Cognitive Foundations of Delay-Discounting*

Ferdinand M. Vieider†[1]

[1] *RISLαβ, Department of Economics, Ghent University*

September 2, 2025

Abstract

I present a model that predicts well-known empirical deviations from exponential discounting to emerge from the optimal Bayesian combination of noisily encoded time delays with prior information about the statistics of the environment. I further show how the model is formally identical to a setup under risk explaining deviations from expected utility maximization. The model thus provides a unified, neuro-biologically founded perspective of the origin of decision "biases" across domains. Biases result from an optimal process of Bayesian inference, which reins in noise in information processing by adaptation to the decision environment. The model is *generative* in nature, and the parameters describe causal drivers of choice behaviour. The model forms part of a series of papers that show how choice patterns traditionally attributed to preferences may emerge from optimal reactions to cognitive frictions in information-processing.

A great deal can be learned about rational decision making by taking into account [...] the limitations upon the capacities and complexity of the organism, and by taking account of the fact that the environments to which it must adapt possess properties that permit further simplication [sic] of its choice mechanisms.

Herbert Simon (1956), p. 129

1 Motivation

I present a theoretical account of how stylized patterns of delay-discounting observed in empirical work can emerge from the noisy neural coding and processing of the numerical quantities describing the choice situation. Commonly observed discounting anomalies emerge from noise arising in mental computations due constraints inherent in the neural architecture. Deviations from exponential discounting are characterized as optimal reactions to such noise in mental processing. I further show that the model is formally identical to models used to characterize decision-making processes under risk and uncertainty.

The model assumes that the mind is hard-wired to implement discounted (expected) utility. Frictions in information processing, however, mean that the mind does not have direct

*I acknowledge financial support from the Research Foundation—Flanders (FWO) under the project "Causal Determinants of Preferences" (G008021N). I am indebted to Carlos Alós-Ferrer, Miguel Ballester, Ranoua Bouchouicha, Olivier L'Haridon, Frank Heinemann, Yuchi Li, Nick Netzer, Ryan Oprea, Pietro Ortoleva, Rafael Polania, Christian Ruff, Peter P. Wakker, and Leeat Yariv for helpful comments and discussions. All errors remain my own.

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

access to the objective choice primitives presented to the decision-maker. Informational frictions in representing and recombining choice primitives are optimally dealt with by exploiting prior information about the probability distribution of different choice primitives in the environment. The decision process thus takes the form of a Bayesian inference problem, whereby observed choice patterns are causally driven by biases in inferences about the true choice primitives. The inferences are nevertheless optimal, in the sense that they minimize the mean squared error of the predictions over the course of many trials.

Relation to the delay-discounting literature. The model organizes many phenomena that have been documented in empirical and theoretical work in psychology and behavioural economics. A key prediction of the model is that discounting should be *subadditive*—i.e. measured discount rates will systematically decrease in the length of the time period used to measure them—as first documented by Read (2001). Other than proposed by Read (2001), however, *present bias*—a specific preference for the sooner outcome when it occurs immediately (Laibson, 1997; Imai et al., 2021; Cheung et al., 2023)—emerges in the model as a separate motive. Apparent hyperbolicity in delays of increasing length from the present then occurs as a consequence of subadditivity. So does the 'as soon as possible effect' (Kable and Glimcher, 2010), which describes apparently-hyperbolic discounting patterns for delays of increasing length from any given up-front delay.

The model also casts a new light on several findings in the delay-discounting literature that may otherwise be considered puzzling. Take the finding that non-trivial discounting is observed in experiments using monetary outcomes. 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 et al., 2020). Cubitt and Read (2007) showed that for discount rates exceeding the market borrowing rate no inferences can be drawn on time preferences at all if one assumes agents to abide by standard principles of economic rationality. Discounting of money is, however, natural based on the noisy cognition perspective I present here, given that anomalies emerge based on the noisy processing of time delays. At the same time, the noisy processing of rewards predicts differences in discounting patterns between money and consumption if such different rewards are encoded with different fidelity. Finally, cognitive frictions in the processing of rewards as well as time delays predicts an absolute magnitude effect (Thaler, 1981) under empirically plausible conditions.

Causal interpretation of parameters. The parameters of the model I present do not have a descriptive, but a *generative* interpretation. They receive their meaning from their role in the mental processing of choice primitives. In other words, they are supposed to *causally determine* choice behaviour, rather than to merely describe and quantify it. The key quantities of the model can indeed be measured neurally using brain activation signals in fMRI studies (Kable and Glimcher, 2007; Barretto-García et al., 2023; de Hollander et al., 2024). They can thus be used to make causal predictions about how behaviour ought to change if the ease of information processing or the statistics of the environment change. In this sense, the model builds bridges between neuro-biological and neurophysiological predictions and behavioural predictions, thus contributing to the *conscilience* between the social and biological sciences (Wilson, 1997; Glimcher, 2010).

Under some restrictions, the core parameters driving behaviour can also be estimated from binary choice data in a static decision environment. While this is not the ultimate goal or main strength of the model, such an implementation nevertheless overcomes issues affecting the estimation of standard behavioural models. The choice of error models in estimations of standard models is often arbitrary, which can adversely affect the inferences one draws from estimations (Alós-Ferrer et al., 2021). Apesteguia and Ballester (2018) showed forcefully how augmenting decision models under risk and time by an independent error term ('white noise') results in non-monotonicities in the levels of risk aversion and impatience predicted as a function of the utility difference between choice options, thereby resulting in systematic distortions of parameters recovered from data. The model I present here overcomes this issue. Given that the model is inherently stochastic, the error and choice model interact, thus disciplining the behaviour of the choice equation.

