the earliest payoff difference for the choice rule to be applicable, and more natural comparison points may well be provided by delays with particularly salient reward differences.

I will further assume that each time delay in the equation above is logged and evaluated in parallel within a neural network. That implies that there will be a comparison between the time delays and the objective reward-difference ratio just like the one discussed in the main text. For an arbitrary comparison j, we will thus observe the following likelihood functions:

$$r_j \sim \mathcal{N}\left(ln\left(t_j\right), \nu^2\right)$$

where $t_j \triangleq \tau_j - \tau_1$ for $\tau_1 > 0$, and $t_j \triangleq \tau_j + \psi$ if $\tau_1 = 0$. We can now follow the same steps as in the main text to derive the solution to the general tradeoff, first deriving the response distribution for each individual time delay, and from there obtaining

the stochastic choice rule. This yield the following stochastic choice rule:

$$Pr[\boldsymbol{x} \succ \boldsymbol{y}] = \Phi\left[\frac{1}{\sqrt{n-1}\,\beta\,\nu}\sum_{j=2}^{n} ln\left(-\mathbb{1} ln\left(\frac{\mathbb{1}(x_j - y_j)}{y_1 - x_1}\right)\right) - \beta\, ln(\alpha\, t_j)\right],$$

where $\mathbb{1}$ is an indicator taking the value of 1 if $x_j - y_j > 0$, and the value of -1 otherwise. The equation as written above makes two additional assumptions. The first is that $y_1 - x_1 > 0$, as will typically be the case based on the assumption that $\boldsymbol{y}$ contains the smaller-sooner amounts, so that the first amount will typically be larger than for the larger-later option. The second is that any $x_j - y_j < y_1 - x_1$. Note, however, that neither of these two assumptions is necessary for the equation to be applicable—all that is needed in case one or the other assumption is violated is some additional indexing to ensure that the equation is well-behaved.

D Derivations and proofs

Combination of signal and prior in (7)

The likelihood and prior take the following form:

$$p(r \mid t) \propto exp\left(-\frac{1}{2\nu^2}\left(r - ln(t)\right)^2\right)$$
$$p(ln(t)) \propto exp\left(-\frac{1}{2\sigma^2}\left(ln(t) - \mu\right)^2\right)$$

We can log the distributions to transform the multiplication into an addition:

$$\begin{aligned} ln\left(p[ln(t)|r]\right) &= ln[p(r|t)] + ln[p(ln(t))] \\ &= -\frac{1}{2\nu^2}\left(r^2 - 2\,r\,ln(t) + [ln(t)]^2\right) - \frac{1}{2\nu^2}\left([ln(t)]^2 - 2\,\mu\,ln(t) + \mu^2\right) \\ &= -\frac{1}{2}\left(\frac{1}{\nu^2} + \frac{1}{\sigma^2}\right)[ln(t)]^2 + ln(t)\left(\frac{r}{\nu^2} + \frac{\mu}{\sigma^2}\right) - \left(\frac{r^2}{2\nu^2} + \frac{\mu^2}{2\sigma^2}\right) \end{aligned}$$

We know that the log-posterior distribution will take the following form:

$$ln\left(p[ln(t)|r]\right) = -\left(\frac{[ln(t)]^2}{2\sigma_p^2} - \frac{2\,ln(t)\,\theta}{2\sigma_p^2} + \frac{\theta^2}{2\sigma_p^2}\right),$$

where $\theta \triangleq \mathbb{E}[ln(t)|r]$ is the posterior mean, and σ_p^2 is the posterior variance. We can thus complete the square by matching the first two expressions in the sums of the last two equations above. We start from the first:

$$-\frac{[ln(t)]^2}{2\sigma_p^2} = -\frac{1}{2}\left(\frac{1}{\nu^2} + \frac{1}{\sigma^2}\right)[ln(t)]^2$$
$$\sigma_p^2 = (\lambda + \xi)^{-1},$$

which is the variance of the posterior in (7), and where $\lambda \triangleq \nu^{-2}$ and $\xi \triangleq \sigma^{-2}$.

We can now match the second element:

$$\frac{2\, ln(t)\, \theta}{2\sigma_p^2} = ln(t)\left(\frac{r}{\nu^2} + \frac{\mu}{\sigma^2}\right)$$
$$\theta = \frac{\lambda\, r + \xi\, \mu}{\lambda + \xi} = \frac{\lambda}{\lambda + \xi} r + \frac{\xi}{\lambda + \xi}\mu,$$

which is the posterior mean from (7). QED.

