

WORKING PAPER

NOISY CODING OF TIME AND REWARD DISCOUNTING

Ferdinand M. Vieider

December 2021
2021/1036

Noisy coding of time and reward discounting*

Ferdinand M. Vieider^{†1,2}

¹*RISL $\alpha\beta$, Department of Economics, Ghent University*

²*RISL $\alpha\beta$ Africa, University Mohammed VI Polytechnic*

December 16, 2021

Abstract

I present a model generating delay-discounting from noisy mental representations of time delays. The optimal combination of noisy signals about time delays with prior information results in a stochastic model predicting discounting to exhibit present-bias, but to be stationary and delay-dependent once up-front delays are introduced. The derivation from an optimal encoding-decoding process sidesteps arbitrariness concerns voiced about earlier models. Data collected in an experiment support the need for separate but interacting parameters to capture present-bias and delay-dependence. The account explains why non-trivial discounting is routinely observed in experiments using monetary rewards instead of consumption.

1 Motivation

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

*I acknowledge financial support from BOF at Ghent University under the project “The role of noise in the determination of preferences”. The experiment has received ethics approval from the ethics committee at the Faculty of Economics and Business, Ghent University, with number *FEB 2021 M*. I am indebted to Ranoua Bouchouicha for help in executing the experiment, and to Benjamin Enke, Xavier Gabaix and Olivier l’Haridon for helpful comments and discussions about earlier versions of this manuscript. All errors remain my own.

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

length of the delay used to measure it (Thaler, 1981; Read, 2001; Dohmen, Falk, Huffman and Sunde, 2017). Psychologists and economists have proposed a variety of functional forms to model such discounting patterns (Mazur, 1987; Loewenstein and Prelec, 1992; Laibson, 1997; Read, 2001; Ebert and Prelec, 2007). It has been remarked, however, that the functional forms proposed do not directly follow from the observed behaviour, and that there is thus a high degree of arbitrariness in such modelling choices (Rubinstein, 2003).

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

The deductive nature of the model contrasts with the inductive process underlying the derivation of most extant discounting functions, which were generally developed to fit aggregate choice data. The derivation of the functional forms from an optimal choice rule implemented through a stochastic neural encoding-decoding process drastically reduces the degrees of freedom in deriving the functions, thus addressing the arbitrariness critique of Rubinstein (2003). Traditional functional forms either confound decreasing impatience and present-bias with delay-dependence, so that one and the same parameter needs to describe both phenomena (Mazur, 1987; Ebert and Prelec, 2007), or focus on one aspect

while neglecting the other (Read, 2001; Laibson, 1997). Functions that can account for both phenomena tend to exhibit a multiplicity of independent parameters (Scholten and Read, 2006; 2010). The noisy coding model, on the other hand, generates predictions of present-bias and delay-dependence based on separate mechanisms, which are however closely intertwined. Its predictions rest on a parsimonious parameter space, and the theoretically-derived interactions between the parameters put tight bounds on their values.

I test the predictions emerging from the model in an experiment geared at detecting the predicted patterns, as well as patterns that are *not* predicted by the model, such as the strongly decreasing impatience entailed by models of generalized hyperbolic discounting. The data provide clear evidence for present bias when the sooner delay is truly immediate, as has been documented in other recent contributions (Balakrishnan, Haushofer and Jakiela, 2020; Imai et al., 2021). There is also strong evidence for delay-dependence, with measured patience increasing systematically in the length of the delay used to measure it. I do, however, find no evidence for strongly decreasing impatience, i.e. once an initial upfront delay has been introduced, impatience does not decrease further as this upfront delay is increased. Testing the predictive performance of the model against standard models augmented by an additive error term, I find the noisy coding model to significantly outperform those models. I can trace this superior performance to the parameter interactions imposed by the theoretical predictions.

The model I present is interesting for several reasons. The origin of discounting patterns in noisy perception of time serves to explain why deviations from optimal discounting are routinely observed in experiments using monetary amounts, instead of only occurring for consumption as predicted by economic theory (Frederick, Loewenstein and O'Donoghue, 2002; Cubitt and Read, 2007). The patterns I document also have policy implications. For one, the model suggests that a welfare maximizing policy maker can safely disregard deviations from exponential discounting, since such deviations are due to the noisy perception of choice parameters rather than to true underlying preferences. The results suggest furthermore

that strong discounting for short time delays can be counteracted by measures that reduce noisiness in time perception by either facilitating the objective perception of time delays, or increasing attention towards a decision problem.