Relation to noisy cognition and error literature. The model I present falls into a class of recent models variously defined as 'noisy cognition' or 'efficient coding' models, and which have emphasized how apparent biases in behaviour can emerge from processes that are optimal from a neuro-biological and evolutionary point of view. Robson (2001a), Robson (2001b), and Netzer (2009) present pioneering models endogenizing behavioural predictions based on evolutionary reasoning. Subsequent papers by Steiner and Stewart (2016), Herold and Netzer (2023) , and Netzer et al. (2024) extend this modelling approach to probability distortions, which are characterized as a 'second best' solution given informational constraints.

Formally, the model I propose falls into the class of Bayesian noisy cognition models. Such models have long been used to model sensory-motor tasks in neuroscience (see Ma et al., 2023, for a book-length treatment). Natenzon (2019) showed how correlation structures in noisy signals could be used to rationalize behaviour that might otherwise seem evolutionarily flawed. Khaw et al. (2021) prominently used a Bayesian noisy coding model to explain small-stake risk aversion (see also Khaw et al., 2023).

The paper also relates to a series of recent empirical papers emphasizing the impact of bounded rationality and neural constraints on choice behaviour (Enke and Graeber, 2023; Oprea, 2024). In particular, Enke et al. (2024) show how many of the empirical phenomena known from the delay-discounting literature can be induced in an atemporal environment, which suggests that they may result from the complexity of the choice situation, rather than from anything inherent to time. This account is perfectly consistent with the model I present here.[1]

Finally, the model is related to a string of recent papers modelling delay discounting as a result of errors or perceptual difficulties. For instance, Lu and Saito (2018) and He et al. (2019a) derive predictions on time discounting based on randomness in choices and preferences. Gabaix and Laibson (2022) predict hyperbolic discounting based on the noisy perception of future *utilities*. These models are based on different intuitions from the one I use here, and produce insights that are complementary to the ones I present. They also make very different predictions from the ones emerging from the present model, generally being aimed at explaining strongly decreasing impatience (or 'true hyperbolicity') in discounting. I will return to this point when detailing the behavioural predictions of the model.

2 Model

Below I will model decisions over time, while drawing parallels to decision-making under risk. I will first discuss the key aspects underlying the Bayesian inference model, before turning to behavioural implications.

[1]One important deviation from what I do here consists in the use of choice lists by Enke et al. (2024). This means that the predictions I derive are not directly applicable to their setting, since the model I present is geared specifically at binary choice. On why the differences between binary choices and such choices collected into lists may matter, see Bouchouicha et al. (2024).

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 $y < x$, paid at a sooner time $\tau_s < \tau_\ell$. The model generalizes to more complex tradeoffs—see Online Appendix B—but is ultimately geared at the simple tradeoffs making up the bulk of the empirical literature.[2] I start from a choice rule devoid of subjective parameters[3], under which the larger-later reward is chosen whenever

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

This choice rule is optimal inasmuch as discounting is stationary, resulting in consistent choice patterns over time.[4] 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 and Ebbesen, 1970; Mischel et al., 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 (see Scholten et al., 2014, for a similar modelling assumption). Taking the natural logarithm of both sides for computational convenience (and without loss of generality—see Online Appendix A), we obtain:

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

where $p \triangleq e^{-\tau_\ell}$ and $q \triangleq e^{-\tau_s}$. The left-hand side then takes the form of log-odds, which shows parallels with decisions under risk by conceiving of the exponentials as probabilities. In Vieider (2024a), I present a parallel model for risk, using exactly the same setup as used here applied to uncertain events. Several papers have indeed argued that log-odds

[2]The model I propose is perceptual in nature, and the presentation format will thus be important in determining inferences and hence choices. The focus on simple tradeoffs that constitute the workhorse of experimental invstigations allows for the derivation of simple closed-form solutions, which will allow me to develop the underlying intuitions.

[3]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.

[4]This does not require the decision-maker to consciously 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.

may underly mental representations of stimuli in general (Gold and Shadlen, 2001; 2002; Zhang and Maloney, 2012; Glanzer et al., 2019).

The central idea underlying the model is that choice primitives cannot be accessed directly, but are instead encoded by signals suitable for neural representation and recombination. This follows from the observation that choice primitives will have to be encoded and manipulated by a finite number of neurons using electrical discharges or *spikes* before they can inform a decision. I here assume that any noise arising from this process will apply to the log-odds, and hence to time delays occurring between two payment options, since $ln\left(\frac{p}{q}\right) = ln\left(\frac{e^{-\tau_\ell}}{e^{-\tau_s}}\right) = -(\tau_\ell - \tau_s)$. This constitutes the natural counterpart to noise being attached to the log-odds under risk (cfr. Khaw et al., 2023; Vieider, 2024b). It captures the idea that much of the cognitive bottleneck affecting processing of real-world information stems from higher-order cognitive information processing, rather than affecting purely perceptual quantities (Drugowitsch et al., 2016; Zheng and Meister, 2024).[5]

Similarly to what happens for time delays, signals about the log reward ratio also ought to be represented as being affected by noise. This captures the idea that errors arise from quick comparative tradeoffs between choice options. Further assuming that the choice rule is mentally implemented on (2) multiplied by -1 and logged once again for computational convenience (this is not necessary, but simplifies things: see Online Appendix A, for an alternative derivation), we obtain the following mental choice rule:

$$\mathbb{E}\left[ln\left(ln\left(\frac{x}{y}\right)\right) \,|\, r_r \right] > \mathbb{E}\left[\, ln(\tau_\ell - \tau_s) \,|\, r_t\right] \tag{3}$$

where r_r is the mental representation of the reward ratio, and r_t the mental representation of the log-odds or time delay. I proceed to discussing the mental representation of time delays and reward ratios in sequence.

