Rationally inattentive intertemporal choice

Samuel J. Gershman[1,2] and Rahul Bhui[1,3]
[1]Department of Psychology, Harvard University
[2]Center for Brain Science, Harvard University
[3]Department of Economics, Harvard University

June 23, 2019

Abstract

Discounting of future rewards is traditionally interpreted as evidence for an intrinsic preference in favor of sooner rewards. However, temporal discounting can also arise from internal uncertainty in value representations of future events, if one assumes that noisy mental simulations of the future are rationally combined with prior beliefs. Here, we further develop this idea by considering how simulation noise may be adaptively modulated by task demands, based on principles of rational inattention. We show how the optimal allocation of mental effort can give rise to well-known phenomena in intertemporal choice, such as the magnitude effect and the interaction between sign and magnitude effects. In a re-analysis of a large data set and in a new experiment, we reveal several behavioral signatures of this novel theoretical account, tying choice stochasticity to the magnitude effect. We conclude that some aspects of temporal discounting may result from a cognitively plausible adaptive response to the costs of information processing.

1 Introduction

The preference for sooner over later rewards is traditionally interpreted as an intrinsic decline in value as outcomes recede into the future. However, recent evidence suggests an alternative (though not mutually exclusive) viewpoint: temporal discounting could arise from *internal uncertainty* in value representations of future rewards. Imagining the future allows an agent to immediately experience anticipated outcomes, helping them to delay gratification, but this prospection may lose its impact when mental simulations are noisy. A number of studies have demonstrated that this mechanism plays an important role in patience, and have shed light on its neural foundations. Discounting is attenuated when people are asked to imagine spending future rewards (Benoit, Gilbert, & Burgess, 2011), when they imagine future outcomes in greater detail (Lebreton et al., 2013), and when episodic tags are provided to facilitate such imagination (Peters & Büchel, 2010).

Thus, agents with imperfect foresight might intrinsically value future rewards just as much as immediate rewards, but the myopic nature of their prospection reduces the weight of the former. This idea has recently been formalized in a Bayesian model of discounting (Gabaix & Laibson,

2017), in which an agent observes noisy simulations of future value and applies Bayes' rule to obtain a posterior estimate. Assuming simulations become noisier the further they reach into the future, the agent increasingly relies on their prior beliefs and discounts the reward value. Gabaix and Laibson (2017) showed how this can lead to hyperbolic discounting while accommodating the effects of experience on intertemporal choice tasks. However, this analysis is predicated on a fixed relationship between reward delay and simulation noise, though there is reason to think that the relationship is not fixed. We extend this perspective by considering how the degree of simulation noise may be adaptively controlled.

Vivid prospection can help agents to delay gratification, but this comes at a cost. Making simulations more precise requires mental effort, and this effort may only be invoked if its benefits outweigh its costs. We formalize this intuition in terms of *rate-distortion theory* (Berger 1971; see also Caplin 2016; Matějka and McKay 2015; C. A. Sims 2003): precise simulations are cognitively costly, and therefore a decision maker must make a trade-off between precision and effort. Most relevant in the present context, larger rewards may be more important to evaluate precisely. In this case, greater magnitudes would be simulated more precisely and, in light of the above argument, would engender more patience.

We posit that such a mechanism contributes to the well-known *magnitude effect* in intertemporal choice, according to which people are disproportionately more patient when judging high-value outcomes (Green, Myerson, & McFadden, 1997; Johnson & Bickel, 2002; Kirby, 1997; Raineri & Rachlin, 1993; Thaler, 1981). This proposal is consistent with several lines of psychological and neural evidence; for instance, Bulganin and Wittmann (2015) found that more rewarding future events are imagined with greater vividness, and produce more activation in the brain's reward system (in the striatum and substantia nigra/ventral tegmental area). Ballard et al. (2017) found that prefrontal cortex (which has been implicated in the exertion of mental effort; e.g., McGuire & Botvinick, 2010) is more heavily recruited when evaluating high-magnitude intertemporal choices. In what follows, we show how our theory qualitatively accounts for past empirical findings about the magnitude effect, quantitatively improves model fit in a large existing data set, and accurately predicts patterns of discounting and stochasticity in a new experiment. These results help sharpen our understanding of the relationship between patience, reward, and mental effort.

2 Theory

In this section, we first describe the Bayesian model of discounting developed by Gabaix and Laibson (2017). We then extend this analysis by endogenizing the simulation noise variance using a rational inattention analysis.

2.1 A Bayesian model of as-if temporal discounting

Following Gabaix and Laibson (2017), we model an agent who is faced with a choice between several rewards that occur at some time in the future. For ease of exposition, we will consider a single reward r delivered after delay t, whose "true value" is denoted by u. We assume that this value is drawn from a Gaussian distribution with mean μ and variance σ_u^2: $u \sim \mathcal{N}(\mu, \sigma_u^2)$. We further

assume that the agent does not directly observe u, but instead observes a noisy signal $s \sim \mathcal{N}(u, \sigma_\varepsilon^2 t)$ generated by some form of mental simulation.

