

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

Daniel R. Cavagnaro^{*1}, Gabriel J. Aranovich^{2,3}, Samuel M. McClure⁴, Mark A. Pitt⁵, and
Jay I. Myung⁵

¹California State University, Fullerton

²University of California, San Francisco

³San Francisco VA Medical Center

⁴Stanford University

⁵The Ohio State University

December 17, 2014

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 first test the method in simulation and compare its performance to several non-adaptive benchmarks. We then conduct an ADO experiment with human subjects, 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

JEL classification: C91, C52, D90

*Correspondence: dcavagnaro@fullerton.edu

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.

However, despite decades of investigation, there is still little consensus on the proper parametric form of the discounting curve (Frederick et al., 2002; van den Bos and McClure, 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 discount-

ing 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).

One 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 (i.e., parametric forms) and estimating 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 estimating and model comparison studies in the discounting literature have often employed fixed-task 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 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.¹ 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; McKerchar 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) have used AIC for model comparison, which does account for model complexity.

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

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).

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 (Cavagnaro et al., 2013b). The present work extends the approach to temporal discounting.

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), 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. We test the method against several benchmark elicitation methods in simulation experiments 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.

Readers who are only interested in the empirical contributions of the paper may wish to skip to section 4, which describes our implementation of ADO in an experiment with human subjects. 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.

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, for reviews). In this section, we describe the particular set of models chosen for our experiment.

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 (Δt). This can be summarized as

$$V = AD_{\Delta t} \tag{1}$$

where $D_{\Delta t}$ is the individual-specific discount factor, with value between 0 and 1, generated by the model, as a function of Δt .

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

$$D_{\Delta t} = e^{-r\Delta t} \tag{2}$$

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:

$$D_{\Delta t} = \frac{1}{1 + k\Delta t} \tag{3}$$

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):

$$D_{\Delta t} = \frac{1}{(1 + k\Delta t)^s} \quad (4)$$

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 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:

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

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 $\Delta t = 0$, 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:

$$D_{\Delta t} = \omega e^{-r\Delta t} + (1 - \omega)e^{-s\Delta t} \quad (6)$$

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):

$$D_{\Delta t} = e^{-(r\Delta t)^s} \quad (7)$$

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 (concavity; values of $s > 1$). In contrast, standard approaches to modeling temporal discounting assume monotonically decreasing impatience. Recent evidence has highlighted the prevalence of increasing impatience (Bleichrodt et al., 2009; Abdellaoui et al., 2010; Attema et al., 2010; Abdellaoui et al., 2013). At the extreme, the 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 Δt until $\Delta t = \frac{1}{r}$, after which discounting decreases as a function of Δt . 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: the Exponential within the Constant-Sensitivity, the Hyperbolic within the generalized Hyperbolic, and both the Exponential and Beta-Delta within the Double-Exponential. Putting these models in competition with one another, using ADO to maximally discriminate them, allowed us to compare families of models while also testing the extent to which additional parameters are justified within each family.

As stated previously, the set of candidate models chosen for this experiment 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 accommodated qualitatively different patterns of behavior and involved varying levels of complexity.