The mental signal and the likelihood. Given limits on neuronal resources, the mental signal r_t (as well as r_r, to be discussed below) will be affected by noise.[6] Noise in mental

[5]This assumption is made without loss of generality. One could indeed easily add noise in the perception of single time delays from the present τ_s and τ_ℓ on top of what I model here. I forego this possibility to keep the model tractable and parsimonious.

[6]Oprea and Vieider (2024) 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 accuracy can only be obtained from a Beta distribution in the limit as its parameters jointly 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

representations and recombinations 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). For $\tau_s, \tau_\ell > 0$, I thus model the mental signal for the time delay between two rewards as a single draw from the following likelihood function:

$$r_t \sim \mathcal{N}(\, ln(\tau_\ell - \tau_s)\, ,\, \nu_t^2\,), \tag{4}$$

where r_t is a shorthand for $r_t \,|\, \tau_\ell, \tau_s$, emphasizing that the signal is conditional on a specific time delay $\tau_\ell - \tau_s$, and where the parameter ν_t quantifies average coding noise of time delays.[7] Representational noise will thus make observed behaviour stochastic, given that the signal r_t from two subsequent draws may be different even conditional on an identical time delay $\tau_\ell - \tau_s$ being presented to a decision-maker. The accuracy with which time delays are perceived will be proportional to $\lambda_t \triangleq \nu_t^{-2}$ — the *precision* of the likelihood function. The error is attached to the *difference* in delays from the present, as discussed above.

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) 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 representations of time delays do, however, incur 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 rationale of

noise decreasing in the allocation of neural resources at a decreasing rate.

[7] The use of a normal distribution is suggested by physiological and topographic evidence of neural activations (Nieder and Miller, 2003; Dehaene, 2003; Harvey et al., 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 et al., 2014, section 5.3). It has the distinctive advantage of allowing for a transparent closed-form solution of the model. Gold and Shadlen (2001) demonstrate the the qualitative conclusions will generalize to a variety of different distributions, as long as they are single-peaked.

logarithmic coding. To avoid this issue, I add a 1 to the objective time delays as suggested in the literature (Petzschner and Glasauer, 2011; Howard and Shankar, 2018).[8] I will thus substitute $s \triangleq 1+\tau_s$ if $\tau_s > 0$, else $s = 0$, and $\ell \triangleq 1+\tau_\ell$ to model the mental representations of the objective time delays in order to prevent the numerical representation from becoming boundless as the time delay approaches 0. While this transformation is inconsequential for $\tau_s > 0$ since $\ell - s = \tau_\ell - \tau_s$, it will have substantive implications for $\tau_s = 0$.

Inference as optimal signal-decoding. In principle, a direct comparison of the delay signal in (4) with the 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 describing the probability distribution of choice primitives in the environment. This captures the insight contained in the inspirational quote by Herbert Simon at the beginning of this article—limitations to computational capacity can be partly compensated by relying on regularities in the environment.

I assume that the following prior captures the statistics of the delay distributions in the environment:

$$ln(\ell - s) \sim \mathcal{N}(\, ln(\eta)\, ,\, \xi_t^{-1}\,), \tag{5}$$

where $\xi_t = \sigma_t^{-2}$ is the precision of the prior and $ln(\eta)$ its mean. 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 et al. (2001) examined 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.[9] This prior is best thought of as learned over a lifetime of experiences (as well as being subject to local modification over the course of an experiment).[10]

Combining the likelihood in (4) with the prior in (5) by Bayesian updating, we obtain the

[8] The choice of 1 is somewhat arbitrary. The main justification for using 1 instead of a more flexible parameter is to contrain the model and to avoid un-necessary inflation in the number of parameters. I will discuss consequences for potential empirical identification farther below.

[9] In terms of the underlying log-odds representation, the distribution is most accurately characterized as *logit-normal*. Normal distributions are very natural for such log-odds representations (cfr. Gelman et al., 2014, section 5.6), and are popular in statistics for their numerical tractability (Atchison and Shen, 1980).

[10] Technically, this can be modelled by higher-level hierarchies, where decision-makers have prior expectations applying across a variety of situations (see Friston, 2005, for neural and phylosophical foundations for such an explanation). When plunged into a new situation such an experiment, decision-makers will then start updating this hyperprior. Given noise in signals and inferences, such updating will generally be slow, and hence conservative.

following posterior distribution:

$$p\,[\,ln(\ell - s)\,|\,r_t\,] = \mathcal{N}\left(\frac{\lambda_t}{\lambda_t + \xi_t}\,r_t + \frac{\xi_t}{\lambda_t + \xi_t}\,ln(\eta)\,,\,\frac{1}{\lambda_t + \xi_t}\right), \tag{6}$$

where $\frac{1}{\lambda_t+\xi_t} = \frac{\nu^2\sigma^2}{\nu^2+\sigma^2}$ is the variance of the posterior. The posterior mean obtains from a simple linear combination of signal and prior mean, with the signal's contribution determined by the Bayesian evidence weight $\beta \triangleq \frac{\lambda_t}{\lambda_t+\xi_t} = \frac{\sigma^2}{\sigma^2+\nu^2}$. I will thus refer to β as *time-discriminability*, or where no ambiguity arises, simply as *discriminability*.

