Published in final edited form as:

J Risk Uncertain. 2016 June ; 52(3): 233–254. doi:10.1007/s11166-016-9242-y.

On the Functional Form of Temporal Discounting: An Optimized Adaptive Test

Daniel R. Cavagnaro,

Mihaylo College of Business and Economics, California State University Fullerton

Gabriel J. Aranovich,

University of California, San Francisco and San Francisco VA Medical Center

Samuel M. McClure,

Stanford University

Mark A. Pitt, and

Department of Psychology, The Ohio State University

Jay I. Myung

Department of Psychology, The Ohio State University

Abstract

The tendency to discount the value of future rewards has become one of the best-studied constructs in the behavioral sciences. Although hyperbolic discounting remains the dominant quantitative characterization of this phenomenon, a variety of models have been proposed and consensus around the one that most accurately describes behavior has been elusive. To help bring some clarity to this issue, we propose an Adaptive Design Optimization (ADO) method for fitting and comparing models of temporal discounting. We then conduct an ADO experiment aimed at discriminating among six popular models of temporal discounting. Rather than supporting a single underlying model, our results show that each model is inadequate in some way to describe the full range of behavior exhibited across subjects. The precision of results provided by ADO further identify specific properties of models, such as accommodating both increasing and decreasing impatience, that are mandatory to describe temporal discounting broadly.

Keywords

temporal discounting; intertemporal choice; adaptive designs; design optimization; model selection

The commonplace and readily understood notion that delaying gratification is unpleasant has become one of the most thoroughly explored constructs in the behavioral sciences. In the scientific literature, the phenomenon is referred to as delay, or temporal, discounting, and is more precisely defined as the amount by which the subjective value of a reward decreases as a function of delay to delivery. Temporal discounting has come to be viewed as

indispensable to the analysis of a wide range of important decision making behaviors, from saving behavior to environmental policy (Laibson, 1997; Frederick et al., 2002; Dasgupta, 2008). The study of temporal discounting has generated promising applications in psychiatry and neuroscience, wherein distinct patterns of discounting behavior have been linked to various types of mental illness, including addiction, gambling, ADHD, and other disorders associated with impaired impulse-control (Koffarnus et al., 2013; Sharp et al., 2012; Reynolds, 2006; Story et al., 2014). Studies examining the neural correlates of temporal discounting have led to significant advances in our understanding of how decision making is realized in the brain (McClure et al., 2004, 2007; Kable and Glimcher, 2007; Rangel et al., 2008). In essence, temporal discounting has been adopted as a behavioral biomarker of impulsivity (Peters and Büchel, 2011), to the extent that targeting temporal discounting directly with cognitive training has become a promising intervention for self-control impairment (Bickel et al., 2011).

In scientific studies, temporal discounting behavior is described quantitatively with a discounting curve, which ostensibly characterizes the extent of decline in value as a function of delay. Precise estimates of the shape of the discounting curve are obtained by fitting a parametric function to choice data from experiments that require participants to choose between reward options that vary in magnitude and delay to delivery. The resulting estimates are subsequently used for tests of group differences in temporal discounting behavior (Reynolds, 2006), for correlation with neural and physiologic data (McClure et al., 2007; Kable and Glimcher, 2007; Moore and Cusens, 2010; Rangel et al., 2008), for treatment targeting (Bickel et al., 2011), and for making behavioral predictions (Dallery and Raiff, 2007; MacKillop and Kahler, 2009; Story et al., 2014). The precise shape of the fitted curve depends on which functional form is assumed, and has important predictive implications for unobserved behavior. Therefore, choosing the right function is critical to the success of the wide-ranging applications of the temporal discounting paradigm.

Despite decades of investigation, the proper parametric form of the discounting curve remains an active area of research. Bodies of literature have emerged across disciplines, proposing different theoretical constructs to account various patterns of observed choice behavior (Frederick et al., 2002; van den Bos and McClure, 2013; Doyle, 2013). Candidate models vary in complexity and include the Exponential, Hyperbolic, generalized Hyperbolic (Green and Myerson, 2004), Constant Sensitivity (Ebert and Prelec, 2007), Double Exponential (McClure et al., 2007), and several others. Among them, the hyperbolic has become the dominant model in the literature, due largely to its success in fitting individual data in experiments and the simplicity of its one-parameter form (van den Bos and McClure, 2013). In published reviews of the temporal discounting literature, various explanations for the elusiveness of consensus have been advanced, including context-dependence and inter-individual variability of temporal discounting preferences, as well as the possibility that tasks designed to elicit time discounting preferences disregard the complexity of temporal discounting decisions, which likely involve several component processes in addition to time preference (Frederick et al., 2002; Berns et al., 2007; Scholten and Read, 2010; Peters and Büchel, 2011; van den Bos and McClure, 2013).

Another possible reason why a consensus has yet to develop is that discriminating among functional forms at the individual level presents formidable methodological challenges. Due to practical limitations on the number of questions that can be asked in a single experiment, the questions that are asked must be finely tuned to the task of discriminating among the models and estimating their parameters as precisely as possible. Moreover, due to individual differences, the right questions to ask to best distinguish among models will differ across participants. Ideally, the set of questions in an assessment would be tuned to each individual, rather than a one-size-fits-all approach, yet parameter estimation and model comparison studies in the discounting literature have often employed one-size-fits-all designs, either using well-studied instruments such as the Kirby Monetary Choice Questionnaire (Kirby et al., 1999), or ad hoc designs developed through a mixture of literature search and personal judgment (Madden and Bickel, 2010).

Additional challenges lie in eliciting temporal discounting preferences, and in the statistical analysis of experiment data for fitting and comparing models. By far the most common method involves presenting subjects with a series of binary choices between smaller-sooner and larger-later monetary rewards, which are systematically varied by amount or delay in order to aid in the identification of indifference points, to which statistical techniques such as nonlinear regression are applied for curve fitting. More recently, maximum likelihood estimation (MLE) has become a popular alternative to analyses based on indifference points (McClure et al., 2004, 2007; Pine et al., 2009, 2010). However, most previous model comparison studies (Myerson and Green, 1995; Rachlin, 2006; Takahashi et al., 2008; McKelcher et al., 2009; Pine et al., 2009, 2010; Peters et al., 2012) have relied on the coefficient of determination (R^2) to determine the goodness-of-fit of the models under consideration, even though this approach can be misleading due to the problem of over-fitting (Pitt and Myung 2002), although several recent studies (Takahashi et al., 2008; Pine et al., 2009; Peters et al., 2012; Abdellaoui et al., 2013) have used AIC for model comparison, which does account for model complexity.¹

The present study demonstrates and implements a Bayesian inference method for discriminating among models of temporal discounting using Adaptive Design Optimization (Cavagnaro et al., 2010; Myung et al., 2013, ADO), which integrates likelihood-based data-modeling with adaptive experimental designs to maximize the efficiency and informativeness of an experiment. In an ADO experiment, stimuli are tailored to each participant by updating model and parameter estimates in real time as data are collected, and using the latest estimates to select stimuli that maximize the expected information gain about the models under consideration. The approach has proved to be effective for discriminating among models of memory retention (Cavagnaro et al., 2011) risky choice (Cavagnaro et al., 2013a) and among functional forms of the probability weighting function in Cumulative Prospect Theory (Cavagnaro et al., 2013b). The present work extends the approach to temporal discounting.

¹In addition, adaptive titration procedures have been developed, which adjust the amount or delay in response to the subject's choices on a trial-by-trial basis, thereby reducing the risk of floor and ceiling effects (Mazur, 1987; Johnson and Bickel, 2002). Several studies have also employed adaptive heuristics in which stimuli depend on subjects' responses. For instance, Abdellaoui et al. (2010) employed a two-stage design in which the first stage measured utilities, separately for each subject, which were then used in the second stage to construct stimuli for measuring time weights.

This work builds on a growing body of research in adaptive designs in experimental economics (e.g., Toubia et al., 2013; Ray et al., 2012; Wang et al., 2010), and the relationship between model evaluation and experimental design (Broomell and Bhatia, 2014). Like ADO, the method called DEEP (Toubia et al., 2013) also updates model estimates in real-time and optimizes stimuli based on expected information gain. The main difference between the two methods is that DEEP was originally formulated for estimating the parameter(s) of an assumed model of preference, whereas ADO was originally formulated for selecting among a group of candidate models (Cavagnaro et al., 2010). However, in the Bayesian framework, model selection and parameter estimation are essentially the same problem, and the natural extension of ADO to the problem of estimating the parameter of an assumed model was made explicit by Myung et al. (2013).

In the present study, we present an ADO method for eliciting temporal discounting preferences and modeling individual-level discounting behavior. The goal of the experiment was to determine which model, among a group of the most prominent in the literature, best accounts for discounting behavior. The data from the ADO-based experiment conclusively discriminate among the six functional forms listed above, at the individual level, but we find that there is substantial heterogeneity in terms of both the qualitative shape of the discounting curve and its functional form. In particular, about 25% of subjects showed increasing impatience (i.e., concavity in the discounting curve), a distinct pattern of behavior that commonly utilized models cannot accommodate. The hyperbolic model was very rarely the preferred model, even using model selection criteria that give it a boost for its simplicity. Exponential and beta-delta discounting also performed poorly. Overall, we found that the Double Exponential and Constant Sensitivity models provide the best explanation for the largest number of subjects. Implications of these findings are discussed in Section 5.

We also tested the ADO method in simulation experiments, pitting it against several benchmark elicitation methods to see how quickly and precisely each can recover a predetermined generating model from among a group of the most prominent models in the literature. We find ADO to be vastly more effective than each benchmark, including the most widely used instrument in applications, the Kirby Monetary Choice Questionnaire (Kirby et al., 1999). Results also suggest that the method is also robust to misspecification of the stochastic error component of the models. Details of the simulation experiments are reported in the online appendix.

