Noisy Time Preference

Lisheng He (hlisheng@sas.upenn.edu)
University of Pennsylvania, Philadelphia, PA

Sudeep Bhatia (bhatiasu@sas.upenn.edu)
University of Pennsylvania, Philadelphia, PA

Abstract

People's desire to be patient or impatient can fluctuate from moment to moment, yet little is known about the effects of variability in time preference on intertemporal choice behavior. We examine this issue through the lens of an exponential discounting model with noisy discount factors. We show that such a model can generate decreasing patience over time, accounting for behavioral patterns typically attributed to hyperbolic discounting, while also making reasonable predictions regarding violations of intertemporal dominance. Additionally, two experiments reveal that many participants do display noise in their discount factors, and that a noisy discount factor model outperforms hyperbolic models in terms of quantitative fit. Ultimately the majority of participants are best described by some type of exponential discounting model (with or without noisy discount factors). These results indicate that it may not be necessary to assume alternate forms of non-exponential discounting, as long as the discount factors in an exponential model are permitted to vary at random. These results also highlight the importance of allowing for different sources of noise in choice modeling.

Keywords: decision making; intertemporal choice; noise; variability; computational modeling

Introduction

Random noise plays a central theoretical role in psychological research on high-level cognition. The assumption of noise not only explains variability in individuals' responses across multiple identical trials; When allowed to interact with attention, memory, and valuation, unsystematic noise is also capable of generating a systematic effect on behavior. In recent years, this type of unsystematic noise has been shown to account for error and response time patterns in perceptual and lexical choice, biases in probability judgment and social judgment, paradoxes in risky decision making, and the appearance of inconsistent or intransitive preferences (Bhatia & Loomes, 2017; Brown & Heathcote, 2008; Costello & Watts, 2014; Denrell, 2015; Erev, Wallsten & Budescu, 1994; Hilbert, 2012; Howes et al., 2016; Ratcliff & Rouder, 1998; Regenwetter, Dana & Davis-Stober, 2011; Tsetsos et al., 2016). In many of these cases unsystematic noise is enough, by itself, to provide a full account of observed behavioral patterns, making additional --more complex-- psychological assumptions about explicit biases in the judgment or decision process unnecessary.

In this paper we provide a formal characterization and analysis of the role of noise in intertemporal choice, that is, choice between payoffs occurring at different points in time. Our approach is motivated by the recent theoretical claims of Bhatia and Loomes (2017), who suggest that there are two key sources of noise in the preferential choice process. The first involves noise in response generation, with decision makers occasionally making mistakes in translating their preferences into choices. The second involves noise in the preferences themselves, with the parameters that characterize these preferences fluctuating from trial to trial. Bhatia and Loomes (2017) apply both sources of noise within a "rational" expected utility theory framework in a set of risky choice tasks, and show that the resulting model can predict (seemingly irrational) violations of EUT, such as choice patterns commonly seen to support Prospect Theory accounts of risk taking.

In intertemporal choice, it is exponential discounting that is considered to be the rational or normative model. In this paper, we propose an exponential discounting model that allows for trial-to-trial variability in discount factors, as well as random mistakes in generating responses, and then examine the properties of this model with both simulations and experiments. Our analysis tests the descriptive boundaries of the exponential discounting model and evaluates when it is and is not necessary to deviate from this rational theory to describe irrational patterns in intertemporal choice data. By performing these tests, we hope to obtain a deeper understanding of the effects of noise in intertemporal choice, complementing the rich existing theoretical literature on variability in cognition and behavior.