Proof. Let $d \triangleq \ell - s$, and write the likelihood as $p(r_t|d) \propto exp\left(-\frac{\lambda_t}{2}(\,r_t - ln(d)\,)^2\right)$ and the prior as $p(d) \propto exp\left(-\frac{\xi_t}{2}(\,ln(d) - ln(\eta)\,)^2\right)$. The posterior will be $p(d|r_t) \propto p(d)\,p(r_t|d)$. Multiplying out the exponentials, rearranging, and completing the square we obtain (6). A step-by-step derivation is shown in online appendix C. □

To make the expression in (6) accessible to the experimenter, we will need to replace the mental signal r_t with an observable quantity. This can be done by means of the *response distribution*, which—given stochasticity in $r_t\,|\,t$—will capture the distribution of the inferred time delay, given the *true* time delay. Let $\theta \triangleq \mathbb{E}[\,ln(\ell - s)\,|\,r_t\,]$ be the expected delay inferred from the noisy signal. We then obtain:

$$p\,(\,\theta\,|\,\tau_\ell, \tau_s\,) = \mathcal{N}\left(\beta\,ln(\ell - s) + (1-\beta)\,ln(\eta)\,,\,\frac{\lambda_t}{(\lambda_t + \xi_t)^2}\right), \tag{7}$$

where $\frac{\lambda_t}{(\lambda_t+\xi_t)^2} = \frac{\nu^2\sigma^4}{(\nu^2+\sigma^2)^2}$ is the variance of the response distribution. 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_t conditional 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.

Proof. 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 (7) follows by substituting $a = (1-\beta)\,ln(\eta)$, $z = r_t$, $\widehat{z} = ln(\ell - s)$, $b = \beta$, and $\sigma_z = \nu_t$. □

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 (7), we can write the expectation of the response distribution as follows:

$$\begin{aligned}\mathbb{E}[\,\theta\,|\,\tau_\ell,\tau_s\,] &= ln(\eta)+\beta\,[\,ln(\ell-s)-ln(\eta)\,]\\ &= ln(\ell-s)+(1-\beta)\,[\,ln(\eta)-ln(\ell-s)\,]\,.\end{aligned} \tag{8}$$

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*. Such prediction errors are, in turn, used to amend the prior in proportion to the average signal-to-noise ratio in the mental representations. The second line emphasizes how regression to the mean of the prior will produce systematic deviations in inferences from the true time delay being presented to the decision-maker. To see this, define the bias as the difference between the average inference and the true time delay, $\mathbb{E}[\,\theta\,|\,\tau_\ell,\tau_s\,]-ln(\ell-s)=(1-\beta)[\,ln(\eta)-ln(\ell-s)\,]$. The average inferred time delay will thus be made up of the true time delay plus bias. The 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 $\lim_{\nu_t\to 0}(\beta)=1$. 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 prior.

Notwithstanding this systematic bias, basing the decision on the posterior mean instead of the maximum likelihood estimate r_t 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:

$$\mathbb{E}\left[\left(\widehat{\theta}-ln(d)\right)^2\right]=\left(\frac{\xi_t}{\xi_t+\lambda_t}\,[ln(\eta)-ln(d)]\right)^2+\frac{\lambda_t}{(\lambda_t+\xi_t)^2}, \tag{9}$$

where $d\triangleq\ell-s$, $\widehat{\theta}\triangleq\mathbb{E}[\,\theta\,|\,d\,]$ is the expectation of the response distribution in (7), the first part on the right-hand side is the squared bias, $(1-\beta)^2[\,ln(\eta)-ln(d)\,]^2$, and the second part is the variance of the response distribution from (7). The variance of the response distribution entering the definition of the mean squared error is smaller than the variance of the signal r_t.[11] Most deviations $ln(\eta)-ln(d)$ 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

[11]This obtains simply from $\frac{\lambda_t}{(\lambda_t+\xi_t)^2}<\lambda_t^{-1}$, which will be the case for any $\xi_t>0$, i.e. for any proper prior variance that is strictly smaller than infinity, $\sigma^2<\infty$ (given that for $\sigma^2=\infty$ we revert to the maximum likelihood estimator).

(4) when aggregating across many stimuli (see Ma et al., 2023, section 4.5, for a detailed discussion). It is in this sense that the biased Bayesian estimator is optimal.[12]

2.2 Noisy coding of rewards

Just like delays, rewards are not directly accessible to the mind and thus need to be represented by neural signals. Here, I assume that such signals directly code log reward-ratios. Let $w \triangleq ln\,({}^x\!/_y)$. Assuming that the log of the reward ratio is encoded in a way similar to time delays, we obtain the following likelihood:

$$r_r \sim \mathcal{N}(\,ln(w)\,,\,\nu_r^2\,). \tag{10}$$

The accuracy of the signal for the reward ratio will be given by the precision $\lambda_r \triangleq \nu_r^{-2}$. I assume the following prior distribution for $ln(w)$:

$$ln(w) \sim \mathcal{N}(\,ln(\rho)\,,\,\sigma_r^2\,), \tag{11}$$

where the prior mean, $ln(\rho)$ is the expected log reward ratio. Intuitively, it thus captures captures expectations of how much better future rewards are compared to present rewards.