The rest of the paper is organized as follows. Section 1 describes the discounting models that we consider in our analysis. Section 2 gives some background to the ADO methodology and the novelties in the current implementation. Section 3 reports the experiment procedure and results, and Section 4 concludes with additional discussion of the results and their implication for both temporal discounting and the ADO methodology.

1 Models

It is beyond the scope of this study to list all of the models that have been advanced to describe temporal discounting choice behavior (see van den Bos and McClure, 2013; Frederick et al., 2002; Doyle, 2013, for reviews). In this section, we describe the particular

set of models chosen for our experiment. The set of candidate models was not intended to be exhaustive. Our aim was to include the most prominent discounting models based on our assessment of the literature, as well as models that accommodate qualitatively different patterns of behavior and involve varying levels of complexity.

As mentioned above, temporal discounting models are designed to capture the extent to which the subjective value (V) of a reward of given magnitude (A) decreases as a function of delay to delivery (t). This can be summarized as

$$V = AD_t \quad (1)$$

where D_t is the individual-specific discount factor, with value between 0 and 1, generated by the model, as a function of t .

The standard discounting model in classical economic theory, introduced by Paul Samuelson (Samuelson, 1937), is the exponential model:

$$D_t = e^{-rt} \quad (2)$$

This model is distinguished by a *constant* rate of discounting indicated by the single parameter, $r > 0$. Though Samuelson specifically disavowed any assertion of the psychological accuracy of this model, it was swiftly adopted as the most popular model in economic theory.

The limitations of the exponential model's psychological plausibility became clear in a series of studies, which demonstrated that neither animals (Ainslie and Herrnstein, 1981; Green et al., 1981), nor humans (Kirby, 1997; Thaler, 1981) appear to discount exponentially. Study after study indicated that the rate of discounting exhibited by subjects' choices was not constant but decreasing, and that so-called preference reversals (Ainslie, 1975), which are not permitted by exponential discounting, were the rule, rather than an exception. In response to these limitations, the now dominant hyperbolic discounting model was proposed (Mazur, 1987), which accommodates declining discount rates, as well as preference reversals:

$$D_t = \frac{1}{1+kt} \quad (3)$$

Like the exponential model, the hyperbolic model has a single parameter, here denoted k . Larger values of $k > 0$ indicate greater impatience. The superiority of the hyperbolic model over the exponential is well-established (Frederick et al., 2002; Green and Myerson, 2004).

The tendency of the hyperbolic model to over-estimate subjective value at short delays, while under-estimating it at longer delays led to the introduction of the generalized Hyperbolic model (Loewenstein and Prelec, 1992; Green and Myerson, 2004):

$$D_t = \frac{1}{(1+kt)^s} \quad (4)$$

It has been suggested that s parameter captures individual differences in the scaling of delay or in time-perception (Rachlin, 2006; Zauberman et al., 2009).

The discrete time quasi-hyperbolic, or $\beta - \delta$, model was introduced by Phelps and Pollak (1968), and popularized by Laibson (1997), in order to summarize the cardinal qualitative feature of the hyperbolic model, the declining discount rate, in a simplified form that was convenient for economic applications:

$$D_t = \begin{cases} 1 & \text{when } t=0 \\ \beta\delta^t & \text{when } t>0 \end{cases} \quad (5)$$

The $\beta \in (0,1)$ parameter in Equation 5 is intended to capture the fact that the “present” is privileged in hyperbolic discounting, while all subsequent periods are discounted at a more moderate rate. The $\beta - \delta$ model was not originally intended to describe laboratory behavior and, given its discontinuous, stylized behavior, we did not expect it to perform well in competition with the other models under consideration in this study.

One obvious limitation of the $\beta - \delta$ model is the discontinuity at $t = 0$, which seems psychologically implausible. In a follow up study, Laibson and colleagues (McClure et al., 2007) generalized the $\beta - \delta$ model by assuming that the delay discounting function can be decomposed into two component functions, representing separate cognitive systems with different levels of patience. The resulting “Double Exponential” model can be expressed as follows:

$$D_t = \omega e^{-rt} + (1-\omega)e^{-st} \quad (6)$$

Here, $\omega \in (0,1)$ signifies the relative contributions of the two components, with each component’s degree of impatience indicated by the magnitude of the parameter in its respective exponential term. In contrast to the other models, the double-exponential has three parameters to be estimated: ω , $r > 0$, and $s > 0$. This model is consistent with psychological models of time preference which posit that an emotional, impulsive system competes with more far-sighted deliberation for behavioral control (Metcalf and Mischel, 1999; Shefrin and Thaler, 1988; Loewenstein, 1996; McClure and Bickel, 2014).

Recent years have seen renewed interest in developing alternative models of temporal discounting. Among these newer models, we selected for inclusion the axiomatically-derived “Constant-Sensitivity” model (Ebert and Prelec, 2007):

$$D_t = e^{-(rt)^s} \quad (7)$$

Though it resembles the exponential model (and includes it as a special case when $s=1$), the additional parameter allows the constant-sensitivity model to accommodate *increasing* impatience. Recent evidence has highlighted the prevalence of increasing impatience at the individual level (Bleichrodt et al., 2009; Abdellaoui et al., 2010; Attema et al., 2010; Abdellaoui et al., 2013). In terms of the discounting curve, increasing impatience means concavity (i.e., the discounting curve gets steeper as the delay to reward increases). In particular, for $s > 1$, the Constant Sensitivity discounting curve is convex for $0 \leq t \leq \frac{1}{r}$. In contrast, standard approaches to modeling temporal discounting typically accommodate only *decreasing* impatience (i.e., convexity) which is the prevailing empirical finding at the aggregate level. At the extreme, the constant sensitivity model accommodates an “extended present” (Ebert and Prelec, 2007), in which a decision-maker is (nearly) indifferent to delay within the “near-future,” while discounting the “far-future” heavily. The length of the near-future is determined by the parameter r . For values of $s > 1$, discounting of subjective value increases as a function of t until $t = \frac{1}{r}$, after which discounting decreases as a function of t . Extreme values of s can lead to a “present-future dichotomy”, where the rewards delivered during the extended present are valued equally, while “future” rewards have zero value. Finally, the constant-sensitivity model permits between-subjects comparison of the length of the “present,” an intriguing quantity that has yet to be analyzed in the literature, to the best of our knowledge.

It should be noted that several of these models are nested within others in the set: Exponential within Constant-Sensitivity, Hyperbolic within generalized Hyperbolic, and both Exponential and Beta-Delta within Double-Exponential. Putting these models in competition with one another, using ADO to maximally discriminate them, allows us to compare families of models while also testing the extent to which additional parameters are justified within each family.

In addition to the six discounting models defined above, we will also consider an additional, “Coin Flip” model in our analysis. The Coin Flip model assumes that, on every trial, regardless of the stimulus, the smaller-sooner and larger-later options are equally likely to be chosen (e.g., the choice is determined by a coin flip). The coin flip model serves as a parameter-free null-hypothesis that allows for all of the substantive discounting models to fail.

2 Method: ADO for discriminating among temporal discounting models

To discriminate among the models described above, we bring to an ADO method that is novel to the study of temporal discounting. Our implementation of ADO in the current study closely follows that of Cavagnaro et al. (2013a) and Cavagnaro et al. (2013b), which used ADO to discriminate among models of risky choice. After giving some historical background and describing the general framework of ADO, the rest of this section reports

the aspects of our implementation that are novel or unique to testing models of temporal discounting. Additional, technical details of the algorithm are provided in the Appendix.

2.1 history and general framework of ADO

ADO is a computer-based experimentation methodology that is very general in its formulation. At the core of ADO is design optimization (DO), which is a statistical technique that selectively chooses design variables (e.g., treatment levels and values, presentation schedule) with the aim of identifying an experimental design that will produce the most informative and useful experimental outcome (Myung and Pitt, 2009). However, DO is a one-shot process in which an optimal design is identified, applied in an experiment, and then data modeling methods are used to aid in interpreting the results. ADO, on the other hand, is a dynamic process that treats the full experiment as a sequence of mini-experiments², and combines design optimization at each mini-experiment with real-time Bayesian updating of parameter estimates and model probabilities between mini-experiments, as shown in Figure 1. This allows the design of the next mini-experiment to be optimized on the fly as each new data point is collected and analyzed. The result is an efficient and informative method of scientific inference.

The optimization of experimental designs has a long history in statistics dating back to the 1950s (e.g. Lindley, 1956). Psychometricians have been doing adaptive experiments for decades in computerized adaptive testing (e.g., Weiss and Kingsbury, 1984), and psychophysicists have developed their own adaptive tools (e.g., Kontsevich and Tyler, 1999; Kujala and Lukka, 2006). Building on these works, as well as the Bayesian adaptive design framework of Chaloner and Verdinelli (1995)), Cavagnaro et al. (2010) introduced Adaptive Design Optimization as a general methodological framework for adaptive designs intended for discriminating among nonlinear mathematical models in cognitive science. The same ADO framework has since been adopted for discriminating among models of risky choice (Cavagnaro et al., 2013a), and applied to problems of discriminating among memory retention functions and among probability weighting functions (Cavagnaro et al., 2011, 2013b). Several related examples of adaptive methodologies have arisen independently in the economics and business literature, including BROAD (Ray et al., 2012), DOSE (Wang et al., 2010), and DEEP (Toubia et al., 2013). While these methods differ in some of the specifics of their implementation, they share the common goal of optimizing experimental design in real-time in order to improve the quality of statistical inference.