The noisy coding model I present is based on an influential theoretical paradigm from neuroscience, according to which the human mind functions like a Bayesian inference machine (Knill and Pouget, 2004; Doya, Ishii, Pouget and Rao, 2007; Vilares and Kording, 2011). In a context where cognitive resources are scarce, such a system may indeed be optimal, in the sense that it can represent wide ranges of sensory information using a contained amount of resources. Applied to human decision-making, the paradigm holds the promise of unifying many well-known stylized facts, which have been modelled using a variety of disparate ad-hoc models. For instance, noisy coding of rewards may result in small-stake risk aversion (Khaw, Li and Woodford, 2020), and optimal adaptation of coding noise to the expected outcome distribution can explain context-dependence in risky choice (Frydman and Jin, 2021). Noisy coding can generate functionals such as described by prospect theory, while at the same time explaining common violations of prospect theory (Vieider, 2021).

The model I present is related in spirit to the model of Gabaix and Laibson (2017). The latter show that if future *utilities* are perceived with noise, and if that noise linearly increases in the time delay from the present, then noisy perception of future utilities may lead to hyperbolic as-if preferences. Such a model, however, cannot separate present-bias from delay-dependence, and results in a prediction of strongly decreasing impatience (i.e., preferences deviate from stationarity even once upfront delays are introduced). This contrasts with the predictions of the noisy coding model, where immediacy effects and delay-dependence emerge from separate mechanisms. Recent empirical findings indeed suggest that delay-dependence is related to cognitive uncertainty while present bias is not (Enke and Graeber, 2021), thus supporting the predictions emerging from the noisy coding model I present. The empirical results I report suggest that such separate mechanisms are indeed needed to properly account for the observed decision patterns,

thus considerably increasing the predictive performance of the model.

The model is less directly related to accounts that have derived predictions of time discounting based on randomness in choices and preferences (Lu and Saito, 2018; He, Golman and Bhatia, 2019). The model furthermore presents some commonalities with accounts of time discounting based on the fundamental uncertainty of the future (Halevy, 2008; Epper and Fehr-Duda, 2018). These models, however, differ from the approach taken here in both their formal approach and in their predictions. I will return to these differences in the discussion.

2 The noisy coding of time model

2.1 Encoding and decoding of time delays

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

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

where ρ is an evolutionarily optimal discount rate of about 2% per year (Rogers, 1994), which we may think of as a constant for all practical purposes.¹ The evolutionary discount rate ρ ensures that the choice rule is optimal, in the sense that it results in low exponential discount rates by rescaling the objective time delay τ . Note, however, that we could equally envision a ‘natural’ discount rate

¹Given the long-term, evolutionary nature of this discount rate, it will be approximately equal across individuals. The main systematic variation we ought to observe is for different age groups, with younger individuals less patient than older individuals. This variation will, however, be small relative to the inter-individual variation that is likely to arise based on the framework derived here in sequence, so that for all practical purposes we may think of ρ as a constant.

while omitting ρ , without affecting the analysis to follow in any way. I will thus omit ρ from the further derivations for parsimony’s sake.

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

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

where $E[\ln(t)|r]$ is the expectation of the posterior distribution of the time delay conditional on the noisy mental signal r , and where I substituted $t \triangleq 1 + \tau$ for the objective delay τ in order to prevent the numerical representation to become boundless as $\tau \rightarrow 0$, which would be hard to reconcile with the resource-saving rationale underlying the logarithmic representation (Petzschner and Glasauer, 2011). While this transformation is inconsequential for large τ , it will have important implications as $\tau \rightarrow 0$.

Neuroscientific evidence suggests that numbers are mentally represented on a log scale to avoid an unbounded increase in the cognitive resources necessary to represent larger numbers (Dehaene and Changeux, 1993; Dayan and Abbott, 2001; Dehaene, 2003). It has further been shown that the dependence of subjectively perceived delays on objective delays may be best fit by a logarithmic function (Zauberman, Kim, Malkoc and Bettman, 2009). I thus model the mental signal

as a single draw from the distribution:

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

where the parameter ν quantifies coding noise, and is assumed to be independent of τ , resulting in a Weber-Fechner law for time perception (Fechner, 1860).