Defining $\gamma \triangleq \frac{\lambda_r}{\lambda_r+\xi_r}$, we obtain the posterior expectation $\mathbb{E}[\,ln(w)\,|\,r_r\,] = \gamma\, r_r + (1-\gamma)\,ln(\rho)$, and thence the following response distribution:

$$p\,(\,ln(w)\,|\,y, x\,) = \mathcal{N}\left(\gamma\,ln\left(ln\left(\frac{x}{y}\right)\right) + (1-\gamma)\,ln(\rho)\,,\,\gamma^2\,\nu_r^2\right). \tag{12}$$

The discussion concerning optimality of this process follows the same lines as above for time delays, and is thus not repeated here. A formal proof is omitted inasmuch as it is identical to the proof for the time delays included above.[13]

---

[12] 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.

[13] Gabaix and Laibson (2022) 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 $u(x_\tau) \sim \mathcal{N}(\widehat{x}, \zeta^2)$. 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 delay from the present τ, i.e. $\omega_\tau^2 = \omega^2 \times \tau$, they obtain $D(\tau) \triangleq \frac{\zeta^2}{\zeta^2+\omega^2\tau} = \frac{1}{1+(\omega^2/\zeta^2)\tau}$, which takes the form of the proportional discount function proposed by Mazur (1987), $D(\tau) = \frac{1}{1+\chi\tau}$, with

2.3 The stochastic choice rule

We can now use the response distributions in (7) and (12) to arrive at the following choice probability of the larger-later option (see Online Appendix C):

$$Pr[(x, \tau_\ell) \succ (y, \tau_s)] = \Phi\left(\frac{\gamma\, ln\left[ln\left(\frac{x}{y}\right)\right] - \beta\, ln\left[\alpha\left(\ell - s\right)\right]}{\sqrt{\beta^2\,\nu_t^2 + \gamma^2\,\nu_r^2}}\right), \tag{13}$$

where $Pr[(x, \tau_\ell) \succ (y, \tau_s)]$ is the probability of the larger-later reward being chosen, and Φ represents the standard normal cumulative distribution function. The probability of choosing the larger-later option thus increases in the reward ratio x/y and decreases in the time delay $\tau_\ell - \tau_s$ as one would expect. The probability of choosing the larger-later option also decreases in $\alpha \triangleq \eta^{\frac{1-\beta}{\beta}}\,\rho^{\frac{\gamma-1}{\beta}}$, which thus captures *impatience*. Impatience is thus a positive function of the length of the expected time delay, η, and an inverse function of expectation about the later relative to the sooner reward. The function has a fixed point at α^{-1}: delays shorter than α^{-1} will be over-estimated, whereas delays longer than α^{-1} will be *under*-estimated.

Choice is inherently stochastic. In contrast to traditional discounting models, which combine a deterministic preference model with an independently chosen stochastic choice model (He et al., 2019b), the probabilistic setup used here produces an inherently stochastic model of inter-temporal choice. The response noise, $\sqrt{\beta^2\,\nu_t^2 + \gamma^2\,\nu_r^2}$, corresponds to the sum of the standard deviations of the response distributions in (7) and (12). This gives us a new sense of the optimality of the Bayesian combination of signal and prior. Whereas variation of the signals r_t and r_r increase unboundedly in coding noise ν_t and ν_r, response noise $\sqrt{\beta^2\,\nu_t^2 + \gamma^2\,\nu_r^2}$ is non-monotonic in coding noise, since $\beta^2\,\nu_t^2$ and $\gamma^2\,\nu_r^2$ increase up to $\nu_t = \sigma_t$ and $\nu_r = \sigma_r$, and decrease again thereafter. Intuitively, as long as the signal-to-noise ratio in r_t and r_r is favourable, the mind predominantly relies on this signal in its inference process, with stochasticity increasing in the noise of the signal. Once the signal-to-noise ratio becomes unfavourable, however, the mind increasingly relies on the prior. This means that as coding noise increases further, stochasticity in behaviour actually declines. Response noise is thus non-monotonic in coding noise.

This addresses issues arising from scale arbitrariness in random utility models, which arise

$\chi \triangleq \frac{\omega^2}{\zeta^2}$. The model thus predicts discounting to be proportional to the time delay from the present, but cannot organize subadditivity or other deviations from proportional discounting as modelled here.

when both numerator and denominator are defined independently on arbitrary scales as documented by Apesteguia and Ballester (2018). The latter show that the level of risk aversion and impatience predicted by a random utility model is non-monotonic in utility curvature. Utility parameters estimated from such a model will thus be systematically biased—a problem that is difficult to solve for behavioural generalizations of discounted expected utility. The model I present here solves this issue due to the tight interlinkage of the numerator and denominator in (13). The non-monotonicity of the response noise in the denominator thereby restores monotonicity in predicted choices.