2.2 Design variable

Although ADO can theoretically be employed to optimize any quantifiable aspect of the experimental design (e.g., the number of participants, the number of treatment groups, the timing of stimulus presentation), in this study we focus on optimizing only the particular set and order of the choice stimuli presented to the participant. As we are focusing on the binary choice paradigm for eliciting temporal discounting preferences, each question will consist of a choice between two amounts of money (x) at different time delays (t). Since all of the models under consideration predict that the discount rate is a decreasing function of t , the

²Each mini-experiment could be a single trial or a block several trials of trials.

choice will always be between a smaller amount of money at a sooner time and a larger amount of money at a later time, denoted by the quadruple (x_S, t_S, x_L, t_L) , where $x_S > x_L$ and $t_S < t_L$. Thus, before presentation of the first stimulus, and after each choice is made by the participant thereafter, ADO will search for an optimal quadruple (x_S, t_S, x_L, t_L) to be presented in each trial.

In principle, the design optimization problem on each trial entails a computationally intensive search of a 4-dimensional Euclidean space, which could not be completed in real-time without a significant delay between trials. Therefore, we restrict the design space to consist of only whole-numbered reward values and delays. This changes the continuous search problem to a discrete one, simplifying the computation and minimizing delay between trials. We further restrict the set of possible rewards and time delays so that $x_S \in \{8, 12, 15, 17, 19, 22\}$, $x_L \in \{12, 15, 17, 19, 22, 23\}$, $t_S \in \{0, 1, 2, 3, 5, 10, 20, 40\}$, and $t_L \in \{1, 2, 3, 5, 10, 20, 40, 80\}$.³

Even with these restrictions on the design space, there are still 756 possible stimuli that could be presented on each trial. ADO considers this same “master set” of possible stimuli for every participant, from which only a handful are actually be presented to obtain choice data. The range of possible values in this design space permits considerable fine-tuning of the stimuli for each participant, affording ADO a resolution not possible with standard, non-adaptive designs. For example, the 27 stimuli composing the Kirby Monetary Choice Questionnaire (Kirby et al., 1999), which is a standard instrument in the applied literature, each consist of an immediate reward (i.e., $t_S = 0$) ranging from \$11 to \$80, and larger reward in the same range at a some delay between 7 and 186 days. In what follows, we will write D to denote the set of possible designs, and $d \in D$ to denote a particular design.

2.3 Objective function

Identifying an optimal design at each stage of an ADO experiment entails solving an optimization problem defined as $d^* = \operatorname{argmax}_d U(d)$, for some real-valued objective function $U(d)$. In the current implementation, we deploy ADO trials with two different objective to match the two simultaneous goals of the experiment: parameter estimation and model discrimination. Both forms are motivated by information theory. For parameter estimation trials, $U(d)$ is set to measure the expected reduction in uncertainty (measured by Shannon entropy) about the values of the parameters that would be provided by the observation of an experimental outcome under design d (Kontsevich and Tyler, 1999; Kujala and Lukka, 2006). For model discrimination trials, $U(d)$ is set to measure the expected reduction in uncertainty about the functional form of the data-generating model that would be provided by the observation of an experimental outcome under design d (Cavagnaro et al., 2010). See Appendix for technical details.

³A geometric spacing of the independent variable (time) has been shown to be beneficial for discriminating between power and exponential decay curves in the study of memory retention (Cavagnaro et al., 2011). These sets of reward values were constructed to maximize the number of distinct ratios between the smaller and larger amounts.

2.4 Priors

ADO must be initialized with prior distributions over models (i.e., model probabilities) and over the parameters of each model. In general, there are many things to take into consideration when defining priors for a model, including theoretical constraints, prior knowledge, and computational convenience (e.g., using conjugate priors). With these in mind, we used equal model probabilities and uniform priors for each parameter in each model, with bounds determined by the range of values that have been reported in previous studies. A survey of the literature⁴ led to the following parameter ranges, which are represented graphically in Figure 2:

Exponential (Exp): $r \in (0.0005, 0.2)$

Hyperbolic (Hyp): $r \in (0.001, 0.1)$

Constant Sensitivity (CS): $r \in (0.0005, 0.1)$, $s \in (0.15, 1.5)$

Green & Myerson (GM): $r \in (0.0001, 1.0)$, $s \in (0.1, 2.0)$

Beta-Delta ($\beta - \delta$): $r \in (0.0, 1.0)$, $s \in (0.0, 1.0)$

Double Exponential (DE): $r \in (0.0, 0.8)$, $s \in (0.8, 1.0)$, $w \in (0.0, 1.0)$

In ADO, priors drive the selection of stimuli in the initial stages of the experiment, but since the parameter distributions are updated sequentially, the data will quickly trump all but the most pathological of prior distributions. Therefore, using informative priors is helpful for speeding up convergence, but not essential to implementing ADO (Myung et al., 2013). Most importantly, the priors should be defined in such a way that when data are generated from any one of the models under consideration, that model can be recovered correctly based on its posterior probability. This property can be verified through model recovery simulations, which are reported in the next section.

3 Experiment

3.1 Procedure

40 subjects were recruited (17 female, mean age = 27.55, SD = 12.69) from a paid participant pool maintained by the Stanford University Psychology Department. Written informed consent was obtained using a consent form and procedures approved by the Institutional Review Board of Stanford University. As described above, the task consisted of a sequence of binary choices between monetary rewards varying by reward magnitude (in dollars) and time of delivery (in days). Each of the 80 delay discounting trials consisted of a choice between a smaller-sooner (x_S, t_S) or larger-later (x_L, t_L) monetary reward simultaneously displayed on either side of the monitor ($x_S \in \{8, 12, 15, 17, 19, 22\}$, $x_L \in \{12, 15, 17, 19, 22, 23\}$, $t_S \in \{0, 1, 2, 3, 5, 10, 20, 40\}$, and $t_L \in \{1, 2, 3, 5, 10, 20, 40, 80\}$). Participants indicated their choice by pressing one of two buttons on a fiber-optic response pad using their right hand. Prior to the experiment, participants completed a set of practice

⁴The following studies were consulted to develop reasonable expectations about the parameter ranges: Ebert and Prelec (2007); Frederick et al. (2002); Green and Myerson (2004); Killeen (2009); Laibson (1997); Loewenstein and Prelec (1992); McClure et al. (2004, 2007); McKelvey et al. (2009); Simpson and Vuchinich (2011).

trials to familiarize themselves with the task. In order to ensure “incentive compatibility,” one trial was randomly chosen following the experiment, and the reward chosen by the participant in that trial was paid to the participant, at the specified delay time, in the form of a (post-dated) check. In addition, all participants received \$5.00 USD for participating in the experiment.

3.2 Analysis

We assume a logistic choice function in fitting and selecting among the models defined above. The choice function is defined as

$$p_i(SS|\theta_m) = \frac{1}{1 + e^{\varepsilon(V(LL) - V(SS))}} \quad (8)$$

where $SS = (x_S, t_S)$ is the smaller-sooner option, $LL = (x_L, t_L)$ is the larger-later option, $V(\cdot)$ are the subjective values determined by the algebraic discounting model, and $\varepsilon > 0$ is a free parameter.

We select among models for each individual subject using the Akaike Information Criterion, computed as $AIC = -2 * \ln L + 2k$, where $\ln L$ is the maximized log-likelihood and k is the number of free parameters in the model (Akaike, 1976). The AIC is interpreted as an information-theoretic distance between the true model and the fitted model (Myung, 2000), so we select the model with the lowest AIC as providing the best explanation of the data.

To select a single best model overall, we aggregate the within-subject results (i.e., the AIC values, not the choice data) using Akaike weights (e.g., Burnham and Anderson, 2004). The Akaike weight of a given model for a given subject, which is obtained as a straightforward transformation of the “raw” AIC value⁵, is the probability that the model is “best” for that subject in the sense of minimizing the information-theoretic distance (Wagenmakers and Farrell, 2004). Therefore, the product of AIC weights for a given model, across subjects, is the joint likelihood of that model being best for every subject. Since there are more than two models under consideration, we will use the Coin Flip model as a baseline against which to compare the joint likelihood of each other model. Therefore, we define the *group AIC factor* (gAF) of a given model m_i as

$$gAF_i = \prod_j \frac{w_{ij}(AIC)}{w_{0j}(AIC)}$$

where $w_i(AIC)$ and $w_0(AIC)$ are the Akaike weights for model m_i and the Coin Flip model, respectively. The model with the highest gAF is the one that is most likely to be the best for every subject.

⁵The Akaike weight for each model is a straightforward transformation of the “raw” AIC values, based on the fact that the raw AIC is an unbiased estimator of minus twice the expected log-likelihood of the model.

The above analysis assumes that every subject has the same model (i.e., functional form) and selects the one that is the most likely to be that shared model. It accounts for individual differences by allowing the parameter of the shared model to vary between subjects. However, this approach may be insufficient if none of the models under consideration is flexible enough to account for the full range of discounting behavior exhibited by subjects. Therefore, under the assumption that different subjects may have different models, we compute the expected probability of each model being best for a randomly selected subject. We do this using an extension of the variational Bayesian method proposed by Stephan et al. (2009). In short, with this method, the expected multinomial distribution over subject-level models is estimated as the expected value of a Dirichlet distribution that is conditioned on the observed data. The Dirichlet distribution has one parameter for each model, and the parameter α_i for each model m_i is approximated as $\alpha_i = 1 + \sigma_j w_j(AIC)^6$. Thus, assuming k candidate models, the expected probability of obtaining the i^{th} model for any randomly selected subject is $\frac{\alpha_i}{\alpha_1 + \dots + \alpha_k}$.

We also conducted a Bayesian version of each of the above analyses. The results are qualitatively identical, so we report them in the online appendix.