Intertemporal Discounting

The simplest intertemporal choice task requires a decision maker to evaluate an option X offering a payoff x with a time delay of t. Discounting models of intertemporal choice assume that these evaluations involve the calculation of a discounted utility, which weighs the payoff based on the magnitude of the time delay. Thus, for a discount function $d(\cdot)$, the utility of X is given by: [1]

[1] In a choice set consisting of many different payoffs with different time delays, the payoff with the highest discounted utility, according to Equation 1, is the one that is chosen. When each option offers multiple payoffs (each with a different time delay), the payoffs are individually discounted based on their time delay, and aggregated into a single utility measure. Note that it is sometimes assumed that payoffs are transformed non-linearly according to a value function, prior to being discounted. However, for expositional clarity, we will avoid this assumption for the purposes of this paper. Our results should not vary with more complex assumptions regarding payoff valuation.

$$U(X) = d(t) \cdot x \quad (1)$$

Exponential discounting, initially introduced by Samuelson (1937), involves a particularly parsimonious discount function. For a discount factor δ, it assumes that the discounted weight on the payoff value is simply given by:

$$d(t) = \delta^t \quad (2)$$

where $0 \leq \delta \leq 1$. Smaller values of δ correspond to increased discounting and lead to smaller weights on later payoffs relative to sooner payoffs. $\delta = 1$ corresponds to a complete absence of discounting of delayed payoffs. This one parameter discount function is the most commonly used discounting function in economics, as well in various applications in psychology, such as reinforcement learning.

Exponential discounting is, however, limited from a descriptive perspective. Notably it is unable to account for observed patterns of *decreasing impatience* for intertemporal choices. Consider, for example, a choice between an option X^P offering \$5 immediately and option Y^P offering \$10 in one month (proximal choice), as well as between option X^R offering \$5 in one month and option Y^R offering \$10 in two months (remote choice). As both the payoffs and the difference in time delays are the same for the proximal and remote choice, exponential discounting, with a fixed discount factor δ, predicts that participants should either select the sooner payoff in both choices or the later payoff in both choices. However, some studies suggest that participants typically select the sooner payoff in the proximal choice, but the later payoff in the remote choice (Green, Fristoe, & Myerson, 1994; Kirby & Herrnstein, 1995; Thaler, 1981; see Frederick, Loewenstein, & O'Donoghue, 2002 for a review). In response to such violations, researchers have suggested that the shape of the discounting function is not exponential but hyperbolic (e.g., Laibson, 1997; Loewenstein & Prelec, 1992; Mazur, 1987; see Table 1 for a representative list of hyperbolic models).

Noise in Intertemporal Choice

One critical issue with the above models is their inability to account for stochasticity inherent in human behavior. In order to use discounting models to describe stochastic choice data, discounting models need to be recast in probabilistic terms. Exponential discounting is generally modelled alongside some assumption of response noise, typically in the form of a logistic choice rule that transforms discounted utilities into choice probabilities (McFadden, 1973). Here, for options X and Y offering payoffs x and y with time delays t and s respectively, the probability of selecting X over Y is given by:

$$\Pr[X \text{ chosen}] = \frac{1}{1 + \exp\{-\theta(d(t) \cdot x - d(s) \cdot y)\}} \quad (3)$$

where $\theta \geq 0$ is a parameter that determines the extent of noise in the choice process. Smaller values of θ correspond to noisier choices, with $\theta = 0$ generating completely random choice (i.e. X and Y equally likely to be chosen, regardless of underlying payoffs and time delays).

There is also, however, another source of randomness in choice: preference noise. The parameters of utility-based models often provide a formal representation of decision makers' preferences. These preferences may not be constant over the time course of an experiment; that is, they may themselves fluctuate in a noisy manner (Becker, DeGroot & Marschak, 1963; Loomes & Sugden, 1995; Regenwetter & Marley, 2001). Within an exponential discounting model, this type of variability would correspond a distribution of discount factors described by a probability density function $f(\delta)$. δ varies from trial to trial, according to f, causing the discount function and thus the option utilities to vary from trial to trial. In a given trial, the option with the higher utility contingent on the sampled δ would be chosen. In a choice between option X offering payoff x with delay t, and option Y offering payoff y with delay s, the probability of choosing X is given by:

$$\Pr[X \text{ chosen}] = \int g(X, Y|\delta)\, f(\delta)\, d\delta$$

$$g(X, Y|\delta) = \begin{cases} 1 & if\ \delta^t x > \delta^s y \\ 0.5 & if\ \delta^t x = \delta^s y \\ 0 & if\ \delta^t x < \delta^s y \end{cases} \quad (4)$$

In such a formulation, $E[\delta]$ is the expected discount factor, and can be seen as characterizing the decision makers' underlying time preference. Although this underlying time preference is stable, trial-to-trial variability in δ could alters decision makers' utilities when they are exposed to the same decision problems repeatedly, thus leading to occasional mistakes in choice.