2.4 Meaning of parameters

The model I present is *generative* in nature, rather than descriptive like typical models in behavioural economics. This means that the parameters have causal meaning in terms of the mental inference processes they are meant to govern. Their meaning for characterizing behaviour thus arises from the observation that they make causal predictions on how behaviour ought to change with the accessibility of the choice attributes or the probability distributions of choice primitives expected by the decision-maker. The parameters thus have intrinsic neural meaning, and can generally be identified based on neuro-physiological data (see e.g. Barretto-García et al., 2023 and de Hollander et al., 2024 for work identifying coding noise signatures in the brain for decisions under risk).[14]

The expression in (13) is a Probit link function, and can thus in principle be directly estimated by mapping it into binary choices by means of a Bernoulli distribution. Separate identification of the parameters γ and β, however, will require additional restrictions on model parameters.[15] Most comparative statics, however, will be driven by the ratio β/γ and by α. These two quantities are straightforward to identify based on structural estimations or in reduced form models, yielding the signal-to-noise ratio of time delays *relative* to rewards. Separate identification of β and γ require additional assumptions on the noise structure (see Vieider, 2024b, for a discussion for the parallel case under risk). As it turns

[14]Here as throughout, I make no distinction between the 'brain' and the 'mind'. In this sense, perception cannot be disentangled from physical processes happening in the brain, and should instead be seen as their direct consequence. See Damasio (2006) for a book-length treatment of the distinction between brain and mind, and why it is largely meaningless. I am grateful to Peter Wakker for pointing out the usefulness of making this aspect of the model explicit.

[15]Identifying the parameters separately will require an assumption of equal coding noise for rewards and time delays, i.e. $\nu = \nu_t = \nu_r$ (see Vieider, 2024b for a discussion of the equivalent identification restrictions under risk). This comes at the cost of interpretability of some of the original parameters of the model.

out, β/γ and by α allow to fully characterize the behavioural predictions of the model, to which I turn next.

2.5 Behavioural implications

To discuss the behavioural implications of the model, it will be convenient to rewrite (13) in the following, mathematically equivalent way:

$$Pr[(x, \tau_\ell) \succ (y, \tau_s)] = \Phi\left(\frac{ln\left(-ln\left(\frac{y}{x}\right)\right) - \gamma^{-1}\,\beta\, ln\left(\alpha\left(\ell - s\right)\right)}{\sqrt{\nu_r^2 + \gamma^{-2}\beta^2\,\nu_t^2}}\right), \tag{14}$$

where the inverse reward ratio y/x takes the form of a nonparametric discount factor.

A key implication of the model is subadditive discounting (Read, 2001) and the as-soon-as-possible effect (Kable and Glimcher, 2010), i.e. measured impatience systematically depends on the time delays used to measure it. Although often ignored by economic models, such subadditivity is pervasive in empirical data (Dohmen et al., 2017). Equation (14) predicts that the degree of subadditivity will depend on the ratio β/γ. In particular, for $\beta < \gamma$ short delays $\ell - s < \alpha^{-1}$ will be over-estimated, while long delays $\ell - s > \alpha^{-1}$ will increasingly be under-estimated. Subadditivity is thereby a consequence of time delays being encoded more noisily than rewards, but it will turn to *superadditivity* (Scholten and Read, 2006) if the signal-to-noise ratio is more favourable for delays than for rewards, i.e. if $\beta > \gamma$. This is not a just-so ex post explanation: the model makes precise causal predictions on how changing the informational content of a signal affects observed choices.

For 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 $\tau_s = 0$.[16] Apparent hyperbolicity in delays of increasing length from the present such as documented e.g. by Thaler (1981) is thus purely a consequence of subadditivity.[17] To see this, let us start by

[16] That is, the model does not predict strongly decreasing patience as defined by Prelec (2004), i.e. systematic decreases in discounting when a delay of given length is pushed farther into the future from an initial up-front delay.

[17] While hyperbolicity is much-discussed in the economics literature, its empirical status is unclear. Most of the historical discussion of decreasing impatience has used delays of increasing lengths from the present (e.g., Thaler, 1981; Ebert and Prelec, 2007). 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. (2019b) 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 et al., 2010; Cavagnaro et al., 2016). See, however, Bleichrodt et al. (2016) for evidence of strongly decreasing impatience for both health and money. There may also exist systematic differences between findings obtained from binary choices versus choice lists to elicit present or future equivalents.

examining delays from the present, where $\tau_s = 0$ and hence $s = 0$. In this special case, (13) provids cognitive microfoundations for the constant sensitivity discount function of Ebert and Prelec (2007). This is most easily seen by zooming in on the stochastic equality condition in the numerator.[18] Let $\delta_\ell \triangleq \frac{y}{x}$. Exponentiating the numerator twice then yields the following expression at the point of stochastic indifference:

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

This equation shows the similarity with the constant sensitivity function of Ebert and Prelec (2007). For upfront delays $\tau_s > 0$, however, the noisy coding model predicts that *delays between the two options* are nonlinearly distorted (i.e. $exp\left(-[\alpha(\ell - s)]^{\beta/\gamma}\right)$). The constant sensitivity model, on the other hand, predicts transformations of the individual time delays from the present, so that $\delta_{\tau_s,\tau_\ell} = exp(-\widehat{\alpha}^{\widehat{\beta}}(\tau_\ell^{\widehat{\beta}} - \tau_s^{\widehat{\beta}}))$, where $\widehat{\beta}$ is the time-sensitivity parameter. Much like other functions from the hyperbolic family, the constant sensitivity model can thus not account for subadditive discounting (Read, 2001). Let me repeat once more, however, that the main edge of the model derives from the causality of its predictions: for instance, the 'dividing line between the near and far future' $\widehat{\alpha}^{-1}$ (which here is a dividing line between short and long delays, given by α), is predicted by the model to be endogenously determined by the prior expectations of the decision-maker.