3.3 Results

Three subjects were dropped from the analysis due to ceiling or floor effects (e.g., choosing the larger later option on every trial), which made the models non-identifiable. Model selection results for the remaining 37 subjects are summarized in the first column of Table 1. Each model was best (i.e., lowest AIC among models under consideration) for at least one subject. Even the Coin Flip model was best for two subjects, indicating a failure of every substantive model under consideration for those two subjects. The Constant Sensitivity model was best for the largest number, 11 out of the 37 subjects, but the Double Exponential model was best for 9 subjects, and the Generalized Hyperbolic model was best for 7 subjects. This heterogeneity is also reflected in qualitative shapes of the fitted discounting curves for each subject, which are shown in Figure 3.

The second column of Table 1 reports the gAF for each model, on a log scale. In contrast to the raw AIC, for which lower numbers are better, a higher gAF means stronger evidence in favor of that model relative to the coin flip model. A gAF of zero indicates equivalent performance to that of the coin flip model. All of the gAF's in the table are large and positive, indicating that they perform overwhelmingly better than the Coin Flip model. The model with the largest gAF is Constant Sensitivity. Thus, besides providing the best explanation for the largest number of subjects, the Constant Sensitivity model is also the best model jointly for all subjects. Although the Exponential and Hyperbolic models are preferred over the Coin Flip model, they perform very poorly relative to more complex alternatives, especially Constant Sensitivity, Double Exponential, and Generalized Hyperbolic.

⁶Stephan et al. (2009) formulate this method in a Bayesian framework, using the posterior model probability in place of the Akaike weight.

A closer look at the parameter estimates for each model/subject reveals the role that increasing impatience played in the relative success of Constant Sensitivity. In our sample of subjects, the median estimate of the s parameter of the Constant Sensitivity model (i.e., the median of the within-subject MLEs) was 0.62, indicating decreasing impatience on average ($s < 1.0$). However, among the 11 subjects for whom Constant Sensitivity was best (according to the AIC), the median MLE of s was 1.26, indicating increasing impatience on average for that group ($s > 1.0$). These results indicate that much of the success of Constant Sensitivity model in our sample was driven by its ability (unique among the models under consideration) to accommodate increasing impatience. That being said, its success was not solely due to its ability to accommodate increasing impatience, as 4 out of the 11 subjects who were best described by Constant Sensitivity actually showed decreasing impatience ($s < 1$). Moreover, there were an additional 4 subjects who appeared to show increasing impatience based on the MLE of the Constant Sensitivity model ($s > 1.0$), but the Constant Sensitivity model did not have the lowest AIC, indicating that their data could be explained better without increasing impatience. This suggests that using the Constant Sensitivity model exclusively may overestimate the number of subjects with increasing impatience.

The above results suggest none of the models under consideration is adequate to explain every subject. Therefore, we use the results of the Dirichlet mixture model, which allows every subject to have a different model. The third column of Table 1 reports the expected probability of each subject-level model under the Dirichlet mixture model. The most probable model is Constant Sensitivity (0.289), followed closely by Double Exponential (0.244), generalized Hyperbolic (0.154), and Hyperbolic (0.109). None of the other models had a probability above 0.100. These probabilities can be interpreted as the expected proportion of subjects who are best described by each model, in the sense of minimizing the information-theoretic distance to the “true” model. They can also be interpreted as the probabilities of a randomly chosen subject being best described by each model.

Finally, to better illustrate the qualitative differences between the models that allowed them to succeed for different subjects, the right panel of Figure 4 shows the median estimate of each model among only the subgroup of subjects for whom that model was best, according to the AIC. For comparison, the left panel of Figure 4 shows the median estimate of each model. These results suggest that each model achieves success at the individual level by capturing a distinct pattern of discounting behavior, which other models are either unable to capture, or can only capture with a more complex functional form. For instance, Constant Sensitivity seems to work best for subjects who show an “extended present,” in which rewards with short delays are treated like immediate rewards, resulting in concavity of the discounting curve near $t=0$. On the other hand, Beta-Delta discounting works best when there is a distinct “present-future dichotomy”, in which every delay greater than zero is discounted by the same amount. Similarly, the Double Exponential works best when there is a short period of steep discounting followed by a longer period of relatively shallow discounting, perhaps reflecting a dual-process origin of the discounting. The Exponential, Hyperbolic, and generalized Hyperbolic models work best when the discounting behavior is somewhere in between these extremes.

4 Discussion and conclusions

The intuitive concept that delaying gratification is aversive has led to an extensive literature on temporal discounting and its applications. Investigation of temporal discounting preferences has generated numerous candidate models, with little agreement on which among them most accurately characterizes human decision making. In the current study, we attempted to identify the best fitting model from a set of popular candidate models using ADO, an approach to experimental design that tailors the task to the individual participant's responses.

Simulations, reported in the online appendix, conclusively demonstrate the capacity of ADO to recover a predetermined generating model in the presence of stochastic error. Model recovery with ADO was demonstrably superior to alternative design strategies, including the widely used Kirby Monetary Choice Questionnaire (Kirby et al., 1999).

In our experiment, ADO successfully converged to the best fitting model at the individual level. As described above, we found very significant individual differences in temporal discounting preferences that extended to the model space. One reason for this heterogeneity could be the aforementioned context-dependence of temporal discounting (Frederick et al., 2002; van den Bos and McClure, 2013). Though numerous studies have demonstrated that discounting exhibits trait-like stability (see Koffarnus et al., 2013, for a review), it has also been shown that individual discount rates may be affected by how a choice is framed, the type of reward involved, as well as by various internal and external state variables, such as the level of emotional arousal, cognitive load, or hunger (Giordano et al., 2002; Li, 2008; Van den Bergh et al., 2008; Wang and Dvorak, 2010; Wilson and Daly, 2004; Peters and Büchel, 2011). Though experimental settings are specifically designed to isolate the construct in question, it may not be possible to adequately control for the internal and external states of individual participants.

Complicating efforts to identify the single “best” model further, it has been suggested that temporal discounting is not a unitary construct, but rather a composite of subprocesses (Frederick et al., 2002). For example, Loewenstein et al. (2001) proposed three constituent “motives” that together generate discounting preferences: impulsivity (acting without forethought), compulsivity (making and sticking to plans), and inhibition (the capacity to override impulses). Subprocesses may exhibit different degrees of intra- and inter-individual variability across states and over time, further contributing to heterogeneity in measured discount rates. It should be noted that neuroimaging studies have consistently shown that multiple distinct neural regions contribute to temporal discounting and intertemporal choice (McClure et al., 2004, 2007; Kable and Glimcher, 2007; Rangel et al., 2008; Carter et al., 2010). Based on these findings, the Double Exponential model was developed specifically to model a mixture of present-oriented and more patient processes (McClure et al., 2007).

Another significant result of the present study was the prevalence of increasing impatience (concavity of the discounting curve) in our sample. This phenomenon challenges the prevailing practice in the literature of modeling temporal discounting as exclusively non-increasing, while providing strong confirmation of the results from a small number of recent

studies, notably by Attema et al. (2010); Abdellaoui et al. (2010) and Abdellaoui et al. (2013). Among the models we analyzed, only the Constant Sensitivity model can accommodate increasing impatience. Perhaps most significant was our finding that a significant number of participants exhibited the “extended present” phenomenon described by Ebert and Prelec (2007) (see Figure 3). In these cases, the discounting curve was essentially flat initially before dropping to zero after $t = \frac{1}{r}$ days, suggesting that this window of time was treated as a uniform “present.” It is surprising to us that this phenomenon has not received greater attention in the literature, as it is intuitive that the duration of the “present” may vary from individual to individual and correlate with neural activity in informative ways. Further research is needed to elucidate whether the extent of the “present” is stable over time, whether it is context-dependent, and how it varies between individuals and groups.

With respect to the motivating question of the present study, though no single model conclusively surpassed the others in describing temporal discounting behavior for all or most subjects, the Constant Sensitivity and Double Exponential models provided the best explanation for the largest number of participants across analyses. From a purely descriptive perspective, this result indicates that the added flexibility these models provide outweighs the “cost” of greater complexity. We believe the success of the Constant Sensitivity model demonstrates that increasing impatience and the extended present are likely to be relatively common behavioral variants, which reinforces the value of utilizing models that accommodate this behavior. The success of the neuroscience-inspired DE model lends further support to the view that multiple processes are involved in evaluating future outcomes, and reinforces the value of informing behavioral modeling with principles from neuroscience (Camerer et al., 2005). We anticipate that analysis of the unique characteristics of the Constant Sensitivity and Double Exponential models may yield important results in future studies. In addition, if increasing impatience, the extended present, and mixture are all important for describing discounting behavior, we propose that a mixture of Constant Sensitivity and Double Exponential would be a logical extension.

It should be noted that some of the most frequently utilized models in the literature were outperformed by our Coin Flip model in several analyses. The single-parameter exponential and hyperbolic models failed often, but even the Constant Sensitivity and Double Exponential models, which performed the best in our sample overall, provided poor explanations of the data for some subjects. We believe this underscores the importance of careful model selection. Widespread adoption of ADO or related approaches should help address this concern, particularly if a Coin Flip model is included among the models under consideration.

As stated previously, our list of candidate models was not intended to be exhaustive, and we cannot rule out the possibility that a model not included in our analysis could better account for discounting preferences than any in our set. Several promising models have been developed in recent years that merit inclusion in future model comparison studies (Bleichrodt et al., 2009; Benhabib et al., 2010; Scholten and Read, 2006, 2010). In addition, it should be noted that all of the models tested assume linear utility, an assumption which has some support at the aggregate level, but could potentially introduce distortions if there is

significant heterogeneity at the individual level (Abdellaoui et al., 2013). However, over the range of reward magnitudes involved in our experiment, any effect of nonlinear utility would likely be small.