Although both response and preference noise do generate stochastic behavior, they are unable, by themselves, to account for violations of exponential discounting, such as decreasing impatience. This is because both types of noise, when applied individually, generate modal choice predictions that are in the direction of the prediction of the corresponding noiseless exponential model.

Of course both preference and response noise can influence intertemporal choice simultaneously. In this setting we would have both variability in discount factors for generating utilities, as well as variability in translating utilities into choice. Choice probabilities with such a model can be obtained by integrating Pr[X chosen] as defined in Equation 3, over the range of feasible values of δ, weighted by their respective probabilities. Thus, in a choice between option X offering payoff x with delay t, and option Y offering payoff y with delay s, the probability of choosing X when both types of noise are present would be given by:

$$\Pr[X \text{ chosen}] = \int g(X, Y|\delta)\, f(\delta)\, d\delta$$

$$g(X, Y|\delta) = \frac{1}{1 + \exp\{-\theta(\delta^t \cdot x - \delta^s \cdot y)\}} \quad (5)$$

These choice probabilities can deviate from the predictions of the corresponding deterministic exponential model (with discount factor of $E[\delta]$). The reason for this is that the utility difference between X and Y, $\delta^t \cdot x - \delta^s \cdot y$ is non-linear in δ. This means that variability in δ distorts the expected differences in utility between the two options, so that the expectation of the utility difference between X and Y, $E[U(X) - U(Y)]$, is not the same as the utility difference of these options, $U(X) - U(Y)$ for $E[\delta]$. When preference noise is applied by itself (as in Equation 4), this distortion does not alter modal choice, as the choice rule is based only on whether $\delta^t \cdot x > \delta^s \cdot y$ or $\delta^t \cdot x < \delta^s \cdot y$, and not on the magnitude of $\delta^t \cdot x - \delta^s \cdot y$. However, when these utility differences are combined with response noise (as in Equation 5) the distorted expected utility differences leads to distorted choice probabilities. Note that this can happen even if the distribution δ is symmetric around $E[\delta]$.

Properties

The exponential discounting model, with both response and preference noise, can account for violations of exponential discounting, such as decreasing impatience. As an illustration of this, consider again the proximal and remote choices in the decreasing impatience example above. If we only allowed for the response noise, and set $\theta = 1$ in Equation 3, we would obtain $\Pr[X^P$ chosen$] < 0.5$ in the proximal choice and $\Pr[X^R$ chosen$] < 0.5$ in the remote choice for all values of $\delta < 0.5$, and $\Pr[X^P$ chosen$] > 0.5$ in the proximal choice and $\Pr[X^R$ chosen$] > 0.5$ in the remote choice for all values of $\delta > 0.5$. This is shown in Figure 1a.

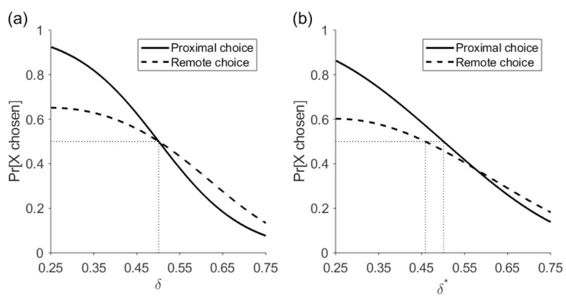

Figure 1: The probability of X chosen as a function of (mean) discount factor in proximal and remote choices. (a) Only response noise (θ = 1) is assumed. (b) Both response (θ = 1) and preference noise (η = 0.25) are assumed.