Noise arises from the agent's limited ability to simulate the event's future value. Gabaix and Laibson (2017) assumed that the variance increases linearly with the delay because events farther in the future are harder to simulate. Combined with the assumption that the prior mean μ is 0 (which we suppose for the remainder of the paper), this leads to the following expression for the posterior mean:

$$\hat{u} = \mathbb{E}[u|s] = D_t s \tag{1}$$

where D_t expresses an "as-if" hyperbolic discount function:

$$D_t = \frac{1}{1 + kt}, \tag{2}$$

with the as-if discount rate k given by:

$$k = \frac{\sigma_\varepsilon^2}{\sigma_u^2}. \tag{3}$$

The discount function is "as-if" because the agent in fact has a neutral time preference, but chooses in accordance with hyperbolic discounting, one of the most broadly supported models of intertemporal choice (see Frederick, Loewenstein, & O'Donoghue, 2002, for a review). Figure 1 illustrates how Bayesian inference in this model produces temporal discounting.

The estimated value of a reward will be regularized towards the mean μ (0 in this case). The strength of this regularization depends on k, which can be thought of as an inverse signal-to-noise ratio. Intuitively, when the simulation noise variance σ_ε^2 is large relative to the prior variance σ_u^2, the simulations are less reliable and the agent will rely more on her prior, whereas when the simulation noise variance is relatively small then agent will rely more on her simulations.

Because we (the experimenters) cannot directly observe the signal s, we use the objective reward r as a proxy. This allows us to link the model directly to experimentally observable variables. We note, however, that this assumption may generate erroneous inferences. For example, we may misinterpret the effects of model misspecification in terms of simulation noise.

2.2 Rational inattention

The Gabaix and Laibson (2017) analysis assumed that the agent has a fixed simulation noise variance. Here we develop the idea that the simulation noise variance is determined by the agent's "attention" to the signal. Intuitively, an agent can improve the reliability of her mental simulations by exerting cognitive effort (i.e., attending more), but pays a cost for this effort.

We approach this problem through the lens of rate-distortion theory (Berger, 1971).[1] The agent is modeled as a communication channel that takes as input the signal and outputs an estimate of the

[1] See C. R. Sims (2016) for an overview of rate-distortion theory applications to human perception, and C. A. Sims (2003), Matějka and McKay (2015), and Caplin (2016) for economic applications of the closely related rational inattention framework.

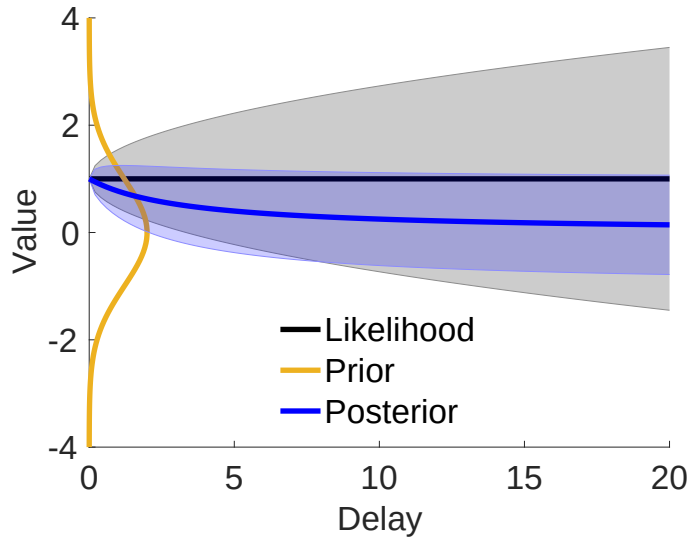

Figure 1: Illustration of discounting as Bayesian inference. Mental simulation generates a noisy value signal ($s_t = 1$), which is combined with a prior distribution to form a posterior distribution over value. Because signal variance (simulation noise) increases with delay, the posterior mean will be regularized more strongly towards the prior mean for longer delays. The curves show means, and the shaded regions show the standard deviation.

value. The agent can select the design of the channel subject to a constraint on the information rate of the channel (the number of bits that can be communicated per signal). Specifically, we define a family of channels parametrized by the simulation noise scaling parameter, σ_ε^2. The optimization problem is to select the value of σ_ε^2 that minimizes the expectation of a distortion function $\mathcal{L}(u, \hat{u})$ that quantifies the cost of estimation error:

$$\bar{\sigma}_\varepsilon^2 = \underset{\sigma_\varepsilon^2}{\operatorname{argmin}}\, \mathbb{E}[\mathcal{L}(u, \hat{u}) | \sigma_\varepsilon^2]. \tag{4}$$