To obtain the posterior expectation underlying the mental choice rule in equation 2, the noisy signal r needs to be combined with a prior indicating the probability of different time delays in the environment. It seems natural to assume a conjugate prior from the normal family, so that:

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

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

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

where $\beta \triangleq \frac{\sigma^2}{\sigma^2 + \nu^2}$ and $\alpha \triangleq e^{(1-\beta)\mu}$. The parameter β governs the degree to which we may expect regression of the posterior mean to the mean of the mental prior, μ (to see this more clearly, one can rewrite equation 5 as $E[\ln(t)|r] = \beta(r - \mu) + \mu$). Unexpected signals falling far from the mean of the prior will thus be discounted more heavily than expected signals. The degree of discounting will be determined by the uncertainty connected to the signal, captured by ν , relative to the uncertainty surrounding the prior mean itself, as captured by σ . The wider the prior distribution and the lower the encoding noise, the larger the weight attributed to the signal will be relative to the weight attributed to the prior mean.

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

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

This threshold equation fully describes the condition under which the larger-later amount will be chosen. To result in testable predictions, however, we need to obtain a probabilistic choice rule devoid of the mental signal r , which cannot be observed by the econometrician. Obtaining the z-score of the likelihood distribution for r from the likelihood in equation 3 and comparing it to the z-score of the threshold equation (see online appendix A for details), we obtain the following probabilistic choice rule:

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

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

2.2 Noisy coding of time delays generates time insensitivity

Leaving decision noise aside momentarily, we can zoom in on the indifference condition underlying equation 7:

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

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

$$\delta_\tau = \exp(-\alpha t^\beta). \quad (9)$$

The right-hand side of this expression corresponds to the constant-sensitivity discount function proposed by [Ebert and Prelec \(2007\)](#). The time-discriminability parameter β is closely connected to the noisiness of time perception in the present setup. This interpretation is highly consistent with empirical evidence showing that time sensitivity can be improved by presenting visual clues for time duration, while it will decrease under time pressure ([Ebert and Prelec, 2007](#)), thus making the parameter susceptible to policy interventions.

The parameter α is mostly related to expected delays in a given environment. It is, however, indirectly influenced by noise in perception as well. Given the definition of the encoded time delay as $t \triangleq 1 + \tau$, the function will show a discontinuous drop at $\tau = 0$ of a size that depends on the parameter α . That is, $\lim_{\tau \rightarrow 0} \delta(t) = \exp(-\alpha)$, where $\delta(t)$ indicates the discount function, and where I assume as usual that $\delta(0) = 1$, so that immediate payments are not discounted. This introduces a drop in the vicinity of $\tau = 0$ which predicts present-bias—a preference for immediately received rewards over even slightly delayed rewards—which is itself proportional to the discount rate. This, in turn, means that changes in time discriminability β will also affect the degree of present-bias through $\alpha \triangleq e^{(1-\beta)\mu}$.

2.3 Discounting with up-front delays

It is important to note that the choice rule being implemented will depend on the type of stimuli faced by the decision-maker. The question of how behaviour may change depending on the choice stimuli is thus central in the present setup, given the *predictive* interpretation of the model. Take a tradeoff between delayed outcomes (y, τ_s) and (x, τ_ℓ) , where the subscripts s and ℓ stand for *sooner* and *later*, respectively. Further define mentally coded delays $s \triangleq 1 + \tau_s$ and $\ell \triangleq 1 + \tau_\ell$. The optimal choice rule now takes the form:

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

which corresponds to the following equation in log-log space:

$$\ln(\tau_\ell - \tau_s) < \ln \left[-\ln \left(\frac{y}{x} \right) \right] - \ln(\rho). \quad (11)$$

This choice rule immediately suggests to mentally transform the time *delay* inside the logarithm on the left, which we can define as $d \triangleq \ell - s = \tau_\ell - \tau_s$, since no analytic solution is possible for the transformation of the individual time delays within the log (this is not unique to the use of the log-log rule; see online appendix C for details). Deriving the mental representation of this choice rule as before (online appendix C), we obtain the following discount factor:

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

This discount factor captures delay-dependence, since the delay between payments is compressed by $\beta < 1$. Delay-dependence aside, however, the function is stationary. Indeed, the function is identical for any delay d independent of the value taken by the up-front delay s . In other words, only the time *delay* between the two payoffs matters, while the up-front delay s is inconsequential. For two time delays τ_s and τ_ℓ from the present, on the other hand, the resulting discount function would take the form $\exp(-\alpha(\ell^\beta - s^\beta))$.