Now consider adding preference noise to this formulation, with $\delta \sim$ Uniform$[\delta^* - 0.25, \delta^* + 0.25]$. Note that $E[\delta] = \delta^*$. In this setting, we find that $\Pr[X^P$ chosen$] < 0.5$ in the proximal choice and $\Pr[X^R$ chosen$] < 0.5$ in the remote choice for $\delta^* < 0.457$, and $\Pr[X^P$ chosen$] > 0.5$ in the proximal choice and $\Pr[X^R$ chosen$] > 0.5$ in the remote choice for $\delta^* > 0.50$. For δ^* in the range (0.457,0.500) we obtain both $\Pr[X^P$ chosen$] < 0.5$ and $\Pr[X^R$ chosen$] > 0.5$, consistent with the finding of decreasing impatience. Note that this asymmetry emerges despite preference noise being unsystematic (i.e. δ distributed symmetrically around δ^*). This is shown in Figure 1b.

A combination of response and preference noise is also necessary for making reasonable predictions for intertemporal dominance. Consider, for example, a choice between an option X offering $10 immediately and option Y^{ND} offering $15 in one month (non-dominance choice), as well as another choice between option X, and option Y^D offering $7.50 immediately (dominance choice). For a decision maker with $\delta = 0.5$ we have $U(X) - U(Y^{ND}) = U(X) - U(Y^D) = 2.50$. As the difference in utilities is the same between X and Y^{ND} and between X and Y^D¸ an exponential choice model with only response noise (as in Equation 3) would predict the same choice probability of X in both cases. In other words, the decision maker would be equally likely to make a mistake and select the less desirable option in the non-dominance choice as in the dominance choice.

In reality decision makers can detect dominance. Although they do occasionally choose dominated options, the likelihood of doing so is much lower than that typically predicted by models equipped with only response noise (e.g., Busemeyer & Townsend, 1993; Loomes & Sugden, 1998).

In order to provide an adequate account of intertemporal dominance, we once again need both response and preference noise. For example, if we allow for $\theta = 1$, as well as $\delta \sim$ Uniform$[\delta^* - 0.25, \delta^* + 0.25]$, with $\delta^* = 0.5$, we obtain $\Pr[X$ chosen$] = 80.0\%$ in the non-dominance choice, but $\Pr[X$ chosen$] = 92.4\%$ in the dominance choice. Thus even though the difference in utilities between X and Y^{ND} and between X and Y^D is the same under $E[\delta] = \delta^*$, the probability of choosing X is higher when it dominates its competitor.

The intuition for the above choice patterns is straightforward: Response noise generates mistakes based on the utility differences between options, implying that dominance is violated too frequently when response noise is applied by itself. Preference noise, in contrast, never violates dominance. When a given δ is applied to the two options in a dominance trial, the utility for the dominating option is always greater than that for the dominated option, leading to a choice probability of 0% for the dominated option. The combination of response and preference noise results in an averaging of these two extreme predictions. Thus, in a model with both response and preference noise, it is possible to choose a dominated option, but the probability of this is smaller than the probability of choosing an equally desirable non-dominated option.

Experiments

We ran two experiments to further test the explanatory score of the exponential model with both response and preference noise.

Methods and Materials

A total of 89 undergraduate students from a university in United States participated in our two experiments: 44 participants (31 female; aged 20.26 ± 1.25) in Experiment 1 and 45 participants (25 female; aged 19.84 ± 1.49) in Experiment 2.

Experiment 1 involved hypothetical binary choices between an option X offering a payoff of x after a time delay t, and an option Y offering a payoff of y after a time delay $s = t + k$. We set x = $100 in all trials and chose t from the set {today, 3 months, 6 months, 9 months} and k from the set {3 months, 6 months, 9 months}. y was determined by applying annual interest rate from the set {-50%, 50%, 100%, 500%, 1000%} to the corresponding time delays. This generated a total of 60 unique choice pairs. Note that the use of a negative interest rate implied that 12 of these choice pairs involved a dominated option (offering a smaller reward with a larger time delay than its competitor).