The response distribution in (8)

Let $z \sim \mathcal{N}(\widehat{z},\, \sigma_z^2)$. From the properties of the normal distribution, we know that $a + b\, z \sim \mathcal{N}(a + b\, \widehat{z},\, b^2\, \sigma_z^2)$. The result in (8) follows by substituting $a = (1 - \beta)\, \mu$, $z = r$, $\widehat{z} = ln(t)$, $b = \beta$, and $\sigma_z = \nu$.

The dynamic equations in (13) and (12)

We start from the following joint-inference problem, concentrating on a single data point r. For the derivation, I will simplify notation by starting from the hyperparameters, indexed by 0, and updating them to the posterior indexed by 1:

$$p(\theta, \mu, \xi|r) = p(\mu, \xi|r) \times p(\theta|r, \mu, \xi).$$

Dropping dependency on the data in the notation to avoid clutter, we can further factorize $p(\mu, \xi) = p(\xi) \times p(\mu|\xi)$. This makes it clear that the prior for mean and precision ought to be modelled through a normal-gamma distribution, which takes into account the interdependence of μ and ξ conditional on the data. Starting from

the parameters of the hyperprior, subscripted by 0, we can then update the prior using a single observation of the inferred time delay, θ_1:

$$\begin{aligned} p(\theta, \mu, \xi|r) &= p(\xi) \times p(\mu|\xi) \times p(\theta|r, \mu, \xi) \\ &\propto \xi^{\zeta_0 - 1} \exp(-\rho_0 \xi) \times \sqrt{\xi}\, exp\left(-\frac{\kappa_0 \xi}{2}(\mu - \widehat{\mu}_0)^2\right) \times \sqrt{\xi} \exp\left(-\frac{\xi}{2}(\theta_1 - \mu)^2\right) \\ &\propto \xi^{\zeta_0} \exp\left(-\frac{\xi}{2}\left[(\kappa_0 + 1)\mu^2 - 2\mu(\kappa_0 \widehat{\mu}_0 + \theta_1) + (\kappa_0 \widehat{\mu}_0^2 + \theta_1^2 + 2\rho_0)\right]\right). \end{aligned}$$

I now write an equation using the quantities we expect to find in the posterior. Starting from the parameters of the hyperprior, as done above, we can denote the posterior after the first observation by subscripting the parameters by 1:

$$\xi^{\zeta_1 - 1/2} \exp\left(-\frac{\xi}{2}\left[\kappa_1 \mu^2 - 2\mu\kappa_1 \widehat{\mu}_1 + \kappa_1 \widehat{\mu}_1^2 + 2\rho_1\right]\right).$$

We can now simply match the quantities in the last two equations. Matching the powers of ξ immediately gives us the updating equation for the shape parameter:

$$\zeta_1 = \zeta_0 + \frac{1}{2},$$

and matching the first element in the exponential tells us that

$$\kappa_1 = \kappa_0 + 1.$$

Next I derive the updating equation for the prior mean in (12). To do this, I match the second, middle element in the exponential of the two equations:

$$\begin{aligned} \kappa_1 \mu_1 &= (\kappa_0 \mu_0 + \theta_1) \\ \widehat{\mu}_1 &= \frac{(\kappa_0 \widehat{\mu}_0 + \theta_1)}{\kappa_0 + 1} \\ &= \widehat{\mu}_0 + \frac{1}{\kappa_0 + 1}(\theta_1 - \widehat{\mu}_0), \end{aligned}$$

where the penultimate step divides both sides by $\kappa_1 = \kappa_0 + 1$, and the last step exploits that $\widehat{\mu}_0 - \frac{\kappa_0}{\kappa_0 + 1} = \frac{1}{\kappa_0 + 1}\widehat{\mu}_0$.