The fact that this function obtains naturally from an intuitive optimal choice rule may explain the pervasiveness of subadditivity (Dohmen et al., 2017), given the model's *predictive* interpretation. Notice that while Ebert and Prelec (2007) also ascribed a causal and hence predictive interpretation to their constant-sensitivity function (p. 1436), the absence of differentiation by stimulus type in their setup confounds delay-dependence with strongly decreasing impatience, since both are captured by one and the same parameter. Although the empirical evidence on strongly decreasing impatience is inconclusive (Scholten and Read, 2010), one may want to ask under what conditions such a discount function *could* be observed. It turns out that such a prediction could be obtained in the present framework only based on a choice rule trading off the relative time delay against the relative cost

of waiting, such as the following:

$$\exp\left(-\rho \times \frac{\tau_\ell}{\tau_s}\right) > \frac{x - y}{x}. \quad (13)$$

Such a choice rule seems rather far-fetched a priori, and we will indeed see that the predictions emerging from a mental transformation of this choice rule—shown in online appendix [D](#)—provides a terrible fit to the data I will present below.

3 Experiment and analysis

3.1 Experiment

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

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

CD , DE) and 12 week delays (AC , CE) with the full delay over 24 weeks (AE) allow for the identification of delay-dependence. Generalized hyperbolicity can be identified from the comparison of BC, CD, and DE.

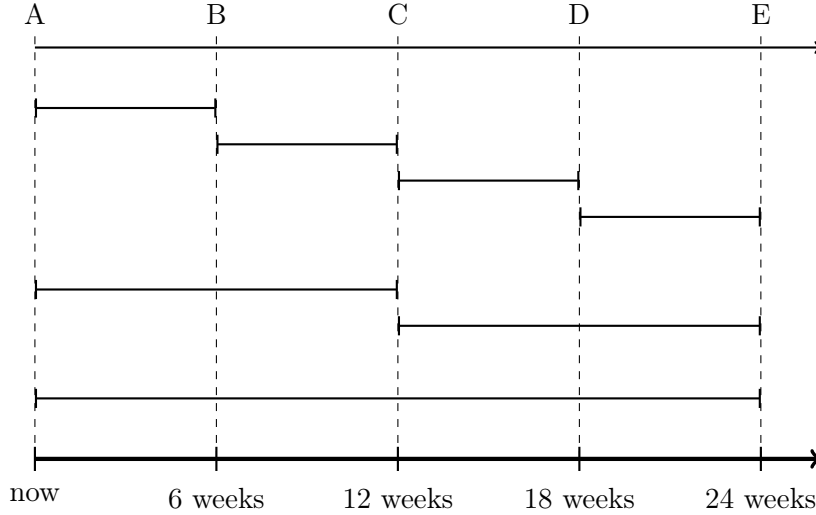


Figure 1: Time horizons used in experiment

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

The future outcome was fixed at €50. The choice of such a round, invariant amount was meant to ensure that outcomes are perceived objectively, rather than being subject to noisy perceptions themselves (see online appendix E for a discussion of robustness to this assumption). The earlier amounts ranged between €33 and €49 inclusive, allowing for discount factors between 0.66 and 1 (discount rates between 0 and 52%) per period. Each screen presented one single choice, and the individual choice pairs were presented completely at random. This design was implemented to fit the discrete binary choice setup predicted by the model. Subjects made a total of 158 choices, which took about 20 minutes on average. All stimuli were presented at least once. In addition, 40 randomly selected stimuli were repeated. The repeated extractions were executed with replacement, so that the same stimulus may recur a number of times. Identification of decision noise, which plays a central role in the model, is thus assured by i) repeat observations

of the same stimuli; ii) monotonicity violations between similar stimuli.