As written out in the Appendix, under a squared error distortion function, $\mathcal{L}(u, \hat{u}) = (u - \hat{u})^2$, the optimal parameter is given by:

$$\bar{\sigma}_\varepsilon^2 = \frac{\sigma_u^2}{(e^{2R} - 1)t}, \tag{5}$$

where R is the information rate constraint in nats. We impose an additional constraint on this formulation, by assuming that information rate increases with reward magnitude (greater incentive to expend cognitive resources) and decreases with delay (simulation of distal events is more cognitively demanding):

$$R = \frac{1}{2} \ln \left(\frac{\beta |r|}{t} + 1 \right), \tag{6}$$

where $\beta > 0$ is a "sensitivity" parameter that governs the relationship between rate, magnitude and delay. As β increases, the rate becomes increasingly sensitive to variations in reward and delay.

The optimal simulation noise parameter is then given by:

$$\bar{\sigma}_\varepsilon^2 = \frac{\sigma_u^2}{\beta|r|}, \tag{7}$$

which implies that the optimal discount parameter is:

$$\bar{k} = \frac{1}{\beta|r|}. \tag{8}$$

Thus, the rate-distortion framework, under the assumptions described above, leads us to a model that predicts the magnitude effect (inverse relation between discount factor and reward magnitude).

We can draw out further implications of this model by connecting it to choice behavior. Let us assume, in the simplest case, that the agent deterministically chooses the option with highest estimated value. In this case, all stochasticity in choice behavior is driven by stochasticity in the agent's simulation process. Marginalizing over these noisy simulations, the choice probability for a standard two-alternative choice (early vs. late) is given by:

$$P(\text{choose early}) = \Phi\left(\alpha\left[D_t r_t - D_{t+\tau} r_{t+\tau}\right]\right), \tag{9}$$

where ϕ is the standard Gaussian cumulative density function, τ is the difference in delay between early (r_t) and late ($r_{t+\tau}$) options, and

$$\alpha = \frac{1}{\sigma_\varepsilon \sqrt{D_t^2 t + D_{t+\tau}^2 (t+\tau)}} \tag{10}$$

is an "inverse temperature" parameter controlling the degree of choice stochasticity (smaller values of α produce greater stochasticity). In the case where the early option is immediate (i.e., $t = 0$), as in many studies of discounting, this simplifies to:

$$\alpha = \frac{1}{\sigma_\varepsilon D_\tau \sqrt{\tau}}. \tag{11}$$

Plugging in the optimal simulation noise parameter gives:

$$\bar{\alpha} = \frac{\sqrt{\beta|r|}(1 + \tau/\beta|r|)}{\sigma_u \sqrt{\tau}}. \tag{12}$$

One can show that

$$\frac{\partial \bar{\alpha}}{\partial |r|} \geq 0 \text{ for } \beta|r| > \tau, \tag{13}$$

which means that for sufficiently large rewards and sufficiently short delays, the model predicts a choice stochasticity magnitude effect: as reward magnitude gets larger, choice stochasticity should get smaller. One can also show that

$$\frac{\partial^2 \bar{\alpha}}{\partial |r| \partial \sigma_u} \leq 0 \text{ for } \beta|r| > \tau, \tag{14}$$

which means that the choice stochasticity magnitude effect declines with reward variance (under the same conditions on reward and delay).

Finally, we can examine what happens to the two magnitude effects when the sensitivity parameter β changes:

$$\frac{\partial^2 \bar{k}}{\partial |r| \partial \beta} \geq 0, \tag{15}$$

$$\frac{\partial^2 \bar{\alpha}}{\partial |r| \partial \beta} \geq 0 \text{ for } \beta |r| > \tau. \tag{16}$$

Because $\frac{\partial \bar{k}}{\partial |r|} \leq 0$, this means that increasing β will *decrease* the discounting magnitude effect (i.e., push it closer to 0). This is somewhat counterintuitive, since one might reason that greater sensitivity to reward should translate into a stronger magnitude effect. This intuition is correct for the choice stochasticity magnitude effect: increasing β will magnify the dependence of choice stochasticity on reward magnitude. The key implication of this analysis is that a change in sensitivity will push the two magnitude effects in opposite directions.

3 Applications to prior experimental results

In this section, we explore the empirical implications of the rational inattention analysis. We begin by examining a pattern of results that serves as a challenging benchmark for models of intertemporal choice: the sign-by-magnitude interaction. Hardisty, Appelt, and Weber (2013) showed that the magnitude effect reverses for losses, with steeper discounting for small losses than for large losses. This finding is problematic for models that assume that losses are simply negative rewards, in which case we would expect the magnitude effect to be monotonic across gains and losses. The rational inattention theory accounts for the sign-by-magnitude interaction (Figure 2) because the optimal agent will reduce simulation noise for large rewards *and* large losses.