The last step will be to derive the updating equation for the rate parameter of the marginal Gamma distribution. To this end, we can equalize the third terms in the exponents of the equations above:

$$\begin{aligned}
\kappa_1\mu_1^2 + 2\rho_1 &= \kappa_0\mu_0^2 + \theta_1^2 + 2\rho_0 \\
\kappa_1\frac{(\kappa_0\mu_0+\theta)^2}{\kappa_1^2} + 2\rho_1 &= \kappa_0\mu_0^2 + 2\rho_0 + \theta^2 \\
\rho_1 &= \rho_0 + \frac{1}{2}\left[\kappa_0\mu_0^2 + \theta_1^2 + \frac{\kappa_0^2\mu_0^2 + 2\kappa_0\mu_0\theta_1 + \theta_1^2}{\kappa_0+1}\right] \\
&= \rho_0 + \frac{1}{2}\frac{\kappa_0}{\kappa_0+1}\left(\mu_0^2 - 2\kappa_0\mu_0\theta_1 + \theta_1^2\right) \\
&= \rho_0 + \frac{1}{2}\frac{\kappa_0}{\kappa_0+1}\left(\theta_1 - \mu_0\right)^2.
\end{aligned}$$

The equations in the main text result from these equations by replacing the subscript 0 with i and the subscript 1 by $i+1$. Q.E.D.

Generalization to repeated stimuli

It is instructive to examine the learning equations derived above to a context where an identical time delay t is repeated n times. This will allow us to derive an expression defined entirely in terms of quantities observed by the experimenter. I now derive such a formulation:

$$\begin{aligned}
p(\theta,\mu,\xi|t) &= p(\xi)\times p(\mu|\xi)\times p(\theta|t,\mu,\xi) \\
&\propto \xi^{\zeta_0-1}\exp(-\rho_0\xi)\times\sqrt{\xi}\,exp\left(-\frac{\kappa_0\xi}{2}(\mu-\widehat{\mu}_0)^2\right)\times\sqrt{\xi}\exp\left(-\frac{\xi}{2}\sum_{i=1}^{n}(\theta_i-\mu)^2\right) \\
&\propto \xi^{\zeta_0}\exp\left(-\xi\left(\frac{1}{2}\sum_{i=1}^{n}(\theta_i-\widehat{\theta})\right)\right)\exp\left(-\frac{\xi}{2}\left(\kappa_0(\mu-\widehat{\mu}_0)^2 + n(\widehat{\theta}-\mu)^2\right)\right),
\end{aligned}$$

where the last line obtains from setting $\theta_i - \mu = \theta_i - \widehat{\theta} + \widehat{\theta} - \mu$, and where $\widehat{\theta} = \beta\, ln(t) + (1-\beta)\,\widehat{\mu}_0$ as derived above.

The derivation now proceeds as above by multiplying out the squared terms, grouping parameters together, and finally completing the square. We thus obtain

the following updating equations:

$$
\begin{aligned}
\zeta_n &= \zeta_0 + \frac{n}{2} \\
\kappa_n &= \kappa_0 + n \\
\widehat{\mu}_n &= \widehat{\mu}_0 + \frac{n}{\kappa_0 + n}(\widehat{\theta} - \widehat{\mu}_0) \\
&= \widehat{\mu}_0 + \frac{n}{\kappa_0 + n}\beta_0 \ (ln(t) - \widehat{\mu}_0) \\
\rho_n &= \rho_0 + \frac{1}{2}\frac{\kappa_0 n}{\kappa_0 + n}(\widehat{\theta} - \widehat{\mu}_0)^2 + \frac{1}{2}\sum_{i=1}^{n}(\theta_i - \widehat{\theta})^2 \\
&= \rho_0 + \frac{1}{2}\frac{\kappa_0 n}{\kappa_0 + n}\,\beta_0 \ (ln(t) - \widehat{\mu}_0)^2 + \frac{n}{2}\,\beta_0^2\,\nu^2,
\end{aligned}
$$

where the substitution in the last term of the last line is based on the definition of coding noise as $\nu^2 \triangleq \frac{1}{n}\sum_{i=1}^{n}(r_i - ln(t_i))^2$.

These equations defined in terms of quantities observable for the experimenter further support the conclusions discussed in the main text based on the learning equations of the mind. The updating equations of the first three parameters above look much like before, except that they are now formulated in terms of n random draws of a signal r for an identical stimulus. The updating equation for the rate parameter of the Gamma, however, has the virtue of further decomposing the squared prediction error in the main text. In particular, we can now see that for repeated signal for an identical stimulus, the sum of squares will add two separate parts—one based on the deviation of the true logged time delay from the mean of the prior, and one based on the deviation of a single posterior inference from the mean of the inferences for an identical stimulus. The latter can furthermore be reformulated in terms of coding noise. This shows that the posterior precision will be inversely proportional to the variation of stimuli in the environment, *and* to the noise with which single observations will be encoded.