2 Methods: How ADO works

2.1 History and General framework

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

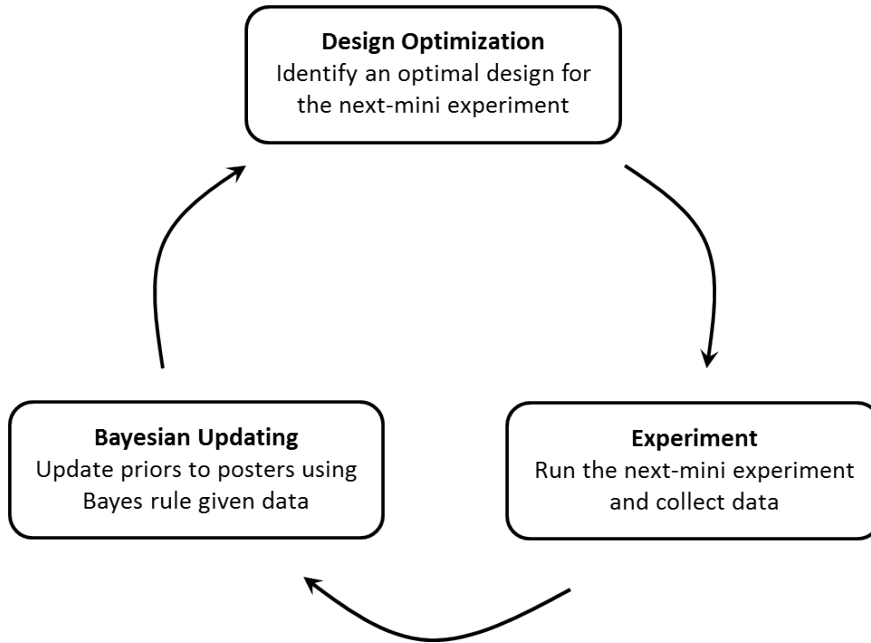


Figure 1: Schematic illustration of the cyclical relationship between Design Optimization, data collection, and Bayesian updating in each mini-experiment (from Myung et al., *ress*).

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. As an added benefit, in contrast to classical approaches that involve post-hoc curve-fitting procedures, ADO essentially performs the data analysis in real time.

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

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

their implementation, they share the common goal of optimizing experimental design in real-time in order to improve the quality of statistical inference.

The implementation of Adaptive Design Optimization requires several key decisions to be made. Most importantly, one must identify a design variable to be optimized in the experiment, such as the stimulus to be presented, and the space of all possible designs must be formulated. One must also choose an objective function on which to evaluate possible designs, the stochastic specification of the model or models under investigation, and the priors with which to initiate the models. In the rest of this section, we will describe the specific implementation of ADO that we will be using to evaluate and discriminate among models of temporal discounting. For a more general, comprehensive tutorial on ADO, see Myung et al. (2013).

2.1.1 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 (Δt). Since all of the models under consideration predict that the discount rate is a decreasing function of t , the 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 $(x_S, \Delta t_S, x_L, \Delta t_L)$, where $x_S > x_L$ and $\Delta t_S < \Delta t_L$. Thus, before presentation of the first stimulus, and after each choice is made by the participant thereafter, ADO will search for an optimal quadruple $(x_S, \Delta t_S, x_L, \Delta t_L)$ 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. However, 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 $x_S \in \{8, 12, 15, 17, 19, 22\}$, $x_L \in \{12, 15, 17, 19, 22, 23\}$, $\Delta t_S \in \{0, 1, 2, 3, 5, 10, 20, 40\}$, and $\Delta t_L \in \{1, 2, 3, 5, 10, 20, 40, 80\}$.³

Even with these restrictions on the design space, there are still 756 possible stimuli that could be presented on each trial. Although only a small subset of these stimuli can be presented over the course of a single experiment, the range of possible values 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

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.

(Kirby et al., 1999), which is a standard instrument in the applied literature, each consist of an immediate reward (i.e., $\Delta t_s = 0$) 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.

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

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

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

$$u(d, \theta_m, y) = \log \frac{p(\theta_m | y, d)}{p(\theta_m)}, \quad (8)$$

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

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

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).

2.3 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

$$p_{s+1}(m) = \frac{p_1(m)}{\sum_{k=1}^K p_1(k) BF_{(k,m)}(z_s|d_s^*)} \quad (10)$$

$$p_{s+1}(\theta_m) = \frac{p(z_s|\theta_m, d_s^*) p_s(\theta_m)}{\int p(z_s|\theta_m, d_s^*) p_s(\theta_m) d\theta_m}. \quad (11)$$

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 emerges as a clear winner under some appropriate stopping criterion, such as $p_s(m) > 0.99$.

2.4 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. Many different forms of the choice function have been suggested, including so-called ‘‘tremble’’ models, Fechnerian random utility, or ‘‘white noise’’ models, and random preference models (see Wilcox, 2008, for a thorough review). While ADO has the flexibility to accommodate any of these forms, it is beyond the scope of this study to discriminate among choice functions and we wish to minimize the dependence of our results on the choice of a particular function. Moreover, for practical purposes, it is important that the choice function can be evaluated quickly so that the ADO calculation time between trials does not become too long. Therefore, 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. Formally, if $d_i = \{(x_S, \Delta t_S), (x_L, \Delta t_L)\}$ is the i^{th} stimulus presented in an experiment, and \prec_{θ_m} is the weak ordering over stimuli determined by the model parameter θ_m , then the probability of choosing $(x_S, \Delta t_S)$ is given by

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

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

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.