To assess the empirical evidence for discounting and choice stochasticity magnitude effects in greater detail, we re-analyzed a large data set ($N = 1284$) of intertemporal choices (Chávez, Villalobos, Baroja, & Bouzas, 2017). Each subject in this study was presented with the same set of 27 choices, taken from Kirby (1997). The rewards for both options and the delay for the larger-later option varied across trials, while the delay for the smaller-sooner option was held fixed at 0 days.

To evaluate the qualitative predictions of the rational inattention theory, we first fit and compared three models with varying degrees of flexibility.

- **M1**: hyperbolic discounting with baseline- and magnitude-dependent discount factor, using the parametrization of Vincent (2016):

$$k_t = \exp(c_k - m_k \log |r_t|), \tag{17}$$

  where c_k is a free parameter capturing baseline discounting, and m_k captures magnitude-

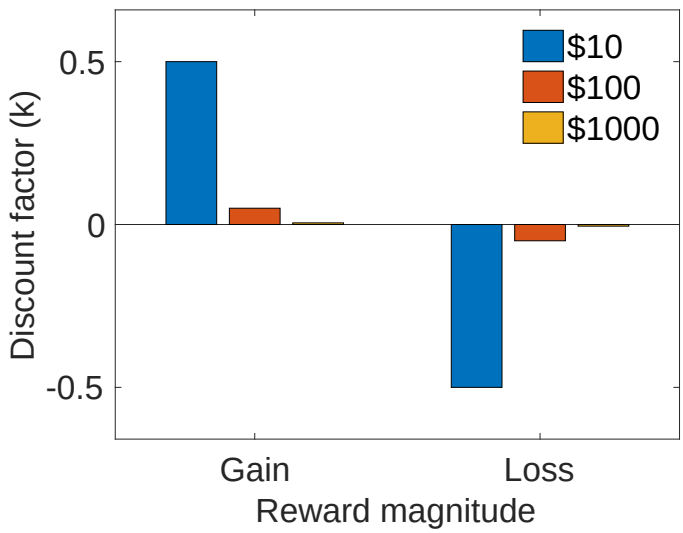

Figure 2: Theoretical discount factor as a function of magnitude and sign, with $\beta = 0.2$ and $\sigma_u^2 = 10$.

dependent discounting.[2] We modeled choices using a generalized version of Eq. 18:

$$P(\text{choose early}) = (1-\omega)\Phi\left(\alpha\left[D_t r_t - D_{t+\tau} r_{t+\tau}\right]\right) + \frac{\omega}{2}, \tag{18}$$

treating α and ω as free parameters. When $\omega > 0$, the model captures "lapses" in which subjects respond at chance. This model has 4 free parameters: c_k, m_k, α, ω.

- **M2**: same as M1, but with baseline- and magnitude-dependent inverse temperature:

$$\alpha = \exp(c_\alpha + m_\alpha \log(|r_t| + |r_{t+\tau}|)). \tag{19}$$

  This model has 5 free parameters: $c_k, m_k, c_\alpha, m_\alpha, \omega$.

- **M3**: same as M2, but without the baseline discounting and inverse temperature parameters. In this case, the parametrization simplifies to $k = |r|^{m_k}$ and $\alpha = (|r_1| + |r_2|)^{m_\alpha}$. This model has 3 free parameters: m_k, m_α, ω.

Note that none of the models defined above are constrained to make the same qualitative predictions as the rational inattention theory. For example, the magnitude scaling parameters might be 0 on average, or might go in a direction opposite what the theory predicts. Thus, fitting these models gives us the opportunity to test whether the parameter estimates are in qualitative alignment with the rational inattention theory, without making a commitment to the specific parametrization of that theory.

To evaluate the *quantitative* predictions of the rational inattention theory, we defined two models using equations 8 and 10 to parametrize the model.

[2]Note that we flipped the sign of m_k relative to what Vincent (2016) used, so that positive values are consistent with the discounting magnitude effect.