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

3.2 Econometric approach

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

those emerging from the model here presented, and which are rather closer to traditional setups—a point to which I will return in the discussion.

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

$$\mathcal{L}_i(\boldsymbol{\theta}_n|z) \propto p(z|\boldsymbol{\theta}_n) \times p(\boldsymbol{\theta}_n), \quad (14)$$

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

$$p(z|\boldsymbol{\theta}_n) = (Pr[(x, \tau) \succ (y, 0)])^z \times (1 - Pr[(x, \tau) \succ (y, 0)])^{1-z}, \quad (15)$$

with $Pr[(x, \tau) \succ (y, 0)]$ taking the form of the choice probability in equation 7 for the noisy coding model. For the standard models estimated below, on the other hand, the choice probability will be defined by the functional form underlying the model (shown below on a case by case basis), plus a normally distributed additive error term, which is usually assumed to take the form of ‘white noise’ (Hey and Orme, 1994; Bruhin, Fehr-Duda and Epper, 2010).

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

$$p(\boldsymbol{\theta}_n) = \mathcal{N}(\bar{\boldsymbol{\theta}}, \boldsymbol{\Sigma}), \quad (16)$$

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

chosen in a way as to be mildly regularizing, thus helping the algorithm to converge, but being wide enough to accommodate any plausible parameter values that may emerge from the data. This follows best practices in Bayesian econometrics (McElreath, 2016), and the estimates reported below are not sensitive to changes in the hyperpriors used, given that the amount of data can easily overpower any prior at the aggregate level.

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

4 Results


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

the introduction of the up-front delay ($p \ll 0.001$).


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

To test delay-dependence, I identify a probabilistic non-parametric discount factor from the sooner amount at which a subject starts choosing the smaller-sooner reward 50% of the time, divided by the later amount. Figure 4 shows two examples of delay-dependence. Panel 4(a) compares a 12 week delay from the present (AC) to its two underlying 6 week delays (AB and BC). Delay-dependence predicts $\delta_{0,12} > \delta_{0,6} \times \delta_{6,12}$, where δ is the discount factor and subscripts indicate time delays in weeks. The great majority of data points falls above the 45° line, indicating delay-dependence. Importantly, this also holds true when we exclude calculated discount factors $\delta_{0,6} \times \delta_{6,12}$ smaller than 0.66, which could otherwise bias the findings due to censoring effects ($p \ll 0.001$).

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



(a) present-bias, 6 weeks



(b) present-bias, 12 weeks

Figure 2: Non-parametric illustration of present bias

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





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

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


dependence in the remaining comparisons CE versus CD times DE, and AE versus AC times CE are similar, and are shown in online appendix F.

I next estimate the noisy coding model, and test its predictive performance against different standard models augmented by a normally distributed, additive error term, as most commonly used in the literature. The noisy coding model easily outperforms the exponential model (ELPD diff: -380.0 ; SE diff: 24.7)². It also outperforms the quasi-hyperbolic model (ELPD difference: -265.1 ; SE difference: 27.7). This is hardly surprising given the strongly delay-dependent patterns documented above, which cannot be captured by these models. The model derived above also clearly outperforms the model derived from the choice rule entailing strongly decreasing impatience in equation 13 (ELPD difference: -5181.2 ; SE difference: 87.4), which is again unsurprising in light of the fact that strongly decreasing impatience is not empirically supported in the data.

²The ELPD difference constitutes a truly Bayesian generalization of the deviance information criterion. Results are virtually identical if I base the model tests on the Watanabe-Aikake Information criterion (WAIC) instead.



(a) subadditivity, 12 weeks from present



(b) subadditivity, 24 weeks vs 6-week delays


Figure 4: Non-parametric illustration of subadditivity

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


The noisy coding model, furthermore, also outperforms the standard version of the constant-sensitivity model, where $\delta_{s,\ell} = \exp(-\widehat{\alpha}(\ell^{\widehat{\beta}} - s^{\widehat{\beta}}))$, with the ‘hat’ distinguishing the parameters from the equivalent parameters in the noisy coding model. (ELPD diff: -265.0 ; SE diff: 27.8). While the constant-sensitivity model can account for delay-dependence, it predicts the co-existence of delay-dependence and strongly decreasing impatience, which are described by one and the same parameter. We have, however, seen that delay-dependence is strong in the data, while strongly decreasing impatience is not supported. This explains the poor predictive performance of the constant-sensitivity model in the present data.