Finally, the model differs from the constant sensitivity function in yet another way, since it produces present bias as a separate motive due to the insertion of an additive 1 into the time delay. When $s = 0$, $\ell \triangleq 1 + \tau_\ell$ implies that the function will show a discontinuous drop in the vicinity of $\tau_\ell = 0$. That is, $\lim_{\tau_\ell \to 0} D(\ell) = exp(-\alpha^{\beta/\gamma})$, where $D(\ell)$ indicates the discount function. Assuming as usual that $D(0) = 1$—where no delay is indicated, no delay needs to be encoded—present bias will be given by $1 - exp(-(\alpha)^{\beta/\gamma})$.[19] Present bias thus results purely from the solution to a numerical overflow problem in mental representations. It is, however, predicted to be strictly linked to the key parameters in the model: present bias increases in impatience α, and depends on β/γ. In particular, the smaller time sensitivity β/γ, the larger the present bias will be, since $\beta/\gamma < 1$ will uplift

[18] I define 'stochastic equality' as the larger-later reward being chosen 50% of the time across many repeated choices.

[19] In practical applications, one ought to think of the additive 1 as being applied on a subjective time scale. This could be modelled by dividing the objective time difference by a constant, $\frac{\tau_\ell - \tau_s}{\kappa}$, which would serve to make α and the additive 1 independent of the time units used to express delays (i.e. days vs weeks vs months vs years). The constant κ is thereby best thought of as a collective normalization constant, rather than an individual parameter, which serves to avoid an inflation in the degrees of freedom.

values of $\alpha < 1$ towards 1. Impatience and present bias are thus strongly intertwined in the model, and both are co-determined by time-sensitivity β/γ.

The model can also account for the absolute magnitude effect—the effect whereby impatience decreases as the magnitude of the rewards increases (Thaler, 1981)—but requires an additional feature to do this. In particular, the magnitude effect will obtain if reward coding noise ν_r increases in the (numerical, nominal) stakes, so that γ *decreases*. Such increases in coding noise are plausible if the stakes used fall outside of the range of habitual amounts experienced on a daily basis (Netzer, 2009; Polania et al., 2019). Garagnani and Vieider (2025) test such accounts of numerical habituation in the context of choices under risk, and find that subjects in the UK indeed make more mistakes if identical value is expressed in larger numerical units. Subjects in Japan, who are habituated to larger numerical units due to the currency units they use in everyday life (180 Yen $\approx$ 1 GBP), make more mistakes when identical stakes are expressed using *small* numerical units.

Such increased reward-coding noise is predicted to have two concomitant effects. One, by increasing β/γ it is predicted to result in a decrease in subadditivity. Two, for values of $\alpha < 1$—implying discounting per period of less than 100%—it is predicted to reduce impatience. That, in turn, implies an increase in choices for the larger-later option, as well as a reduction in present bias. Importantly, the effect is predicted to be linked to the numerical magnitudes used to represent the rewards, rather than to any inherent economic value, which makes for an interesting empirical hypothesis.[20]

3 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, and ambiguity, I have shown the model to

[20] Gershman and Bhui (2020) present a model based on Gabaix and Laibson (2022) which explains the absolute magnitude effect as a reduction in the simulation noise of future utilities when the outcome magnitues are large enough to make this worthwhile. Using increases in numerical magnitudes detached from economic value seems a promising way of disentangling these opposite explanations, given that in their model the reduction in noise unambiguously arises from the increase in economic magnitudes in a "rational inattention" optic.

reconcile several findings that remain challenging for any *one* descriptive model to organize. More importantly, the model makes causal predictions on how behaviour should change if one were to manipulate the underlying parameters, either behaviourally or neurally.

The model predictions 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 logarithmic transformations of time delays incorporated in the model present commonalities with behavioural models incorporating logarithmic time *perception*. In two 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 subadditivity thereby emerge based on different patterns of subjective time perception, and are enshrined in two formally different models. Zauberman et al. (2009) and Kim and Zauberman (2009) document how "hyperbolic discounting" can emerge from logarithmic perception of time, and empirically show that average subjective time ratings bear a logarithmic relationship to objective time delays from the present. Cooper et al. (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 papers nevertheless remain behavioural in nature, in the sense that given their perceptual assumptions, the parameters that govern behaviour are added as exogenous variables. To reiterate one last time, the strength of the current model consists in making directional causal predictions on how behaviour ought to change when different choice attributes are made more or less accessible (i.e. by manipulating the cognitive frictions that constitute the 'constraint'), and when decision-makers are made to expect different types of stimuli. A further key element is that the model creates a direct link between neurophysiologically measurable brain activations and observable choice behaviour, with a key role taken up by the dispersion or noise in such measurements. Importantly, this happens without the need to pass via intermediate concepts such as 'subjective perception'.

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., 2019b). 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 subadditivity and present-bias based on different mechanisms, these error models predict that decreasing impatience (hyperbolicity) may arise from errors or 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 stable preferences.

Note also that the absence of a prediction of truly hyperbolic behaviour in the model I presented does not necessarily imply that such behaviour does not exist (even though its existence is empirically contended). For instance, Gabaix and Laibson (2022) present a formally similar approach based on Bayesian statistics, which predicts hyperbolic (proportional) discounting to emerge from difficulties in simulating future utilities. Halevy (2008) and Chakraborty et al. (2020) predict decreasing impatience to emerge from the inherent uncertainty of the future. Such accounts are highly complementary to the predictions emerging from the model presented here, and different mechanisms may well jointly contribute to determining choice behaviour.

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 et al., 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 et al., 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.