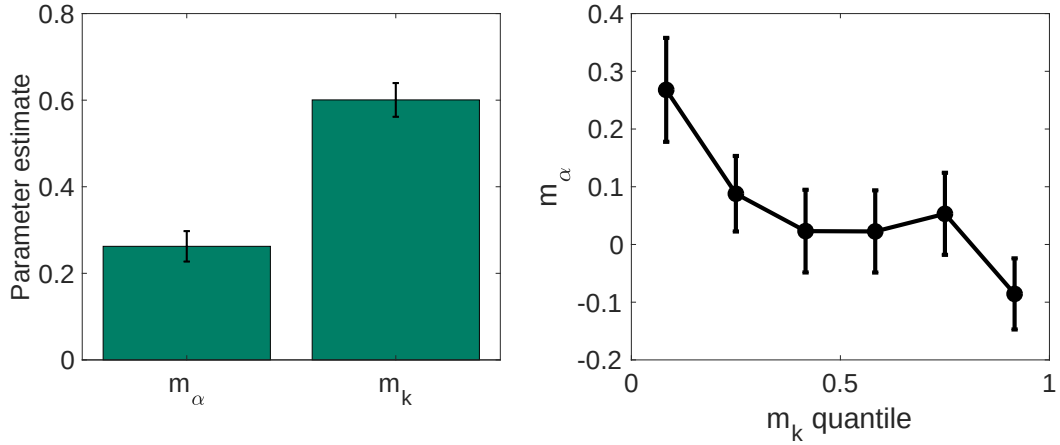

Figure 3: (Left) Magnitude scaling parameter estimates for the 3-parameter hyperbolic discounting model (M3). Error bars show standard error of the mean. (Right) The magnitude scaling parameter for choice stochasticity declines as a function of the magnitude scaling parameter for discounting.

- **M4**: hyperbolic discounting with endogenized discount factor, using Eq. 8, fitting β as a free parameter. We approximated σ_u^2 as the empirical variance of the rewards each individual subject observed in the experiment. As in the other models, we use Eq. 18 to model choices, treating α and ω as free parameters. This model has 3 free parameters: β, α, ω.
- **M5**: hyperbolic discounting with endogenized discount factor and inverse temperature. This model uses the same formulation as M4, but sets α using Eq. 10. The model has 2 free parameters: β, ω.

All models were fit using maximum likelihood estimation. To compare models, we computed the *protected exceedance probability* (PXP; Rigoux, Stephan, Friston, & Daunizeau, 2014), the probability that each model has higher model evidence than all the other models, taking into account the probability that the data may have arisen from a null (chance) model. To approximate model evidence, we used the Bayesian information criterion.

Restricting the comparison to the flexibly parametrized models 1-3, we found that M3 was decisively favored (PXP > 0.99). Thus, as in the rational inattention theory, all variation in choice stochasticity and discounting is due to variations in reward magnitude. Consistent with the rational inattention theory's qualitative predictions, the magnitude scaling parameter for inverse temperature (m_α) was significantly greater than 0 [$t(1283) = 7.47, p < 0.0001$], indicating that choice stochasticity decreases with reward magnitude, whereas the magnitude scaling parameter for discounting (m_k) was significantly less than 0 [$t(1283) = 15.42, p < 0.0001$], indicating that myopia decreases with reward magnitude (Figure 3, left). Finally, we observed that the two magnitude scaling effects are negatively correlated ($r = -0.21, p < 0.0001$; Figure 3, right), consistent with the rational inattention model's predictions.

We next compared the two variations of the rational inattention theory (M4 and M5), finding decisive support for the fully endogenized version (M5; PXP > 0.99). Finally, we compared the

full set of models, finding that M5 received the strongest support (PXP > 0.99). Thus, the data support the theory both qualitatively and quantitatively.

4 The effect of reward variance on discounting and choice stochasticity

The rational inattention model predicts that the choice stochasticity magnitude effect should decrease with reward variance, because the noisy simulations become increasingly down-weighted as the reward variance increases, and this down-weighting interacts multiplicatively with the reward magnitude. The model also predicts that there should be no effect of reward variance on the discounting magnitude effect. We tested these predictions in a new experiment in which the reward variance was manipulated while holding the mean and range of rewards fixed.

4.1 Methods

Two hundred and twenty-one people were recruited from Amazon Mechanical Turk via TurkPrime (Litman, Robinson, & Abberbock, 2017), and paid $1.25 for their participation. To elicit time preferences, we used a choice titration task in which participants made a series of binary choices between a smaller-sooner reward and a larger-later reward. They faced 40 "titrator" trials, each consisting of 6 hypothetical binary choices between fixed smaller-sooner and larger-later options, distinguished by larger-later delays which varied from 1 to 6 months. The smaller-sooner reward was always $1 in every trial. The larger-later rewards were drawn from a Gaussian distribution with mean $5 truncated to be above $1 and below $9, rounded to the nearest cent. Participants were randomly assigned to one of two conditions: in the low variance condition, the larger-later distribution had (untruncated) standard deviation 1, while in the high variance condition, the larger-later distribution had (untruncated) standard deviation 5. Empirically, the former condition had variance 1.03 and the latter had variance 4.97.

4.2 Results

To evaluate the variance predictions, we fit M3 to the choice data. The key parameter estimates are shown in Figure 4, broken down by variance condition. Replicating our prior results with the Chavez data set, we found a significant discounting magnitude effect $[m_k < 0 : t(220) = 4.10, p < 0.0001]$ and a significant choice stochasticity magnitude effect $[m_\alpha > 0 : t(220) = 19.28, p < 0.0001]$ when collapsing across conditions. Critically, the choice stochasticity magnitude effect was significantly lower in the high variance condition, $[t(219) = 3.82, p < 0.0005]$, whereas there was no effect of variance on the discounting magnitude effect ($p = 0.28$). Using a Bayesian t-test with a scaled JZS prior (Rouder, Speckman, Sun, Morey, & Iverson, 2009), we found a posterior probability > 0.99 favoring the null hypothesis that variance does not modulate the discounting magnitude effect. These results collectively provide evidence consistent with our rational inattention model.