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

The standard deviation of the encoding distribution, ν , is an essential building block of the noisy coding model. Given this parameter’s interpretation as reflecting noisiness in numerical and time perception, we would expect it to be



(a) time sensitivity in CS vs NCT models



(b) time sensitivity in SA vs NCT model

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

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

correlated with mathematical ability. I obtained two measures of mathematical ability. One is the high school graduation grade in mathematics, measured out of 10, with higher scores indicating better grades. The second consists of the answers to a question measuring self-declared math affinity, with answers ranging from 1 ('no affinity with maths') to 11 ('math wizard') on a Likert scale. Individual-level noisiness in coding, ν_n , indeed shows a negative correlation with self-declared math affinity (Bayesian 95% credibility interval $[-0.273, -0.021]$), as well as with the final high-school grade in mathematics (Bayesian 95% credibility interval $[-0.210, -0.005]$). This further supports the channel of cognitive ability as a driver of coding noise and hence time distortions.

Discussion

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

The noisy coding model shares common foundations with recent models applied

to decision making under risk and uncertainty. In particular, [Khaw et al. \(2020\)](#) showed that a model in which outcomes are subject to noisy mental coding and subsequent Bayesian decoding can account for small stake risk aversion, which may otherwise be hard to reconcile with reasonable levels of risk aversion over larger stakes. [Frydman and Jin \(2021\)](#) generalized their setup to a model in which coding noise can efficiently adapt to the variability of stimuli expected in the environment, and experimentally documented the occurrence of the context effect predicted by their model. In [Vieider \(2021\)](#), I showed that if probabilities are noisily encoded as well as outcomes, then a noisy coding setting will naturally generate prospect-theory like patterns, as well as predicting violations of prospect theory. While still in its infancy, this approach thus appears to hold the promise of a unification of descriptive approaches to individual decision-making.

The model I presented is closely related in spirit to the model of [Gabaix and Laibson \(2017\)](#). The latter propose a model whereby future *utilities* are perceived with noise. If this noise in the perception of utilities is linearly increasing in the time delay attached to the utility, then the combination of these noisy perceptions with a prior results in posterior expectations taking the form of the hyperbolic discounting function proposed by [Mazur \(1987\)](#). Their modelling setup thus formalizes the intuition that hyperbolic discounting patterns may derive from difficulties people have imagining their future selves. Notice, however, that the predictions emerging from their model agree with those emerging from traditional hyperbolic discounting models. For instance, the same mechanism must explain present-bias, strongly decreasing impatience, and delay-dependence. The data I presented, however, indicate the need for different albeit interacting mechanisms to explain these phenomena, such as predicted by the noisy coding model.

The noisy coding model also has some connections to accounts of discounting based on the inherent uncertainty of the future. In particular, [Halevy \(2008\)](#) showed that if agents have some uncertainty about future payoffs, and if they exhibit nonlinear probability distortions, then this will result in present-biased behaviour. [Epper and Fehr-Duda \(2018\)](#) further showed how the uncertainty of

the future could result in generalized hyperbolic patterns, as well as accounting for a number of other stylized facts that have been documented in the literature. The uncertainty in these models concerns the contract survival probability, i.e. the probability that a future payment will really take place and can be enjoyed. The uncertainty in the model I have presented, on the other hand, concerns the mental representation of the time delay itself. My experimental procedures were specifically devised to exclude future uncertainty, which may have plausibly worked due to the regular interactions between the students and the lecturer. In more general settings, however, the two accounts may well complement each other.