As shown below, Experiment 1 involved fairly high choice probabilities for the delayed option *Y*. Although this should not alter our key conclusions, we wished to replicate our tests with stimuli generating roughly equivalent choice probabilities for *X* and *Y*. Thus we ran a second experiment, with stimuli generated using the methods above, but with annual interest rates in the set {-25%, 25%, 75%, 125%, 175%}. By using smaller interest rates we obtained smaller values of y for corresponding values of t and k, leading to higher choice proportions for *X*.

We excluded data from five participants in Experiment 1 and three participants in Experiment 2 because they constantly chose either *X* or *Y* in all the non-dominance choices. This left 39 participants in Experiment 1 and 42 participants in Experiment 2 for our analysis.

Model Fitting

The main goal of the two experiments was to test whether the exponential discounting model with both response and preference noise is able to provide a good quantitative account of choice data. For this purpose, we fit the core exponential discounting model embedded in a logistic choice rule for response noise, with a variable discount factor $\delta \sim$ Uniform$[\delta^* - \eta, \delta^* + \eta]$ (Equation 5), with $\delta^* - \eta \geq 0$ *and* $\delta^* + \eta \leq 1$. We refer to this model as the *noisy exponential model* for the remainder of this paper.

Table 1: Alternate hyperbolic discounting functions.

Function Name	Function Form	Domain
Mazur-1 hyperbolic	$d(t) = (1 + \alpha t)^{-1}$	$\alpha > 0$
Mazur-2 hyperbolic	$d(t) = (1 + \alpha t^{\tau})^{-1}$	$\alpha, \tau > 0$
LP hyperbolic	$d(t) = (1 + \alpha t)^{-\beta/\alpha}$	$\alpha, \beta > 0$
Quasi-hyperbolic	$d(t) = \begin{cases} 1, \text{when } t = 0 \\ \beta\delta^t, \text{when } t > 0 \end{cases}$	$0 \leq \beta \leq 1$ $0 \leq \delta \leq 1$

We also wished to contrast the predictions of the noisy exponential model with the various hyperbolic models proposed in prior work. We considered the one parameter hyperbolic discounting model proposed by Mazur (1987), which we refer to as *Mazur*-1 *hyperbolic*, as well as the two parameter hyperbolic model proposed by Mazur (1987), which we refer to as *Mazur*-2 *hyperbolic*. We also considered the two-parameter generalized hyperbolic discounting model proposed by Lowenstein and Prelec 1992), which we refer to as the *LP hyperbolic.* Finally, we considered the *quasi-hyperbolic* model proposed by Laibson (1997). All these hyperbolic discounting models are presented in Table 1. We also tested the predictive power of the baseline exponential model (Equation 2). All hyperbolic discounting models and the baseline exponential discounting model were embedded within the logistic choice function (Equation 4) to allow for response noise.

Results

Summary of Choice Data Among the non-dominance choices, we found that option *Y* was chosen 66.9% of the time in Experiment 1, and 48.6 % of the time in Experiment 2. The frequency of choosing the dominated options was comparable across the two experiments: Participants in Experiment 1 and Experiment 2 chose the dominated option (option *Y* offering both the smaller and the more delayed reward) 3.9% of the time and 4.7% of the time respectively.