Although ADO can significantly improve the efficiency of data collection, and thus the quality of inference, it is not without limitations that may constrain its full potential. Here, we discuss two such limitations. Firstly, not every design variable of an experimental task can be optimized in ADO. The design variables being optimized must be “quantitative” such that the likelihood function depends upon the specific design values. Secondly, ADO may not be robust to imprecise formulations of the models under study. As a concrete example, ADO makes the assumption that the set of models under consideration includes the one that actually generated the data (i.e., the “true” model). This technical assumption is almost certain to be violated in practice because our understanding of the topic being modeled is sufficiently incomplete to make any model only a first order approximation of the true model. Ideally, one would like to optimize a design for an infinite set of models representing all conceivable realities. To our knowledge, no implementable computational approach is currently available to solve a problem of this scope.

The ADO methodology, as currently formulated, is designed to optimize the selection of designs over a sequence of stages *within* a single experimental session with one participant. It would be desirable if ADO could somehow be extended to optimize the problem of design selection *across* multiple experimental sessions as well as within one session. To this end, Kim et al. (2014) have recently proposed a hierarchical Bayesian extension of ADO, which provides a judicious way to exploit information gained from previous sessions of the same experiment conducted with other participants.

In conclusion, although the set of models we considered in this study was not exhaustive, we believe that the current results contribute to the goal of achieving clarity on the issue of which parametric assumptions to make in fitting temporal discounting curves, in addition to demonstrating the significant heterogeneity of discounting behavior that can be expected in even a comparatively homogeneous group of healthy subjects. Finally, we feel that the results argue persuasively for the adoption of methods such as ADO in eliciting temporal discounting preferences.

References

- Abdellaoui M, Attema AE, Bleichrodt H. Intertemporal tradeoffs for gains and losses: An experimental measurement of discounted utility*. *The Economic Journal*. 2010; 120(545):845–866.
- Abdellaoui M, Bleichrodt H, IHaridon O. Sign-dependence in intertemporal choice. *Journal of Risk and Uncertainty*. 2013; 47(3):225–253.
- Ainslie G. Specious reward: a behavioral theory of impulsiveness and impulse control. *Psychological bulletin*. 1975; 82(4):463. [PubMed: 1099599]
- Ainslie G, Herrnstein RJ. Preference reversal and delayed reinforcement. *Animal Learning & Behavior*. 1981; 9(4):476–482.
- Akaike H. Canonical correlation analysis of time series and the use of an information criterion. *Mathematics in Science and Engineering*. 1976; 126:27–96.
- Attema AE, Bleichrodt H, Rohde KI, Wakker PP. Time-tradeoff sequences for analyzing discounting and time inconsistency. *Management Science*. 2010; 56(11):2015–2030.

- Benhabib J, Bisin A, Schotter A. Present-bias, quasi-hyperbolic discounting, and fixed costs. *Games and Economic Behavior*. 2010; 69(2):205–223.
- Berns GS, Laibson D, Loewenstein G. Intertemporal choice—toward an integrative framework. *Trends in cognitive sciences*. 2007; 11(11):482–488. [PubMed: 17980645]
- Bickel WK, Yi R, Landes RD, Hill PF, Baxter C. Remember the future: working memory training decreases delay discounting among stimulant addicts. *Biological psychiatry*. 2011; 69(3):260–265. [PubMed: 20965498]
- Bleichrodt H, Rohde KI, Wakker PP. Non-hyperbolic time inconsistency. *Games and Economic Behavior*. 2009; 66(1):27–38.
- Broomell SB, Bhatia S. Parameter recovery for decision modeling using choice data. *Decision*. 2014; 1(4):252.
- Burnham KP, Anderson DR. Multimodel inference understanding aic and bic in model selection. *Sociological methods & research*. 2004; 33(2):261–304.
- Camerer C, Loewenstein G, Prelec D. Neuroeconomics: How neuroscience can inform economics. *Journal of economic Literature*. 2005:9–64.
- Carter RM, Meyer JR, Huettel SA. Functional neuroimaging of intertemporal choice models. *Journal of Neuroscience, Psychology, and Economics*. 2010; 3(1):27–45.
- Cavagnaro DR, Gonzalez R, Myung JI, Pitt MA. Optimal decision stimuli for risky choice experiments: An adaptive approach. *Management science*. 2013a; 59(2):358–375. [PubMed: 24532856]
- Cavagnaro DR, Myung JI, Pitt MA, Kujala JV. Adaptive design optimization: A mutual information-based approach to model discrimination in cognitive science. *Neural computation*. 2010; 22(4): 887–905. [PubMed: 20028226]
- Cavagnaro DR, Pitt MA, Gonzalez R, Myung JI. Discriminating among probability weighting functions using adaptive design optimization. *Journal of risk and uncertainty*. 2013b; 47(3):255–289. [PubMed: 24453406]
- Cavagnaro DR, Pitt MA, Myung JI. Model discrimination through adaptive experimentation. *Psychonomic bulletin & review*. 2011; 18(1):204–210. [PubMed: 21327352]
- Chaloner K, Verdinelli I. Bayesian experimental design: A review. *Statistical Science*. 1995:273–304.
- Dallery J, Raiff BR. Delay discounting predicts cigarette smoking in a laboratory model of abstinence reinforcement. *Psychopharmacology*. 2007; 190(4):485–496. [PubMed: 17205320]
- Dasgupta P. Discounting climate change. *Journal of risk and uncertainty*. 2008; 37(2–3):141–169.
- Doyle JR. Survey of time preference, delay discounting models. *Judgment and Decision Making*. 2013; 8(2):116–135.
- Ebert JE, Prelec D. The fragility of time: Time-insensitivity and valuation of the near and far future. *Management Science*. 2007; 53(9):1423–1438.
- Frederick S, Loewenstein G, O'donoghue T. Time discounting and time preference: A critical review. *Journal of economic literature*. 2002:351–401.
- Gelman, A., Carlin, JB., Stern, HS., Dunson, DB., Vehtari, A., Rubin, DB. *Bayesian data analysis*. CRC press; 2013.
- Giordano LA, Bickel WK, Loewenstein G, Jacobs EA, Marsch L, Badger GJ. Mild opioid deprivation increases the degree that opioid-dependent outpatients discount delayed heroin and money. *Psychopharmacology*. 2002; 163(2):174–182. [PubMed: 12202964]
- Green L, Fisher E, Perlow S, Sherman L. Preference reversal and self control: Choice as a function of reward amount and delay. *Behaviour Analysis Letters*. 1981
- Green L, Myerson J. A discounting framework for choice with delayed and probabilistic rewards. *Psychological bulletin*. 2004; 130(5):769. [PubMed: 15367080]
- Johnson MW, Bickel WK. Within-subject comparison of real and hypothetical money rewards in delay discounting. *Journal of the experimental analysis of behavior*. 2002; 77(2):129–146. [PubMed: 11936247]
- Kable JW, Glimcher PW. The neural correlates of subjective value during intertemporal choice. *Nature neuroscience*. 2007; 10(12):1625–1633. [PubMed: 17982449]

- Kass RE, Raftery AE. Bayes factors. *Journal of the american statistical association*. 1995; 90(430): 773–795.
- Killeen PR. An additive-utility model of delay discounting. *Psychological review*. 2009; 116(3):602. [PubMed: 19618989]
- Kim W, Pitt MA, Lu Z-L, Steyvers M, Myung JI. A hierarchical adaptive approach to optimal experimental design. *Neural Computation*. 2014; (26):2463–2492.
- Kirby KN. Bidding on the future: evidence against normative discounting of delayed rewards. *Journal of Experimental Psychology: General*. 1997; 126(1):54.
- Kirby KN, Petry NM, Bickel WK. Heroin addicts have higher discount rates for delayed rewards than non-drug-using controls. *Journal of Experimental Psychology: General*. 1999; 128(1):78. [PubMed: 10100392]
- Koffarnus MN, Jarmolowicz DP, Mueller ET, Bickel WK. Changing delay discounting in the light of the competing neurobehavioral decision systems theory: a review. *Journal of the experimental analysis of behavior*. 2013; 99(1):32–57. [PubMed: 23344987]
- Kontsevich LL, Tyler CW. Bayesian adaptive estimation of psychometric slope and threshold. *Vision research*. 1999; 39(16):2729–2737. [PubMed: 10492833]
- Kujala JV, Lukka TJ. Bayesian adaptive estimation: The next dimension. *Journal of Mathematical Psychology*. 2006; 50(4):369–389.
- Laibson D. Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics*. 1997:443–477.
- Li X. The effects of appetitive stimuli on out-of-domain consumption impatience. *Journal of Consumer Research*. 2008; 34(5):649–656.
- Lindley DV. On a measure of the information provided by an experiment. *The Annals of Mathematical Statistics*. 1956:986–1005.
- Loewenstein G. Out of control: Visceral influences on behavior. *Organizational behavior and human decision processes*. 1996; 65(3):272–292.
- Loewenstein G, Prelec D. Anomalies in intertemporal choice: Evidence and an interpretation. *The Quarterly Journal of Economics*. 1992:573–597.
- George, Loewenstein, Weber, R., Flory, J., Manuck, S., Muldoon, M. Dimensions of time discounting. Paper Presented at Conference on Survey Research on Household Expectations and Preferences; Ann Arbor. Nov. 2–3; 2001.
- MacKillop J, Kahler CW. Delayed reward discounting predicts treatment response for heavy drinkers receiving smoking cessation treatment. *Drug and alcohol dependence*. 2009; 104(3):197–203. [PubMed: 19570621]
- Madden, GJ., Bickel, WK. Impulsivity: The behavioral and neurological science of discounting. American Psychological Association; 2010.
- Mazur JE. An adjusting procedure for studying delayed reinforcement. *Quantitative analyses of behavior*. 1987; 5:55–73.
- McClure SM, Bickel WK. A dual-systems perspective on addiction: contributions from neuroimaging and cognitive training. *Annals of the New York Academy of Sciences*. 2014; 1327(1):62–78. [PubMed: 25336389]
- McClure SM, Ericson KM, Laibson DI, Loewenstein G, Cohen JD. Time discounting for primary rewards. *The Journal of Neuroscience*. 2007; 27(21):5796–5804. [PubMed: 17522323]
- McClure SM, Laibson DI, Loewenstein G, Cohen JD. Separate neural systems value immediate and delayed monetary rewards. *Science*. 2004; 306(5695):503–507. [PubMed: 15486304]
- McKerchar TL, Green L, Myerson J, Pickford TS, Hill JC, Stout SC. A comparison of four models of delay discounting in humans. *Behavioural Processes*. 2009; 81(2):256–259. [PubMed: 19150645]
- Metcalf J, Mischel W. A hot/cool-system analysis of delay of gratification: dynamics of willpower. *Psychological review*. 1999; 106(1):3. [PubMed: 10197361]
- Moore SC, Cusens B. Delay discounting predicts increase in blood alcohol level in social drinkers. *Psychiatry research*. 2010; 179(3):324–327. [PubMed: 20494455]
- Myerson J, Green L. Discounting of delayed rewards: Models of individual choice. *Journal of the experimental analysis of behavior*. 1995; 64(3):263–276. [PubMed: 16812772]