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

Alós-Ferrer, Carlos, Ernst Fehr, and Nick Netzer (2021) 'Time will tell: Recovering preferences when choices are noisy.' *Journal of Political Economy* 129(6), 1828–1877

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

Atchison, J, and Sheng M Shen (1980) 'Logistic-normal distributions: Some properties and uses.' *Biometrika* 67(2), 261–272

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

Barretto-García, 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* 7(9), 1551–1567

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

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

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

Bouchouicha, Ranoua, Ryan Oprea, Ferdinand M. Vieider, and Jilong Wu (2024) 'Is prospect theory really a theory of choice?' Technical Report, Ghent University Discussion Papers

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*

Cohen, Jonathan, Keith Marzilli Ericson, David Laibson, and John Myles White (2020)

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

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

Damasio, Antonio R (2006) *Descartes' error* (Random House)

de Hollander, Gilles, Marcus Grueschow, Franciszek Hennel, and Christian C. Ruff (2024) 'Rapid changes in risk preferences originate from bayesian inference on parietal magnitude representations.' *bioRxiv*

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'

Drugowitsch, Jan, Valentin Wyart, Anne-Dominique Devauchelle, and Etienne Koechlin (2016) 'Computational precision of mental inference as critical source of human choice suboptimality.' *Neuron* 92(6), 1398–1411

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 (2023) 'Cognitive uncertainty.' *Quarterly Journal of Economics*

Enke, Benjamin, Thomas Graeber, and Ryan Oprea (2024) 'Complexity and time.' *Journal of the European Economics Association*

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 (2022) 'Myopia and discounting.' Technical Report, National bureau of economic research

Garagnani, Michele, and Ferdinand M. Vieider (2025) 'Economic consequences of numerical adaptation.' *Psychological Science, forthcoming*

Gelman, Andrew, John B Carlin, Hal S Stern, David B Dunson, Aki Vehtari, and Donald B Rubin (2014) *Bayesian data analysis,* vol. 3 (CRC press Boca Raton, FL)

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

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

Glimcher, Paul W (2010) *Foundations of neuroeconomic analysis* (Oxford University Press)

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

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 (2019a) 'Variable time preference.' *Cognitive psychology* 111, 53–79

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

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

Kable, Joseph W, and Paul W Glimcher (2007) 'The neural correlates of subjective value during intertemporal choice.' *Nature neuroscience* 10(12), 1625–1633

_ (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

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

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, 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

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

Netzer, Nick, Arthur Robson, Jakub Steiner, and Pavel Kocourek (2024) 'Risk perception: measurement and aggregation.' *Journal of the European Economic Association* p. jvae053

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 (2024) 'Decisions under risk are decisions under complexity.' *American Economic Review*

Oprea, Ryan, and Ferdinand M. Vieider (2024) 'Minding the gap: On the origins of the description-experience gap.' *Working Paper*

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

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

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

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

_ (2001b) 'Why would nature give individuals utility functions?' *Journal of Political Economy* 109(4), 900–914

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

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

Simon, Herbert A (1956) 'Rational choice and the structure of the environment.' *Psychological Review* 63(2), 129

Steiner, Jakub, and Colin Stewart (2016) 'Perceiving prospects properly.' *American Economic Review* 106(7), 1601–31

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

Vieider, Ferdinand M. (2024a) 'Decisions under uncertainty as bayesian inference on choice options.' *Management Science, forthcoming*

_ (2024b) 'Decisions under uncertainty as bayesian inference on choice options.' *Management Science* pp. 9014–9030

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

Wilson, Edward O (1997) *Consilience* (Knopf)

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

Zheng, Jieyu, and Markus Meister (2024) 'The unbearable slowness of being: Why do we live at 10 bits/s?' *Neuron*

ONLINE APPENDIX

A Choice rule on orignal scale

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 $d \triangleq \ell - s$ on the logarithmic scale just like above. To obtain the posterior expectation of the time delay d 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 (subscripts dropped for notational conveninence). We thus obtain

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

Substituting this mental quantity for t into a choice rule $exp(-\mathbb{E}[d \,|\, r]) > \frac{y}{x}$ 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 (7). 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 $\tilde{\alpha} \triangleq exp\left(\frac{(1-\beta)}{\beta}\widehat{\mu}\right)$.

B 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 some more complex tradeoffs in a straightforward way. Note, however, that generalizing the model further would require a departure from the closed form solutions used in this present paper. Such a generalization may also be questionable, since highly complex tradeoffs with many outcomes and time delays are likely to trigger different cognitive processes from those modelled here.

Take for instance choice options having payouts at different time horizons, such as the ones used by Abdellaoui et al. (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.

C Derivations and proofs

Combination of signal and prior in (6). The likelihood and prior take the following form:

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

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

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

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

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

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:

$$\begin{aligned} -\frac{[ln(t)]^2}{2\sigma_p^2} &= -\frac{1}{2}\left(\frac{1}{\nu^2} + \frac{1}{\sigma^2}\right)[ln(t)]^2 \\ \sigma_p^2 &= (\lambda_t + \xi_t)^{-1}, \end{aligned}$$

which is the variance of the posterior in (6), and where $\lambda_t \triangleq \nu^{-2}$ and $\xi_t \triangleq \sigma^{-2}$.

We can now match the second element:

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

which is the posterior mean from (6). QED.

Joint distribution of time and reward signals. The proof for the combination of the two response functions in 13 is straightforward. The derivations in the main text distribute the signals r_t and r_r independently. The joint distribution simply exploits the assumption of independence in the signals, which I maintain throughout (see Natenzon, 2019, on modelling correlated signals). The difference in signal will then again follow a normal distribution with as its mean the difference of the two means of the signals, and its variance given by the sum of the two variances of the signals.