We also tested for decreasing impatience in the two experiments. In the non-dominance trials of Experiment 1, participants chose option *Y* (offering the payoff with the larger delay) 65.9% when t = 0 months, 66.0% when t = 3 months, 67.7% when t = 6 months, and 67.8% when t = 9 months. Formally, the probability of choosing *Y* in the non-dominance trials increased with t in a mixed-effect logistic regression model with random intercepts for participants intervals between options and the implied interest rates in the trial (β = 0.028, z = 2.397, p = .017). A similar test applied individually to each participant found that the effect of t on the choice probability of *Y* was positive and significant ($p <$ 0.05) for six participants, positive and non-significant ($p >$ 0.05) for 13 participants, negative and significant for one participant, and negative and non-significant for 19 participants. These results provide some evidence of decreasing impatience on the aggregate level, but also suggest substantial heterogeneity in this choice pattern across participants, with roughly half the sample showing significant or non-significant decreasing impatience, and the other half showing significant or non-significant increasing impatience.

We obtained more ambiguous results in Experiment 2. In this experiment we observed a choice frequency for option *Y* of 47.6% when t = 0 months, 49.0% when t = 3 months, and 48.9% when t = 6 months, and 49.1% when t = 9 months. Although this choice frequency is increasing in t, it did not reach statistical significance in a mixed-effect model (β = 0.015, z = 1.447, p = 0.148). On the individual level, the effect of t on the choice probability of *Y* was positive and significant for three participants, positive and non-significant for 21 participants, negative and significant for one participant, and negative and non-significant for 17 participants. This time, a little bit more than half sample showed significant or non-significant decreasing impatience, and the other half showed significant or non-significant increasing impatience.

Overall, the findings of Experiments 1 and 2 suggest that decreasing impatience is not as robust as is widely held, but is consistent with some recent experiments with similar inconclusive effects for decreasing impatience (e.g., Kable & Glimcher, 2010; Read, 2001).

Best-fit Parameters. Of key interest to the tests in this paper is the preference noise η in the noisy exponential model. Figure 2 shows the individual-level and group-level estimates of η, as well as the estimates of δ^* from the noisy exponential model. As the baseline exponential model is nested in the noisy exponential model and η is the additional specification, we can evaluate the statistical significance of η using the likelihood ratio test. The variable being evaluated here is the ratio between the maximum likelihood values of the two models. This ratio has a chi-square distribution with a degree of freedom of 1. The likelihood ratio test reveals that the noisy exponential model has a significantly higher likelihood than the baseline exponential model (implying that restricting $\eta = 0$ results in significantly worse fits) on the group level ($\chi^2(1) > 3.84$, $p < .05$) as well as for 46% participants for Experiment 1, and for 42.9% of participants for Experiment 2.

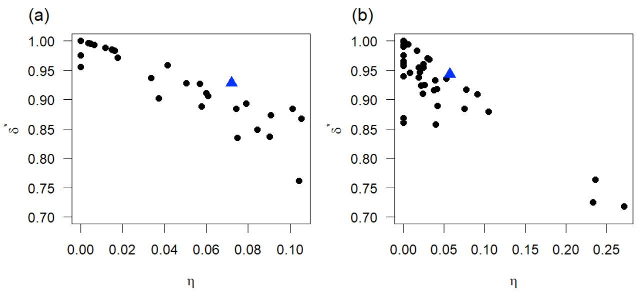

Figure 2: Distributions of individual-level (black dots) and group-level (blue triangles) estimates of η and δ^* from the noisy exponential model. (a) Estimates of Experiment 1.(b) Estimates of Experiment 2.

We can also perform a similar set of tests for the quasi-hyperbolic model and the LP hyperbolic model, both of which embed the baseline exponential discounting model. However, on an individual level only 17.9% of participants in Experiment 1 and 19.0% of participants in Experiment 2 are better fit by the quasi-hyperbolic model than the exponential model. These proportions are 15.4% and 38.1% for LP hyperbolic model in Experiments 1 and 2 respectively.