- Myung IJ. The importance of complexity in model selection. *Journal of Mathematical Psychology*. 2000; 44(1):190–204. [PubMed: 10733864]
- Myung JI, Cavagnaro DR, Pitt MA. A tutorial on adaptive design optimization. *Journal of mathematical psychology*. 2013; 57(3):53–67. [PubMed: 23997275]
- Myung, JI., Cavagnaro, DR., Pitt, MA. Model evaluation and selection. In: Batchelder, WH. Colonius, H. Dzhafarov, E., Myung, JI., editors. *New Handbook of Mathematical Psychology*. Cambridge University Press; London: in press
- Myung JI, Pitt MA. Optimal experimental design for model discrimination. *Psychological review*. 2009; 116(3):499. [PubMed: 19618983]
- Peters J, Büchel C. The neural mechanisms of inter-temporal decision-making: understanding variability. *Trends in cognitive sciences*. 2011; 15(5):227–239. [PubMed: 21497544]
- Peters J, Miedl SF, Büchel C. Formal comparison of dual-parameter temporal discounting models in controls and pathological gamblers. *PloS one*. 2012; 7(11):e47225. [PubMed: 23226198]
- Phelps ES, Pollak RA. On second-best national saving and game-equilibrium growth. *The Review of Economic Studies*. 1968:185–199.
- Pine A, Seymour B, Roiser JP, Bossaerts P, Friston KJ, Curran HV, Dolan RJ. Encoding of marginal utility across time in the human brain. *The Journal of Neuroscience*. 2009; 29(30):9575–9581. [PubMed: 19641120]
- Pine A, Shiner T, Seymour B, Dolan RJ. Dopamine, time, and impulsivity in humans. *The Journal of Neuroscience*. 2010; 30(26):8888–8896. [PubMed: 20592211]
- Rachlin H. Notes on discounting. *Journal of the experimental analysis of behavior*. 2006; 85(3):425–435. [PubMed: 16776060]
- Rangel A, Camerer C, Montague PR. A framework for studying the neurobiology of value-based decision making. *Nature Reviews Neuroscience*. 2008; 9(7):545–556. [PubMed: 18545266]
- Ray, D., Golovin, D., Krause, A., Camerer, C. California Institute of Technology working paper. 2012. Bayesian rapid optimal adaptive design (broad): Method and application distinguishing models of risky choice.
- Reynolds B. A review of delay-discounting research with humans: relations to drug use and gambling. *Behavioural pharmacology*. 2006; 17(8):651–667. [PubMed: 17110792]
- Samuelson PA. A note on measurement of utility. *The Review of Economic Studies*. 1937; 4(2):155–161.
- Scholten M, Read D. Beyond discounting: the tradeoff model of intertemporal choice. *Operations research working papers*. 2006
- Scholten M, Read D. The psychology of intertemporal tradeoffs. *Psychological Review*. 2010; 117(3): 925. [PubMed: 20658858]
- Sharp C, Barr G, Ross D, Bhimani R, Ha C, Vuchinich R. Social discounting and externalizing behavior problems in boys. *Journal of Behavioral Decision Making*. 2012; 25(3):239–247.
- Shefrin HM, Thaler RH. The behavioral life-cycle hypothesis. *Economic inquiry*. 1988; 26(4):609–643.
- Simpson CA, Vuchinich RE. Reliability of a measure of temporal discounting. *The Psychological Record*. 2011; 50(1):1.
- Stephan KE, Penny WD, Daunizeau J, Moran RJ, Friston KJ. Bayesian model selection for group studies. *Neuroimage*. 2009; 46(4):1004–1017. [PubMed: 19306932]
- Story GW, Vlaev I, Seymour B, Darzi A, Dolan RJ. Does temporal discounting explain unhealthy behavior? a systematic review and reinforcement learning perspective. *Frontiers in behavioral neuroscience*. 2014:8. [PubMed: 24478655]
- Takahashi T, Oono H, Radford MH. Psychophysics of time perception and intertemporal choice models. *Physica A: Statistical Mechanics and its Applications*. 2008; 387(8):2066–2074.
- Thaler R. Some empirical evidence on dynamic inconsistency. *Economics Letters*. 1981; 8(3):201–207.
- Toubia O, Johnson E, Evgeniou T, Delquié P. Dynamic experiments for estimating preferences: An adaptive method of eliciting time and risk parameters. *Management Science*. 2013; 59(3):613–640.

- Van den Bergh B, Dewitte S, Warlop L. Bikinis instigate generalized impatience in intertemporal choice. *Journal of Consumer Research*. 2008; 35(1):85–97.
- van den Bos W, McClure SM. Towards a general model of temporal discounting. *Journal of the experimental analysis of behavior*. 2013; 99(1):58–73. [PubMed: 23344988]
- Wagenmakers EJ, Farrell S. Aic model selection using akaike weights. *Psychonomic bulletin & review*. 2004; 11(1):192–196. [PubMed: 15117008]
- Wang, SW., Filiba, M., Camerer, CF. California Institute of Technology working paper. 2010. Dynamically optimized sequential experimentation (dose) for estimating economic preference parameters.
- Wang X, Dvorak RD. Sweet future fluctuating blood glucose levels affect future discounting. *Psychological Science*. 2010; 21(2):183–188. [PubMed: 20424042]
- Weiss DJ, Kingsbury G. Application of computerized adaptive testing to educational problems. *Journal of Educational Measurement*. 1984; 21(4):361–375.
- Wilson M, Daly M. Do pretty women inspire men to discount the future? *Proceedings of the Royal Society of London. Series B: Biological Sciences*. 2004; 271(Suppl 4):S177–S179. [PubMed: 15252976]
- Zauberman G, Kim BK, Malkoc SA, Bettman JR. Discounting time and time discounting: Subjective time perception and intertemporal preferences. *Journal of Marketing Research*. 2009; 46(4):543–556.

Appendix

A Additional details on the implementation of ADO

A.1 Objective Function

Design optimization at each stage of an ADO experiment entails solving an optimization problem defined as $d^* = \operatorname{argmax}_d U(d)$, for some real-valued objective function $U(d)$. Suppose there are K models under consideration, denoted by $m = \{1, 2, \dots, K\}$, let θ_m denote a parameter vector of model m , and let $p(m)$ and $p(\theta_m)$ denote the prior model and parameter probabilities, respectively. Also, let y_m denote the outcome vector resulting from a hypothetical experiment conducted with design d under model m . Then, the general form of the objective function to be maximized is

$$U(d) = \sum_m p(m) \int \int u(d, \theta_m, y_m) p(y_m | \theta_m, d) p(\theta_m) dy_m d\theta_m$$

In the above equation, $u(d, \theta_m, y_m)$ is the so-called "local utility function."⁷ In the current implementation, we deploy ADO trials with two different local utility functions to match the two simultaneous goals of the experiment: parameter estimation and model discrimination. For parameter estimation trials, we will set

$$u(d, \theta_m, y) = \log \frac{p(\theta_m | y, d)}{p(\theta_m)}, \quad (9)$$

⁷The local utility function measures the "utility" of the design d when the true model is m with parameter θ_m and the outcome y_m is observed. Thus, $U(d)$ can be viewed as the "expected utility" of the design d , where the expectation is taken over possible models, parameters and experiment outcomes. Myung et al. (see 2013, for details).

which makes $U(d)$ equivalent to the expected reduction in uncertainty (measured by Shannon entropy) about the values of the parameters that would be provided by the observation of an experimental outcome under design d (Kontsevich and Tyler, 1999; Kujala and Lukka, 2006). In other words, the optimal design is the one that is expected to extract the maximum information about the model's parameters. For model discrimination trials, we will set

$$u(d, \theta_m, y) = \log \frac{p(m|y, d)}{p(m)}, \quad (10)$$

which gives $U(d)$ a similar, information theoretic interpretation as the expected amount of information about the data-generating model that would be provided by the observation of an experimental outcome under design d (Cavagnaro et al., 2010).

Formally, following the Bayesian decision theoretic framework of Chaloner and Verdinelli (1995), the general form of the objective function to be maximized is

$$U(d) = \sum_m p(m) \int \int u(d, \theta_m, y_m) p(y_m | \theta_m, d) p(\theta_m) dy_m d\theta_m$$

where $m = \{1, 2, \dots, K\}$ is one of a set of K models being considered, and y_m is the outcome vector resulting from a hypothetical experiment conducted with design d under model m , θ_m is a parameter vector of model m . In the above equation, $p(m)$ and $p(\theta_m)$ are the prior model and parameter probabilities, respectively.

The motivation for this objective function can be summarized as follows. If the so-called “local utility function,” $u(d, \theta_m, y_m)$, measures “utility” of the design d when the true model is m with parameter θ_m and the outcome y_m is observed, then $U(d)$ can be viewed as the “expected utility” of the design d , where the expectation is taken over possible models, parameters and experiment outcomes.