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

References

- Abdellaoui, Mohammed, Han Bleichrodt, Olivier L'Haridon, and Corina Paraschiv (2013) 'Is There One Unifying Concept of Utility? An Experimental Comparison of Utility under Risk and Utility over Time.' *Management Science* 59(9), 2153–2169
- Ainslie, George (1975) 'Specious reward: A behavioral theory of impulsiveness and impulse control.' *Psychological Bulletin* 82(4), 463–496
- Baillon, Aurélien, Han Bleichrodt, and Vitalie Spinu (2020) 'Searching for the reference point.' *Management Science* 66(1), 93–112

- Balakrishnan, Uttara, Johannes Haushofer, and Pamela Jakiela (2020) ‘How soon is now? evidence of present bias from convex time budget experiments.’ *Experimental Economics* 23(2), 294–321
- Bruhin, Adrian, Helga Fehr-Duda, and Thomas Epper (2010) ‘Risk and Rationality: Uncovering Heterogeneity in Probability Distortion.’ *Econometrica* 78(4), 1375–1412
- 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
- Conte, Anna, John D. Hey, and Peter G. Moffatt (2011) ‘Mixture models of choice under risk.’ *Journal of Econometrics* 162(1), 79–88
- Cubitt, Robin P, and Daniel Read (2007) ‘Can intertemporal choice experiments elicit time preferences for consumption?’ *Experimental Economics* 10(4), 369–389
- Dayan, Peter, and Laurence F Abbott (2001) *Theoretical neuroscience: computational and mathematical modeling of neural systems* (Computational Neuroscience Series)
- Dehaene, Stanislas (2003) ‘The neural basis of the weber–fechner law: a logarithmic mental number line.’ *Trends in cognitive sciences* 7(4), 145–147
- Dehaene, Stanislas, and Jean-Pierre Changeux (1993) ‘Development of elementary numerical abilities: A neuronal model.’ *Journal of cognitive neuroscience* 5(4), 390–407
- Dohmen, Thomas, Armin Falk, David Huffman, and Uwe Sunde (2017) ‘The robustness and pervasiveness of sub-additivity in intertemporal choice.’ In ‘Working Paper’
- Doya, Kenji, Shin Ishii, Alexandre Pouget, and Rajesh PN Rao (2007) *Bayesian brain: Probabilistic approaches to neural coding* (MIT press)
- Ebert, Jane E. J., and Drazen Prelec (2007) ‘The Fragility of Time: Time-Insensitivity and Valuation of the Near and Far Future.’ *Management Science*

- 53(9), 1423–1438
- Enke, Benjamin, and Thomas Graeber (2021) ‘Cognitive uncertainty in intertemporal choice.’ Technical Report, Mimeo
- Epper, Thomas, and Helga Fehr-Duda (2018) ‘Risk in time: The intertwined nature of risk taking and time discounting.’ *Working Paper*
- Fechner, Gustav Theodor (1860) ‘Elements of psychophysics, 1860.’
- 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
- Frydman, Cary, and Lawrence J Jin (2021) ‘Efficient coding and risky choice.’ *Quarterly Journal of Economics, forthcoming*
- Gabaix, Xavier, and David Laibson (2017) ‘Myopia and discounting.’ Technical Report, National bureau of economic research
- Gelman, Andrew, Jessica Hwang, and Aki Vehtari (2014) ‘Understanding predictive information criteria for bayesian models.’ *Statistics and computing* 24(6), 997–1016
- Gold, Joshua I, and Michael N Shadlen (2001) ‘Neural computations that underlie decisions about sensory stimuli.’ *Trends in cognitive sciences* 5(1), 10–16
- Halevy, Yoram (2008) ‘Strotz Meets Allais: Diminishing Impatience and the Certainty Effect.’ *The American Economic Review* 98(3), 1145–1162
- He, Lisheng, Russell Golman, and Sudeep Bhatia (2019) ‘Variable time preference.’ *Cognitive Psychology* 111, 53–79
- Hey, John D., and Chris Orme (1994) ‘Investigating Generalizations of Expected Utility Theory Using Experimental Data.’ *Econometrica* 62(6), 1291–1326
- 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
- Khaw, Mel Win, Ziang Li, and Michael Woodford (2020) ‘Cognitive imprecision and small-stakes risk aversion.’ *The Review of Economic Studies, forthcoming*
- Knill, David C, and Alexandre Pouget (2004) ‘The bayesian brain: the role of