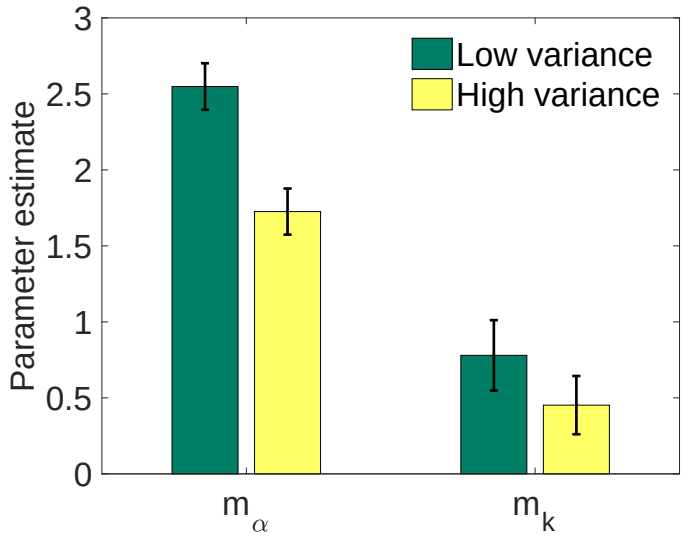

Figure 4: Magnitude scaling parameter estimates for the 3-parameter hyperbolic discounting model (M3), aggregated separately for low and high variance conditions. Error bars show standard error of the mean.

5 Discussion

Temporal discounting may stem partly from the inability of decision makers to perfectly simulate future outcomes. In this paper, we develop a theoretical account of prominent regularities in intertemporal choice, based on the idea that mental simulation of the future is noisy but controllable. Our approach connects the Bayesian model of discounting by Gabaix and Laibson (2017) with the information-theoretic framework of rate-distortion theory (Berger 1971; see also Caplin 2016; Matějka and McKay 2015; C. A. Sims 2003). Supposing the prospective value of a reward becomes noisy when it is internally projected into the future, Bayesian agents should compensate for this uncertainty by relying more heavily on their prior beliefs—if priors are centered near zero, this leads to discounting of value. However, the degree of noise in the simulation may be controlled by the agent, at a cost. If it is more important to accurately evaluate larger rewards, the agent should spend extra mental effort to make their simulations more precise when dealing with greater magnitudes. This mechanism could lead to reduced temporal discounting when dealing with large rewards, a commonly observed phenomenon known as the magnitude effect. Our model can also account for the reversal of this effect in the loss domain (i.e., its interaction with the sign effect), and makes new predictions about how it is related to stochasticity in choice and variability in the reward context, which we validate in the re-analysis of a large data set and a new experiment.

Note that the uncertainty we are dealing with is *internal*. This contrasts with theories of discounting based on objective risk in the arrival of rewards (e.g., Dasgupta & Maskin, 2005; Sozou, 1998). In the present framework, discounting can occur even when the decision maker has no innate preference for earlier rewards and there is no uncertainty in the environment or in the provision of the reward. Of course, all of these pathways are not mutually exclusive, and we do not claim the others are

inconsequential. Our goal is rather to clearly describe how apparent anomalies of intertemporal choice could arise from a cognitively plausible adaptive response to limits on information processing.

Our proposal is supported by a range of neural and behavioral evidence. Psychologically speaking, the allocation of attention in our framework may manifest as *cognitive control*—the set of mechanisms required to pursue a goal in the face of distractions and competing responses. Shenhav, Botvinick, and Cohen (2013) argued that the exertion of control depends on its expected value, the combination of its effort costs and payoff benefits in a given task, and that this plays a role in many decisions including intertemporal choices. Our model formally draws out the implications of this cost-benefit logic, providing a high-level normative perspective that complements more mechanistic approaches to studying attention in this context (e.g., Amasino, Sullivan, Kranton, & Huettel, 2019).

From a neuroscientific perspective, the exertion of cognitive control is known to rely on a network of regions in prefrontal cortex, which some studies have linked directly to temporal discounting (e.g., Benoit et al., 2011; Figner et al., 2010). Shenhav et al. (2013) proposed that the expected value of control is computed by such regions and guides the investment of attention into each task, and Ballard et al. (2017) demonstrated that frontal executive-control areas of the brain are particularly engaged in intertemporal choices with high-magnitude rewards. Moreover, disruption of activity in these areas via transcranial magnetic stimulation reduces the magnitude effect (Ballard, Aydogan, Kim, & McClure, 2018). This mechanism may also be reflected in the reward system, and Bulganin and Wittmann (2015) showed that more rewarding events are indeed imagined with greater vividness and produce more activation in the striatum and substantia nigra/ventral tegmental area. Taken together, these studies indicate that the brain adaptively modulates simulation noise and this plays a meaningful role in temporal discounting.