The stochastic choice rule in (14)

Let $z \sim \mathcal{N}(\widehat{z},\, \sigma_z^2)$. From the properties of the normal distribution, we know that $c - a - b\,z \sim \mathcal{N}(c - a - b\,\widehat{z},\, b^2\,\sigma_z^2)$. Substituting (9) into (3), rearranging, substituting

$c = ln\left[-ln\left(\frac{y}{x}\right)\right]$, $a = ln(\alpha)$, $z = r$, $b = \beta$, $\widehat{z} = ln(t)$, $\sigma_z = \nu$, and normalizing by the standard deviation gives us (14).

E Econometrics

The noisy coding of time model is inherently stochastic, so that it can be directly implemented without any further need for separate assumptions about the error structure. 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 $p_i(\boldsymbol{\theta}_n|z)$ given choice data z over the individual parameter vector $\boldsymbol{\theta}_n$ from

$$p_i(\boldsymbol{\theta}_n|z) \propto p(z|\boldsymbol{\theta}_n) \times p(\boldsymbol{\theta}_n), \tag{17}$$

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:

$$p(z|\boldsymbol{\theta}_n) = (Pr[(x, \tau_\ell) \succ (y, \tau_s)])^z \times (1 - Pr[(x, \tau_\ell) \succ (y, \tau_s)])^{1-z}, \tag{18}$$

with $Pr[(x, \tau_\ell) \succ (y, \tau_s)]$ taking the form of the choice probability in (14).

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

$$p(\boldsymbol{\theta}_n) = \mathcal{N}(\overline{\boldsymbol{\theta}}, \boldsymbol{\Sigma}), \tag{19}$$

where $\overline{\boldsymbol{\theta}}$ is a vector containing the aggregate parameter means, and $\boldsymbol{\Sigma}$ is a covariance matrix of the individual-level parameters. Both $\overline{\boldsymbol{\theta}}$ and $\boldsymbol{\Sigma}$ are endogenously estimated from the aggregate data, and serve as priors for the individual-level es-

timates. The hyperpriors for the parameters in $\overline{\boldsymbol{\theta}}$ and $\boldsymbol{\Sigma}$ are chosen 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 reportedare 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).

F Classroom experiment

F.1 Experiment

The classroom experiment is meant to serve as a proof of concept. I pursue two main goals. One, I want to nonparametrically quantify both the patterns predicted by the model (present bias, delay-dependence), and those not predicted by it (strongly decreasing impatience). While none of these patterns are new, they have rarely been presented together and based on one and the same experiment. Some are further disputed, so that establishing these patterns is an important first step. Two, I want to create an incentivized benchmark to which to compare the non-incentivized online experiments. As can be seen from the patterns reported below, the qualitative patterns across subject pools and incentivization protocols are indeed identical.

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 discounting yet. 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, as well as for research 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 9. They were chosen to allow for optimal identification of all model components. 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 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 with AC, of CD and DE with CE, and of all the 6 week delays (AB, BC , CD , DE) and 12 week delays (AC , CE) with

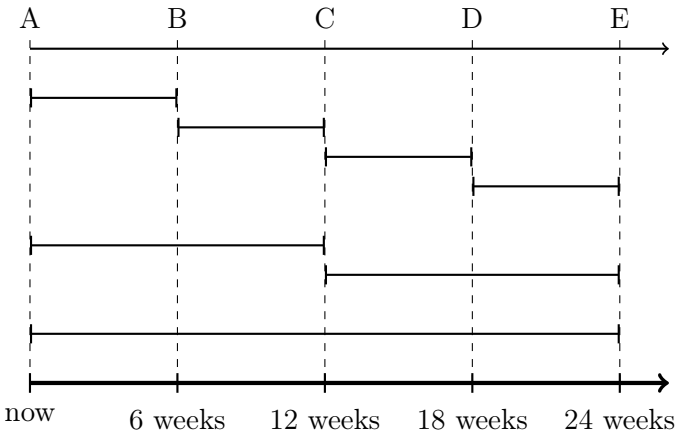

Figure 9: Time horizons used in experiment

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. Notice that, given 4 delays of 6 weeks, 2 of 12 weeks, and 1 of 24 weeks, the stimuli are well-fitted by a log-normal distribution.

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.