- uncertainty in neural coding and computation.’ *TRENDS in Neurosciences* 27(12), 712–719
- Laibson, David (1997) ‘Golden Eggs and Hyperbolic Discounting.’ *The Quarterly Journal of Economics* 112(2), 443–478
- Loewenstein, George, and Drazen Prelec (1992) ‘Anomalies in Intertemporal Choice: Evidence and an Interpretation.’ *The Quarterly Journal of Economics* 107(2), 573–597
- Lu, Jay, and Kota Saito (2018) ‘Random intertemporal choice.’ *Journal of Economic Theory* 177, 780–815
- Mazur, James E (1987) ‘An adjusting procedure for studying delayed reinforcement.’ *Commons, ML.; Mazur, JE.; Nevin, JA* pp. 55–73
- McElreath, Richard (2016) *Statistical Rethinking: A Bayesian Course with Examples in R and Stan* (Academic Press)
- Petzschnner, 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
- Read, Daniel (2001) ‘Is Time-Discounting Hyperbolic or Subadditive?’ *Journal of Risk and Uncertainty* 23(1), 5–32
- Rogers, Alan R (1994) ‘Evolution of time preference by natural selection.’ *The American Economic Review* pp. 460–481
- Rubinstein, Ariel (2003) “‘Economics and Psychology’? The Case of Hyperbolic Discounting*.’ *International Economic Review* 44(4), 1207–1216
- Samuelson, Paul A. (1937) ‘A Note on Measurement of Utility.’ *The Review of Economic Studies* 4(2), 155–161
- Scholten, Marc, and Daniel Read (2006) ‘Discounting by intervals: A generalized model of intertemporal choice.’ *Management science* 52(9), 1424–1436
- (2010) ‘The psychology of intertemporal tradeoffs.’ *Psychological Review* 117(3), 925–944
- Stan Development Team (2017) ‘RStan: the R interface to Stan.’ R package version 2.17.2

- Thaler, Richard (1981) ‘Some empirical evidence on dynamic inconsistency.’ *Economics Letters* 8(3), 201–207
- Vehtari, Aki, Andrew Gelman, and Jonah Gabry (2017) ‘Practical bayesian model evaluation using leave-one-out cross-validation and waic.’ *Statistics and computing* 27(5), 1413–1432
- Vieider, Ferdinand M. (2021) ‘Noisy neural coding and decisions under uncertainty.’ Working Papers of Faculty of Economics and Business Administration, Ghent University, Belgium 21/1022, Ghent University, Faculty of Economics and Business Administration
- Vilares, Iris, and Konrad Kording (2011) ‘Bayesian models: the structure of the world, uncertainty, behavior, and the brain.’ *Annals of the New York Academy of Sciences* 1224(1), 22
- Zauberman, Gal, B Kyu Kim, Selin A Malkoc, and James R Bettman (2009) ‘Discounting time and time discounting: Subjective time perception and intertemporal preferences.’ *Journal of Marketing Research* 46(4), 543–556

ONLINE APPENDIX

A Model derivation for delays from the present

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

$$P[\ln(t)|r] \sim \mathcal{N}\left(\frac{\sigma_t^2}{\sigma_t^2 + \nu_t^2} \times r + \frac{\nu_t^2}{\sigma_t^2 + \nu_t^2} \times \mu, \frac{\nu_t^2 \sigma^2}{\nu_t^2 + \sigma^2}\right), \quad (17)$$

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

B Alternative choice rule in log space

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

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

C Derivation with up-front delays

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

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

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

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

D A choice rule for generalized hyperbolicity

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

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

$$P[(x, \ell) \succ (y, s)] = \Phi \left[\frac{\beta^{-1} [\ln [-\ln(\frac{x-y}{x})] - \ln(\alpha)] - [\ln(\ell) - \ln(s)]}{\sqrt{2\nu}} \right], \quad (18)$$

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

E Generalization to include outcome distortions


I have assumed outcomes to be perceived without noise. If outcomes are also subject to noisy coding, the two dimensions will need to be derived jointly to accurately represent the decision-making process. Starting from the linear in log-

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


$$Pr[(x, \ell) \succ (y, 0)] = \Phi \left(\frac{\ln(\theta) + \beta \times \ln(t) - \gamma \times \ln[-\ln(\frac{y}{x})]}{\sqrt{\nu^2 \beta^2 + \nu_o^2 \gamma^2}} \right). \quad (19)$$

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

F Additional results



(a) subadditivity, 12 weeks from 12 weeks



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

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

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

G Experimental Instructions

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

Make your choice:

€49 immediately €50 in 24 weeks

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

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

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