The precise form of $u(d, \theta_m, y)$ should be motivated by the goal of the experiment. In the current implementation we will deploy ADO trials with two different local utility functions to match the two simultaneous goals of the experiment: parameter estimation and model discrimination. For parameter estimation trials, we will set

$$u(d, \theta_m, y) = \log \frac{p(\theta_m | y, d)}{p(\theta_m)}, \quad (11)$$

which makes $U(d)$ equivalent to the expected reduction in uncertainty (measured by Shannon entropy) about the values of the parameters that would be provided by the observation of an experimental outcome under design d (Kontsevich and Tyler, 1999; Kujala

and Lukka, 2006). In other words, the optimal design is the one that is expected to extract the maximum information about the model's parameters (Myung et al., 2013). For model discrimination trials, we will set

$$u(d, \theta_m, y) = \log \frac{p(m|y, d)}{p(m)}, \quad (12)$$

which gives $U(d)$ a similar, information theoretic interpretation as the expected amount of information about the data-generating model that would be provided by the observation of an experimental outcome under design d (Cavagnaro et al., 2010).

A.2 Bayesian updating

On stage s of an ADO experiment, the design d_s^* to be implemented in the next stage is chosen by maximizing $U(d)$. Upon the observation of a specific experimental outcome z_s in that stage, the prior distributions to be used to find an optimal design for the next stage are updated via Bayes' rule and Bayes factor calculation (e.g., Gelman et al., 2013) according to the following equations

$$p_{s+1}(m) = \frac{p_1(m)}{\sum_{k=1}^K p_1(k) BF_{(k,m)}(z_s | d_s^*)} \quad (13)$$

$$p_{s+1}(\theta_m) = \frac{p(z_s | \theta_m, d_s^*) p_s(\theta_m)}{\int p(z_s | \theta_m, d_s^*) p_s(\theta_m) d\theta_m}. \quad (14)$$

In the equation $BF_{(k,m)}(z_s | d_s^*)$ is the Bayes factor that is defined as the ratio of the marginal likelihood of model k to that of model m given the outcome z_s and optimal design d_s^* Kass and Raftery, 1995). To recap, the ADO process involves, in each stage of experimentation, finding the optimal design d_s^* by maximizing the utility function $U(d)$, conducting a mini-experiment with the optimized design, observing an outcome z_s , and updating the model and parameter priors to the corresponding posteriors through Bayes' rule, as illustrated in Figure 1. This process continues until one model emerge

A.3 Stochastic specification

The objective function defined above requires each model under consideration to have a probabilistic likelihood function $p(y | \theta_m, d)$. Since the discounting models we have defined are algebraic in nature, we must equip them with a choice function, which translates utilities into choice probabilities. Following Cavagnaro et al. (2013a) and Cavagnaro et al. (2013b), we employ a parameter-free weak utility model that makes minimal assumptions about the structure of stochastic errors. The model assumes only that there is an unknown probability, between 0 and 0.5, of a "choice error" on any given trial, and that these choice errors are

independent from trial to trial. See Appendix for technical details. Formally, if $d_i = \{(x_S, t_S), (x_L, t_L)\}$ is the i^{th} stimulus presented in an experiment, and $\prec_{\theta_m}$ is the weak ordering over stimuli determined by the model parameter θ_m , then the probability of choosing (x_S, t_S) is given by

$$p_i((x_S, t_S)|\theta_m) = \begin{cases} \varepsilon_i & \text{if } (x_S, t_S) \prec_{\theta_m} (x_L, t_L) \\ \frac{1}{2} & \text{if } (x_S, t_S) \sim_{\theta_m} (x_L, t_L) \\ 1 - \varepsilon_i & \text{if } (x_S, t_S) \succ_{\theta_m} (x_L, t_L) \end{cases}$$

where $\varepsilon_i \in [0, 0.5]$.

The biggest advantage of this choice function for use in ADO is that there are no parameters to estimate, since ε_i is independent of the choices made on trials 1 through $i - 1$. This greatly simplifies the calculations that must be done between trials, which is essential when using the method in an experiment with human participants. Moreover, this model can be viewed as a generalization of many, commonly used choice functions. For instance, it reduces to a constant error model when ε_i is constant across trials, while a random utility model is a special case in which ε_i is a parametric function of d_i . In the next section, we will demonstrate that an ADO experiment assuming this general choice function can correctly recover the discounting model and parameters in a simulation experiment, even when the data are generated with a more specific logistic choice function.

B Model recovery simulations

We conducted model-recovery simulations to assess how quickly and conclusively a known data-generating model could be identified correctly using ADO, as compared to three different non-adaptive benchmarks. Each simulated experiment consisted of a 81 trials in which data (i.e., choices between delayed rewards) were generated according to a prespecified “generating” model (e.g., Hyperbolic with $r=0.02$). In each simulated experiment, we assessed whether the generating model and parameters could be identified correctly and conclusively from among the models defined above, on the basis of the generated data.

The stimuli on which choices were generated depended on the design strategy employed in each simulation. ADO is one such design strategy, and in the ADO simulations, stimuli were selected adaptively according to the ADO algorithm described above.⁸ We compared this design strategy to three non-adaptive benchmarks. In one of these benchmarks, stimuli were drawn uniformly at random, with replacement, from the design space D that is used in the implementation of ADO. The purpose of this benchmark is to assess the extent to which the efficiency of ADO can be attributed to the geometric spacing of the delays in D . Geometric

⁸ADO was implemented with two different objective functions. The first half of the trials using the objective function for parameter estimation (Equation 11), and the second half using the objective function for model discrimination (Equation 12). After each trial, posterior distributions and model probabilities were estimated based on thinned samples of size 5000, 20,000, and 80,000 from the joint posterior distribution of the one-parameter, two-parameter, and 3-parameter models, respectively, which were obtained via the independent Metropolis Hastings algorithm implemented in C++.

spacing of the independent variable is known to be beneficial for discriminating among decay models in general (e.g., Cavagnaro et al., 2011), so selecting stimuli at random from D may be sufficient to discriminate among models of temporal discounting. Therefore, we will refer to this design strategy as “Geometric.”

A second benchmark that we evaluated was to draw smaller-sooner larger-later stimuli at random, with integer values of x_S , t_S , x_L and t_L drawn uniformly from the following intervals: $x_S \in [1, 20]$, $t_S \in [0, 40]$, $x_L \in [x_S + 1, 30]$, $t_L \in [t_S + 1, 80]$. This benchmark will be referred to simply as “Random.”

The final benchmark that we evaluated utilizes a fixed design from the literature. In particular, in the so-called “Kirby” simulations, stimuli were selected sequentially from the Kirby Monetary Choice Questionnaire (Kirby et al., 1999), which is a fixed set of 21 questions, each consisting of an immediate reward (i.e., no delay) ranging from \$11 to \$80 and a delayed reward ranging from \$25 and \$85. The delays range from 7 to 186 days. To extend the simulation to 81 trials, each question in the questionnaire was repeated three times.

Whichever strategy was used to select stimuli, the generated data were analyzed as though the generating model were unknown. That is, analyses were initiated with the uninformative priors defined above, which were then updated upon observation of each data point according to Bayes rule as given in Equation A.2.⁹ The ADO simulations utilized these updated estimates to select stimuli at each stage. The other design strategy did not use them but they were computed anyway for the record. At the conclusion of each simulated experiment, model recovery was assessed based on the posterior probability of the generating model. That is, the higher the posterior probability of the generating model, the more conclusively it was identified by the data, and therefore the more effective the design strategy was.

We conducted simulations with various generating models and parameters, and the results are generally similar. Therefore, for the purposes of exposition, we will present the results of one particular model: CS with $r = 0.025$ and $s = 0.4$. To simulate human performance, the data will be generated according to a probabilistic choice function. In particular, we will use a logistic (softmax) choice function defined by

$$p_i(SS|\theta_m) = \frac{1}{1 + e^{\varepsilon(V(LL) - V(SS))}} \quad (15)$$

where $SS = (x_S, t_S)$ is the smaller-sooner option, $LL = (x_L, t_L)$ is the larger-later option, $V(\cdot)$ are the subjective values determined by the algebraic discounting model, and $\varepsilon > 0$ is a

⁹All simulations were coded to update model probabilities and parameter distributions in real-time, just as ADO does. However, since the calculations must be done in real-time, Monte Carlo sample sizes must be small, which can lead to biased estimates. Therefore, at the conclusion of each simulation, model probabilities were recalculated based on the entire sequence of data, using Monte Carlo integration with sample sizes of 1,000, 100,000, and 1,000,000 for the one, two, and three parameter models, respectively. In addition, to provide richer estimates of the posterior parameter estimates new samples of size 50,000, 250,000, and 1,000,000 were drawn from the joint posterior distributions for the one, two, and three parameter models, respectively.

sensitivity parameter. When ϵ equals zero, $p(SS) = 0.5$ regardless of the subjective values of the delayed rewards, so choice behavior is random. As ϵ increases, choice behavior is determined more and more by the difference in subjective value for the two choice options. The actual “error rate” across a set of choices for a given value of ϵ depends on the particular choice stimuli. If the utility differences are large then even very small values of ϵ will yield essentially errorless data, but if the utility differences are small then even large values of ϵ will yield quite noisy data. This means that ADO stimuli will tend to induce more errors for a given value of ϵ , since ADO stimuli are constructed to put maximal pressure on the models under consideration. In the simulations reported below, we set $\epsilon = 2.0$, which yielded errors on approximately 30% of trials in the ADO simulation, 10% of trials in each of the Geometric and Random simulations, no errors in the Kirby simulation.