Another perspective from neuroscience is provided by studies of patients with Parkinson's disease, who are known to have systemically low levels of dopamine. Foerde et al. (2016) observed that patients on medication (with putatively higher dopamine levels) exhibit both more patience (higher estimated values of the k parameter) and a weaker magnitude effect compared to patients off medication. Both of these findings are consistent with the idea that higher levels of dopamine correspond to higher values of the sensitivity parameter β. Higher sensitivity means that reward will induce a greater willingness to exert cognitive effort, which in this case means reducing simulation noise and thereby reducing discounting. At the same time, increases in sensitivity will actually make the magnitude effect smaller, because of the concave relationship between the discount parameter and reward magnitude. Our interpretation of dopamine in terms of sensitivity is consistent with other work on Parkinson's patients showing that high levels of dopamine produce greater reward sensitivity (Chong et al., 2015; Manohar et al., 2015). More broadly, it has been suggested that dopamine may control allocation of cognitive effort (Westbrook & Braver, 2016). We conjecture that dopamine may play a specific role in mediating the relationship between reward and information rate, but further research will be required to directly test this hypothesis.

Data Availability

The data and code for all statistical analyses are available online at:
https://github.com/sjgershm/rational-discounting

References

Amasino, D. R., Sullivan, N. J., Kranton, R. E., & Huettel, S. A. (2019). Amount and time exert independent influences on intertemporal choice. *Nature Human Behaviour*, *3*, 383-392.

Ballard, I. C., Aydogan, G., Kim, B., & McClure, S. M. (2018). Causal evidence for the dependence of the magnitude effect on dorsolateral prefrontal cortex. *Scientific Reports*, *8*.

Ballard, I. C., Kim, B., Liatsis, A., Aydogan, G., Cohen, J. D., & McClure, S. M. (2017). More is meaningful: The magnitude effect in intertemporal choice depends on self-control. *Psychological Science*, *28*(10), 1443–1454.

Benoit, R. G., Gilbert, S. J., & Burgess, P. W. (2011). A neural mechanism mediating the impact of episodic prospection on farsighted decisions. *Journal of Neuroscience*, *31*(18), 6771–6779.

Berger, T. (1971). *Rate distortion theory: A mathematical basis for data compression.* Prentice-Hall.

Bulganin, L., & Wittmann, B. C. (2015). Reward and novelty enhance imagination of future events in a motivational-episodic network. *PloS ONE*, *10*(11), e0143477.

Caplin, A. (2016). Measuring and modeling attention. *Annual Review of Economics*, *8*, 379–403.

Chávez, M. E., Villalobos, E., Baroja, J. L., & Bouzas, A. (2017). Hierarchical Bayesian modeling of intertemporal choice. *Judgment and Decision Making*, *12*, 19.

Chong, T. T.-J., Bonnelle, V., Manohar, S., Veromann, K.-R., Muhammed, K., Tofaris, G. K., ... Husain, M. (2015). Dopamine enhances willingness to exert effort for reward in Parkinson's disease. *Cortex*, *69*, 40–46.

Dasgupta, P., & Maskin, E. (2005). Uncertainty and hyperbolic discounting. *American Economic Review*, *95*(4), 1290–1299.

Figner, B., Knoch, D., Johnson, E. J., Krosch, A. R., Lisanby, S. H., Fehr, E., & Weber, E. U. (2010). Lateral prefrontal cortex and self-control in intertemporal choice. *Nature Neuroscience*, *13*(5), 538.

Foerde, K., Figner, B., Doll, B. B., Woyke, I. C., Braun, E. K., Weber, E. U., & Shohamy, D. (2016). Dopamine modulation of intertemporal decision-making: Evidence from Parkinson disease. *Journal of Cognitive Neuroscience*, *28*, 657–667.

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

Gabaix, X., & Laibson, D. (2017). *Myopia and discounting* (Tech. Rep.). National Bureau of Economic Research.

Green, L., Myerson, J., & McFadden, E. (1997). Rate of temporal discounting decreases with amount of reward. *Memory & Cognition*, *25*, 715–723.

Hardisty, D. J., Appelt, K. C., & Weber, E. U. (2013). Good or bad, we want it now: Fixed-cost present bias for gains and losses explains magnitude asymmetries in intertemporal choice. *Journal of Behavioral Decision Making*, *26*, 348–361.

Johnson, M. W., & Bickel, W. K. (2002). Within-subject comparison of real and hypothetical money rewards in delay discounting. *Journal of the Experimental Analysis of Behavior*, *77*, 129–146.

Kirby, K. N. (1997). Bidding on the future: evidence against normative discounting of delayed rewards. *Journal of Experimental Psychology: General*, *126*, 54–70.