Model Comparisons The likelihood ratio tests shown above suggest that the noisy exponential model does better than the quasi-hyperbolic and LP hyperbolic models, as it provides a greater improvement over the baseline exponential model both on the group and on the individual level. However this test cannot be used to directly compare the noisy exponential model with the quasi-hyperbolic and LP hyperbolic models, as these models are not nested within each other. To perform such a test, we thus need to use the Bayes Information Criterion, calculated as BIC = $-2 \ln(L) + k \times \ln(n)$, where L is the maximum likelihood value, k is the number of free parameters and n is the number of data points. Lower values of BIC indicate better fits, controlling for model flexibility (quantified by the total number of free parameters). Table 2 shows the BIC values for the models fit to the group-level data in Experiments 1 and 2. Group-level fits impose the same parameters to all participants, not allowing for individual differences. It also shows aggregate BIC values for individual-level fits. As can be seen in Table 2, the best performing model according to the group-level BIC and the aggregate individual-level BIC, for both experiments, is the noisy exponential model. We can also examine the proportion of participants best fit by each of the models when all models are compared simultaneously. Here we find that the best performing model is the baseline exponential model, which has the lowest BIC values for 41.0% and 46.5% of participants in Experiments 1 and 2. This is followed by the noisy exponential model, which provides the lowest BIC values for 30.8% and 23.3% of participants in the two experiments. Thus, around 70% of the participants in our two experiments are best fit by either the baseline exponential or the noisy exponential model, according to BIC.

Summary and Discussion

This paper has examined the role of noise in intertemporal decision making through the lens of the exponential discounting model. We propose a modification to this model that allows for time preference to vary from moment to moment. Formally, this involves distribution over the discount factors that quantify time preference. We have shown how this noisy exponential model can be used to predict seemingly irrational patterns of behavior, such as decreasing impatience. A noisy exponential model also provides a better account of violations of intertemporal dominance.

Empirically, we have tested the quantitative properties of our proposed model and the model fitting exercise has revealed a number of novel insights regarding the effect of noise in intertemporal discounting. Firstly, an examination of best fitting parameters has shown that the level of noise in

Table 2: Summary of model fit in Experiments 1 and 2.

Model Name	Group-level BIC		Aggregate Individual-level BIC		Percentage of Best Fits	
	Exp 1	Exp 2	Exp 1	Exp 2	Exp 1	Exp 2
Mazur-1 hyperbolic	9423	11633	6425	7357	17.9%	20.9%
Mazur-2 hyperbolic	9410	11560	6118	6772	2.6%	2.3%
LP hyperbolic	9408	11546	6043	6690	5.1%	0.0%
Quasi-hyperbolic	9413	11547	6068	6706	2.6%	7.0%
Exponential	9414	11556	5973	6610	**41.0%**	**46.5%**
Noisy exponential	**8996**	**11382**	**5867**	**6587**	30.8%	23.3%

discount factors, on the group level, is significantly greater than zero. This is also the case on the individual level for more than 40% of our participants, across the two experiments (we observe positive noise for most the remaining participants, but this does not reach statistical significance). We also find that the noisy exponential model outperforms the four hyperbolic models considered in this paper, in terms of quantitative fit. This emerges on both on the individual and the group level, for both the experiments. Conversely, on the individual level, hyperbolic discounting models do not provide better fits even relative to the baseline exponential model, although the data provide some evidence for decreasing impatience. This suggests that hyperbolic discounting models either over predict decreasing impatience or, if flexible enough to reduce to exponential discounting, the premium in model fits does not overcome the penalty of model complexity in model selection.

The results of this paper indicate that a rational model of intertemporal decision making that permits (unsystematic) variability in the degree of time preference has tremendous explanatory power. By doing so, it complements a rich existing literature in psychology on the descriptive role of random noise in cognition and behavior. Noise is not just useful only for accommodating observed variability in peoples' behavior. Rather, it occupies a central theoretical position of our understanding of this behavior.

Acknowledgments

We thank Wenjia Zhao and Zijin Zhang for their assistance for data collection.

References

Becker, G. M., DeGroot, M. H., & Marschak, J. (1963). An experimental study of some stochastic models for wagers. *Systems Research and Behavioral Science*, *8*(3), 199-202.