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 B for a discussion of robustness to this assumption). The earlier amounts ranged between €33 and €49 inclusive in steps of €1. Each screen presented one single choice in random order. Subjects made 158 choices, which took 20 minutes on average. There were 40 repeated choices. Identification of decision noise, which plays a central role in the model, is thus assured by i) repeated observations of the same stimuli; and ii) monotonicity violations between similar stimuli.

Before the start of the experiment, the lecturer presented the instructions, to make sure that everybody had an understanding of the tasks. The lecturer 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, 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 was signed by the lecturer. The certificate also contained the address and telephone number of the lecturer, and students were encouraged to get in touch 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.

F.2 Results

I start from an examination of the nonparametric evidence , shown in figure 10. Panel A plots decumulative choice proportions for the larger-later reward as the sooner-smaller reward increases from €33 through €49. 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 are even stronger, and again indicate an increase in patience following the introduction of the up-front delay ($p \ll 0.001$; figure in online appendix **??**). Further comparing 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) yields no evidence for strongly decreasing impatience. 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.

Panel B in figure 10 examines delay-dependence by identifying a probabilistic non-parametric discount factor from the sooner amount at which a subject starts

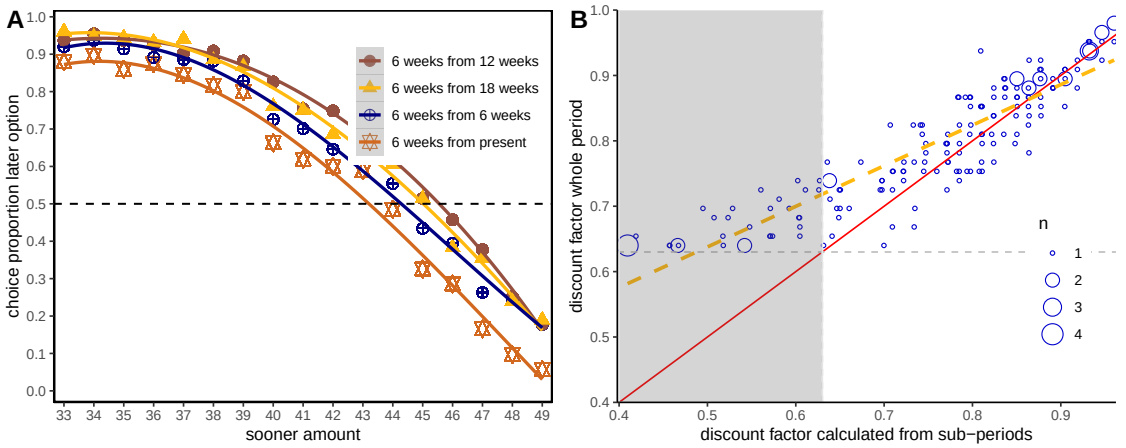

Figure 10: Non-parametric illustration of present-bias and delay-dependence

Panel A compares decumulative choice proportions for 6 week delays from different initial times. Panel B plot the discount factor for a 12 week delay from the present, $\delta_{0,12}$, against the discount factor calculated based on the two underlying 6-week discount factors, $\widehat{\delta}_{0,12} = \delta_{0,6} \times \delta_{6,12}$.

choosing the smaller-sooner reward 50% of the time, divided by the later amount. The particular case shown 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, indicating delay-dependence. Importantly, this also holds true when we exclude calculated discount factors $\delta_{0,6} \times \delta_{6,12}$ smaller than 0.64, which could otherwise bias the findings due to censoring effects ($p \ll 0.001$). This effect is even stronger for the longest 24-week delay (AE) against the product of the discount factors for the 2 underlying 12-week periods (AC and CE), and is 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 0.64 ($p \ll 0.001$).

Figure 11 further shows the comparison of decumulative choice proportions for the delay AC compared to the delay CE, i.e. for 12 week delays where the first delay occurs from the present, and the second from an upfront delay of 12 weeks.