A summary of the results of each simulation is shown in Table 2. Despite the fact that the ADO stimuli induced more choice errors in the simulated data, the posterior probability of the data generating model (CS) at the conclusion of the ADO simulation was 0.999. In contrast, none of the benchmark simulations were able to correctly recover CS as the data generating model. The Kirby simulation came the closest, with a final posterior probability of 0.351, although in that simulation the Exponential model actually came in ahead with a posterior probability of 0.366. In the Random simulation, the Hyperbolic model had the highest posterior probability. Only in the ADO and Geometric simulations did the true generating model (CS) end up with the highest posterior probability, and only in the ADO simulation was that probability conclusive.

To assess parameter recovery, we computed the root-mean-squared-error (RMSE) of the estimates of r and s based on the posterior parameter distributions. For example, the RMSE for r , which had a true value of 0.025 in the simulations, was computed as

$$RMSE(r) = \sqrt{\frac{\sum_{i=1}^{10000} (r_i - 0.025)^2}{10000}},$$

where $\{r_i\}_{i=1, \dots, 10000}$ are 10,000 sampled values of r from its posterior distribution. A smaller RMSE indicates a better estimate of the parameter. The RMSE for r and s at the conclusion of each simulation are given second and third columns of Table 2, respectively. Based on these numbers, it seems that the key to ADO’s success in recovering the generating model was its ability to estimate the s parameter precisely. All four simulations estimated r fairly precisely, but the ADO simulation achieved by far the most precise estimate of s .

After completing the above analyses, we recalculated the posterior probabilities and parameter distributions for the same simulation data assuming the logistic choice function in Equation 15 for each model. This means that each model had an additional free parameter ϵ to be estimated from the data. The prior on ϵ was set to be uniform between 0.0 and 3.0. The resulting model probabilities and RMSEs are shown in Table 3, which has an extra column for the RMSE of the ϵ parameter. The RMSE for ϵ was similar across the four design conditions, meaning that the ADO simulation did not show any advantage in estimating ϵ . This is not unexpected since the stimuli were not specifically intended to estimate the free parameter in the choice function. Nevertheless, the posterior probability of the generating model is still much higher in the ADO condition (0.988), than in any of the benchmark conditions (0.373, 0.563, and 0.441 for Kirby, Random, and Geometric, respectively). The

precision of the parameter estimates of r and s in each design condition are comparable or slightly better in the reanalysis than what they were in the original analysis.

The parameter estimates obtained in each design condition can also be compared graphically by plotting the posterior distribution of each parameter given the simulated data. The graph on the top-left of Figure 5 overlays four posterior distributions of r , one for each design condition, while the graph on the top-right overlays the distributions of s and the graph on the bottom-left overlays the distributions of ε . Each distribution was estimated using MATLAB's kernel smoothing function (`ksdensity`) based on 10,000 draws from the joint posterior distribution of r , s , and ε , given the simulated data in the corresponding condition. The vertical line in each graph represents the true (generating) values of r , s , and ε in the simulations.

The graph in the bottom-right of Figure 5 displays the best estimate of the discounting curve based on the maximum a-posteriori probability (MAP) estimates of the parameters r and s in each condition. The MAP is obtained as the mode of the posterior distribution, and can be viewed as the Bayesian equivalent of the Frequentist maximum likelihood (ML). The “true” discounting curve is plotted in red, but it may be difficult to see it because the estimate obtained in the ADO simulation (solid black line) sits right on top of it. That is to say, the discounting curve was recovered almost perfectly in the ADO simulation. The estimates obtained in the Geometric and Random simulations are also both quite close to the true curve, despite the fact that the parameter estimates are not very precise. In theory, no design should consistently yield biased parameter estimates unless the priors are biased and the designs cannot extract enough information to overcome those biases. These results demonstrate that obtaining unbiased parameter estimates is not necessarily sufficient to recover the functional form of the generating model.

The current results are based on just one set of simulations, with one particular generating model. Results will vary across replications depending on the generating model and parameters, as well the frequency and timing of errors in the generated data. Model and parameter recovery is particularly variable with the Geometric and Random designs, since their stimuli are generated at random as well. Sometimes they chance upon stimuli that happen to be diagnostic for recovering whatever the generating model happens to be, but since the design space is large and the “sweet spot” of stimuli that are diagnostic for any given model is small, such instances are rare and very often the model recovery results will be much worse than they were in this instance. The stimuli in the Kirby design are fixed, and the effectiveness of this design depends the generating model, and which stimuli are diagnostic for that generating model. It will be effective only if the generating model happens to be one for which the Kirby stimuli are diagnostic. On the other hand, ADO actively seeks out the sweet spot in the design space by learning the generating model as data points are observed. In that way, ADO is able to achieve a level of model discrimination not possible with the benchmark designs, while also being sensitive to individual differences.

A more extensive analysis could systematically investigate the relative effectiveness of each design strategy across a wide range generating models and parameters, and across different

sets of candidate models. MATLAB code for running additional simulations is available upon request and we will leave such an ambitious project for future work. From the current simulation results, we conclude that ADO is capable of recovering the generating model quite conclusively, even when the data are generated with stochastic error. Therefore, we move forward with confidence to an experiment using ADO with human subjects, in which the “true” data-generating model is unknown.

C Bayes factor results

Table 4 reports our analyses of the experiment data based on the within-subject normalized posterior likelihood (i.e., probability) of each of the seven models given the observed data. As shown in the first column of the table, the model that was most probable for the largest number was Constant Sensitivity (8 out of 37 subjects). The Double Exponential, Hyperbolic, and Coin Flip models were best for 7 subjects each. In sum, it seems that there were dramatic individual differences in terms of the best-fitting model.

The second column of Table 4 reports the group Bayes factor (gBF) of each model. The gBF of a model m_j is computed as the ratio of the joint likelihood of that model versus the coin flip model:

$$gBF(m_i) = \prod_{j=1}^{37} p_j(m_i) / p_j(m_0).$$

Where $p_j(m_i)$ and $p_j(m_0)$ are the posterior probabilities of models m_i and the coin flip model, respectively, for subject j . In the table, $\log(\text{gBF})$ greater than zero indicates evidence in favor of a model over the Coin Flip model, while a negative $\log(\text{gBF})$ indicates evidence against the model. Interestingly, only the Constant Sensitivity and Double Exponential models have a positive $\log(\text{gBF})$. The others have a negative $\log(\text{gBF})$, indicating that the Coin Flip model provides a more likely explanation than any of them. Between Constant Sensitivity and Double Exponential, the Double Exponential has the higher $\log(\text{gBF})$, indicating that it is the preferred model of the group. More precisely, the Double Exponential is approximately $10^{15.575-11.539} \approx 10,000$ times more likely than the CS model to have generated the group data.

The third column of Table 4 reports the expected probability of each model based on the Dirichlet mixture model. By this metric, the Double Exponential model is actually the most likely model for a randomly selected subject, followed closely by Constant Sensitivity. These were the only two models with an expected probability greater than 0.2.

Fig. 1. Schematic illustration of the cyclical relationship between Design Optimization, data collection, and Bayesian updating in each mini-experiment. Reprinted from Myung et al. (ress, Figure 10).

Fig. 2.

Depiction of the prior predictive distribution of discounting curves in each family. In each graph, the x-axis measures time and the y-axis measures the discount factor. Curves were generated by randomly selecting 100 parameter vectors from the priors for each model.

Fig. 3.
MLE of the model with the lowest AIC for each participant

Fig. 4.

Comparison of median estimates of each model in the entire sample versus the subgroups in which each model was best.

Fig. 5.

The four overlaid distributions in each of panels r , s and ε are the marginal posterior distributions obtained in each model recovery simulation. The vertical line in each graph shows the true value of the parameter, so a tighter distribution around that line indicates a better estimate. In the bottom-right panel, the five overlaid discounting curves are the generating curve (red) and the best estimate of that curve based on each of the four simulations. Estimates are based on the CS model using the maximum a-posteriori probability (MAP) estimates of r and s . It may be difficult to differentiate all of the curves visually because the estimate obtained in the ADO simulation (solid black line) sits right on top of the generating curve (red line). That is to say, the discounting curve was recovered almost perfectly in the ADO simulation. The estimates obtained in the Geometric and Random simulations are also both quite close to the true curve, despite the fact that the parameter estimates are not very precise.

Table 1

Model selection results based on AIC.

Model	# Best (Lowest AIC)	log(gAIC)	Expected Probability
Exponential	3	50.433	0.088
Hyperbolic	4	47.548	0.109
Constant Sensitivity	11	106.290	0.289
Generalized Hyperbolic	7	83.002	0.154
Beta-Delta	1	41.403	0.068
Double Exponential	9	90.245	0.244
Coin Flip	2	0.000	0.049

Table 2

Results of the model recovery simulation.

Design	p(CS)	RMSE(r)	RMSE(s)
Kirby	0.351	0.0133	0.4251
Random	0.123	0.0246	0.1720
Geometric	0.319	0.0141	0.2947
ADO	0.999	0.0103	0.0675

Table 3

Reanalysis of the model recovery simulation assuming a logistic choice function.

Design	p(CS)	RMSE(r)	RMSE(s)	RMSE(<i>e</i>)
Kirby	0.373	0.0114	0.4435	0.6866
Random	0.563	0.0201	0.0992	0.5159
Geometric	0.441	0.0119	0.1364	0.5527
ADO	0.988	0.0061	0.0388	0.5220

Table 4

Model selection results based on posterior probability.

Model	# Best (Highest Prob.)	log(gBF)	Expected Probability
Exponential	2	−9.584	0.098
Hyperbolic	7	−6.990	0.146
Constant Sensitivity	8	11.535	0.187
Generalized Hyperbolic	4	−4.248	0.125
Beta-Delta	2	−21.041	0.077
Double Exponential	7	15.579	0.200
Coin Flip	7	0.000	0.166