Bhatia, S. & Loomes, G. (2017). Noisy preferences in risky choice: A cautionary note. *Psychological Review,124*(5), 678-687.

Brown, S. D., & Heathcote, A. (2008). The simplest complete model of choice response time: Linear ballistic accumulation. *Cognitive Psychology*, *57*(3), 153-178.

Busemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, *100*(3), 432-459.

Costello, F., & Watts, P. (2014). Surprisingly rational: Probability theory plus noise explains biases in judgment. *Psychological Review, 121*(3), 463-480.

Denrell, J. C. (2015). Reference-dependent risk sensitivity as rational inference. *Psychological Review*, *122*(3), 461-484.

Erev, I., Wallsten, T. S., & Budescu, D. V. (1994). Simultaneous over-and under-confidence: The role of error in judgment processes. *Psychological Review*, *101*(3), 519-527.

Frederick, S., Loewenstein, G., & O'Donoghue, T. (2002). Time discounting and time preference: A critical review. *Journal of Economic Literature, 40*(2), 351–401.

Green, L., Fristoe, N., & Myerson, J. (1994). Temporal discounting and preference reversals in choice between delayed outcomes. *Psychonomic Bulletin & Review, 1*(3), 383-389.

Hilbert, M. (2012). Toward a synthesis of cognitive biases: how noisy information processing can bias human decision making. *Psychological Bulletin*, *138*(2), 211- 237.

Howes, A., Warren, P. A., Farmer, G., El-Deredy, W., & Lewis, R. L. (2016). Why contextual preference reversals maximize expected value. Psychological review, 123(4), 368-391.

Kable, J. W., & Glimcher, P. W. (2010). An "as soon as possible" effect in human intertemporal decision making: behavioral evidence and neural mechanisms. *Journal of Neurophysiology*, *103*(5), 2513-2531.

Kirby, K. N. & Herrnstein, R. J. (1995). Preference reversals due to myopic discounting of delayed reward. *Psychological Science, 6*(2), 83–89.

Laibson, D. (1997). Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics, 112*(2), 443–477.

Loewenstein, G. & Prelec, D. (1992). Anomalies in intertemporal choice: evidence and an interpretation. *Quarterly Journal of Economics, 107*(2), 573-597.

Loomes, G., & Sugden, R. (1995). Incorporating a stochastic element into decision theories. *European Economic Review, 39,* 641-648.

Loomes, G., & Sugden, R. (1998). Testing different stochastic specifications of risky choice. *Economica, 65,* 581-598.

Mazur, J. E. (1987). An adjusting procedure for studying delayed reinforcement. In M. L. Commons, J. E. Mazur, J. A. Nevin, & H. Rachlin (Eds.), *The Effect of Delay and of Intervening Events on Reinforcement Value: Quantitative Analyses of Behavior* (pp. 55-73). Hillsdale, NJ: Lawrence Erlbaum.

Ratcliff, R., & Rouder, J. N. (1998). Modeling response times for two-choice decisions. *Psychological Science*, *9*(5), 347-356.

Read, D. (2001). Is time-discounting hyperbolic or subadditive? *Journal of Risk and Uncertainty*, *23*(1), 5-32.

Regenwetter, M., & Marley, A. A. J. (2001). Random relations, random utilities, and random functions. *Journal of Mathematical Psychology*, *45*(6), 864-912.

Regenwetter, M., Dana, J., & Davis-Stober, C. P. (2011). Transitivity of preferences. *Psychological Review*, *118*(1), 42-56.

Samuelson, P. A. (1937). A note on measurement of utility. *The Review of Economic Studies, 4*(2), 155-161.

Thaler, R. H. (1981). Some empirical evidence on dynamic inconsistency. *Economics Letters, 8*(3), 201-207.

Tsetsos, K., Moran, R., Moreland, J., Chater, N., Usher, M., & Summerfield, C. (2016). Economic irrationality is optimal during noisy decision making. *Proceedings of the National Academy of Sciences, 113*(11), 3102-3107.