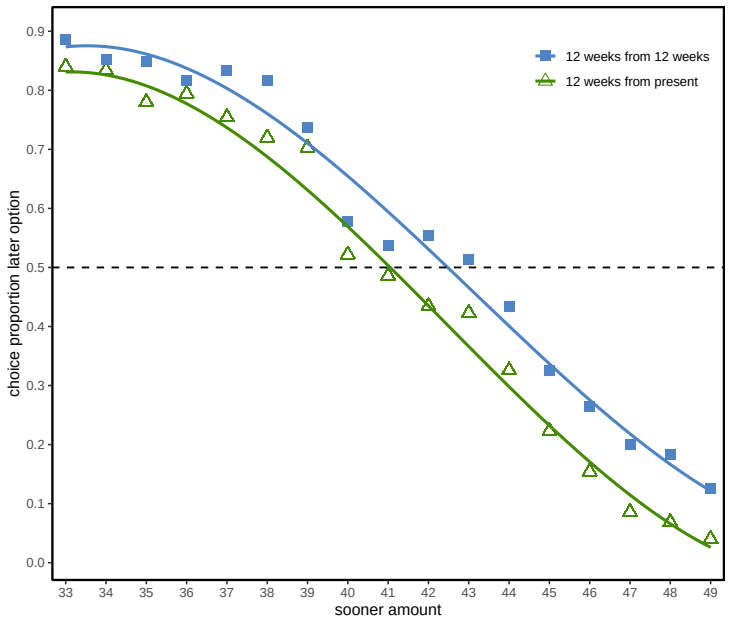

Figure 11: present-bias for 12 week delays—additonal comparisons

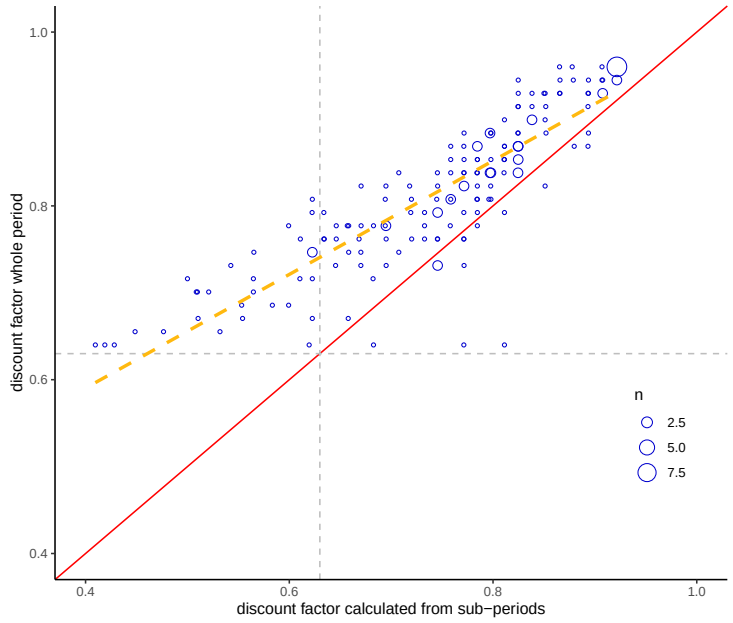

(a) subadditivity, 12 weeks from 12 weeks

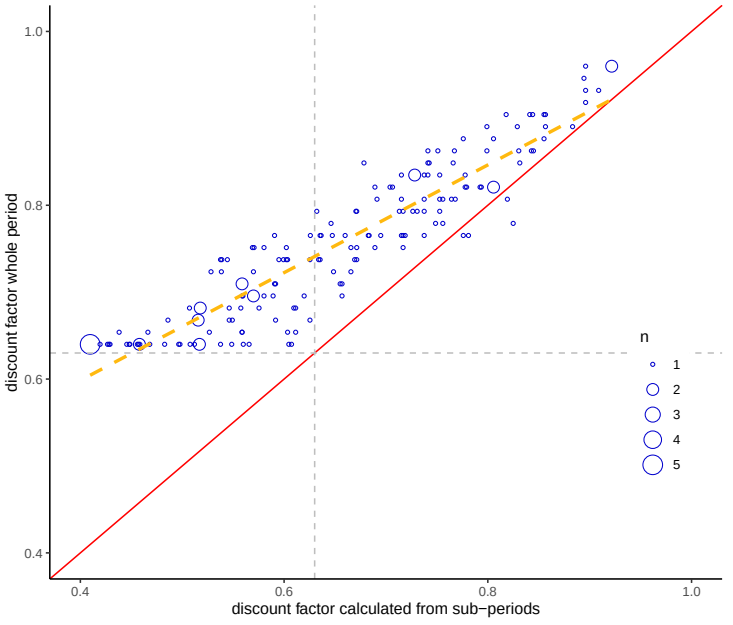

(b) subadditivity, 24 weeks vs 12-week delays

Figure 12: 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 12(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 12(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.