Lebreton, M., Bertoux, M., Boutet, C., Lehericy, S., Dubois, B., Fossati, P., & Pessiglione, M. (2013). A critical role for the hippocampus in the valuation of imagined outcomes. *PLoS Biology*, *11*(10), e1001684.

Litman, L., Robinson, J., & Abberbock, T. (2017). Turkprime.com: A versatile crowdsourcing data acquisition platform for the behavioral sciences. *Behavior Research Methods*, *49*(2), 433–442.

Manohar, S. G., Chong, T. T.-J., Apps, M. A., Batla, A., Stamelou, M., Jarman, P. R., ... Husain, M. (2015). Reward pays the cost of noise reduction in motor and cognitive control. *Current Biology*, *25*, 1707–1716.

Matějka, F., & McKay, A. (2015). Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review*, *105*, 272–98.

McGuire, J. T., & Botvinick, M. M. (2010). Prefrontal cortex, cognitive control, and the registration of decision costs. *Proceedings of the National Academy of Sciences*, *107*, 7922–7926.

Peters, J., & Büchel, C. (2010). Episodic future thinking reduces reward delay discounting through an enhancement of prefrontal-mediotemporal interactions. *Neuron*, *66*(1), 138–148.

Raineri, A., & Rachlin, H. (1993). The effect of temporal constraints on the value of money and other commodities. *Journal of Behavioral Decision Making*, *6*, 77–94.

Rigoux, L., Stephan, K. E., Friston, K. J., & Daunizeau, J. (2014). Bayesian model selection for group studies–revisited. *Neuroimage*, *84*, 971–985.

Rouder, J. N., Speckman, P. L., Sun, D., Morey, R. D., & Iverson, G. (2009). Bayesian t tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review*, *16*, 225–237.

Shenhav, A., Botvinick, M. M., & Cohen, J. D. (2013). The expected value of control: An integrative theory of anterior cingulate cortex function. *Neuron*, *79*(2), 217–240.

Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, *50*, 665–690.

Sims, C. R. (2016). Rate-distortion theory and human perception. *Cognition*, *152*, 181–198.

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

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

Vincent, B. T. (2016). Hierarchical bayesian estimation and hypothesis testing for delay discounting tasks. *Behavior Research Methods*, *48*(4), 1608–1620.

Westbrook, A., & Braver, T. S. (2016). Dopamine does double duty in motivating cognitive effort. *Neuron*, *89*, 695–710.

Appendix

In order to derive the optimal precision $\bar{\sigma}_\varepsilon^2$, the expected quadratic loss is computed as follows. Conditioning on u and s,

$$
\begin{aligned}
(u-\hat{u})^2 &= (u - D_t s)^2 && (20)\\
&= (D_t(u-s) + u(1-D_t))^2 && (21)\\
&= D_t^2(u-s)^2 + 2D_t(1-D_t)u(u-s) + u^2(1-D_t)^2. && (22)
\end{aligned}
$$

Taking the expectation over $p(s|u)$, and subsequently over $p(u)$,

$$
\begin{aligned}
\mathbb{E}_s[(u-\hat{u})^2] &= \sigma_\varepsilon^2 t D_t^2 + 0 + u^2(1-D_t)^2 && (23)\\
\mathcal{L} = \mathbb{E}_u[\mathbb{E}_s[(u-\hat{u})^2]] &= \sigma_\varepsilon^2 t D_t^2 + \sigma_u^2(1-D_t)^2 && (24)\\
&= \frac{\sigma_\varepsilon^2 t(\sigma_u^2)^2}{(\sigma_u^2+\sigma_\varepsilon^2 t)^2} + \frac{\sigma_u^2(\sigma_\varepsilon^2 t)^2}{(\sigma_u^2+\sigma_\varepsilon^2 t)^2} && (25)\\
&= \frac{\sigma_u^2\sigma_\varepsilon^2 t}{\sigma_u^2+\sigma_\varepsilon^2 t}. && (26)
\end{aligned}
$$

We then plug this into the rate-distortion function for a Gaussian source (which reflects the rate-distortion frontier, that is, the minimal achievable information rate for a given distortion level, or equivalently the minimal achievable distortion for a given rate):

$$
\begin{aligned}
R &= \frac{1}{2}\ln\left(\frac{\sigma_u^2}{\mathcal{L}}\right) && (27)\\
&= \frac{1}{2}\ln\left(\sigma_u^2\left(\frac{\sigma_u^2+\sigma_\varepsilon^2 t}{\sigma_u^2\sigma_\varepsilon^2 t}\right)\right) && (28)\\
&= \frac{1}{2}\ln\left(1+\frac{\sigma_u^2}{\sigma_\varepsilon^2 t}\right), && (29)
\end{aligned}
$$

which can be rearranged to yield equation 5 in the main text.