2.4.1 Priors

The full, Bayesian specification of each model includes a prior distribution over its parameters. 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 defined a uniform prior 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 literature 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): $r \in (0.0, 0.8)$, $s \in (0.8, 1.0)$, $w \in (0.0, 1.0)$

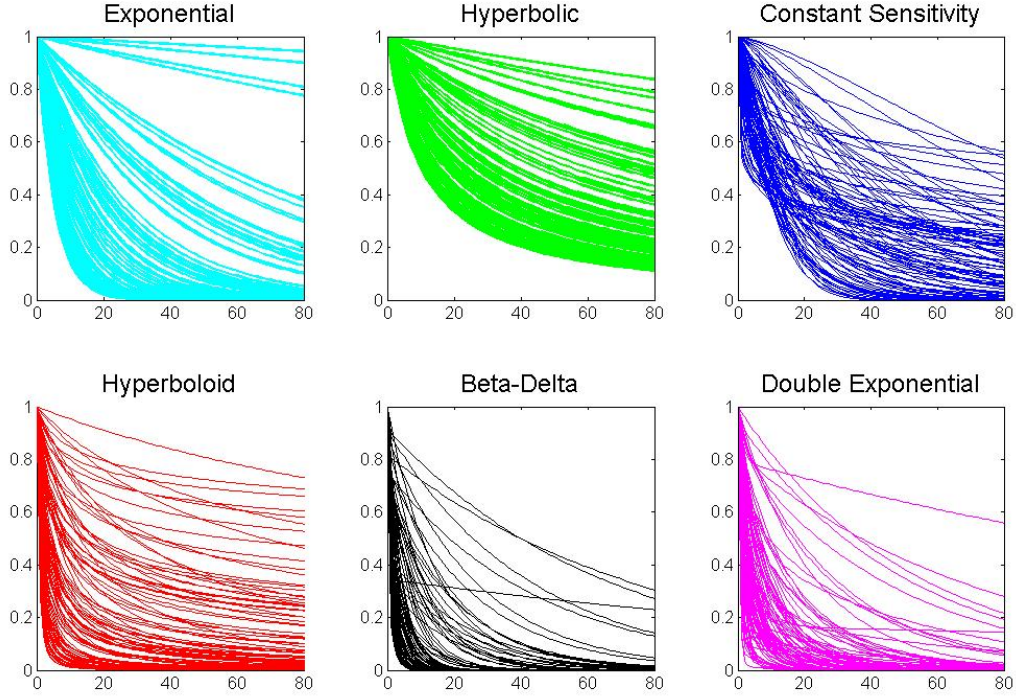


Figure 2: Depiction of the prior predictive distribution of discounting curves in each family. In each graph, the x-axis measures time and the y-axis measures the discount factor. Curves were generated by randomly selecting 100 parameter vectors from the priors for each model.

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 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.⁴ 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 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 , Δt_S , x_L and Δt_L drawn uniformly from the following intervals: $x_S \in [1, 20]$, $\Delta t_s \in [0, 40]$, $x_L \in [x_S + 1, 30]$, $\Delta t_L \in [t_S + 1, 80]$. 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 2.3.⁵ 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.

⁴ADO was implemented with two different objective functions. The first half of the trials using the objective function for parameter estimation (Equation 8), and the second half using the objective function for model discrimination (Equation 9). 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++.

⁵All 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

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

$$p_i(SS|\theta_m) = \frac{1}{1 + e^{\epsilon(V(LL)-V(SS))}} \quad (12)$$

where $SS = (x_S, \Delta t_S)$ is the smaller-sooner option, $LL = (x_L, \Delta t_L)$ is the larger-later option, $V(\cdot)$ are the subjective values determined by the algebraic discounting model, and $\epsilon > 0$ is a 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 1. 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

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.

Table 1: 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

of each simulation are given second and third columns of Table 1, 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 12 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 2, 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.

Table 2: Reanalysis of the model recovery simulation assuming a logistic choice function.

Design	p(CS)	RMSE(r)	RMSE(s)	RMSE(ϵ)
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

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 3 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 3 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.

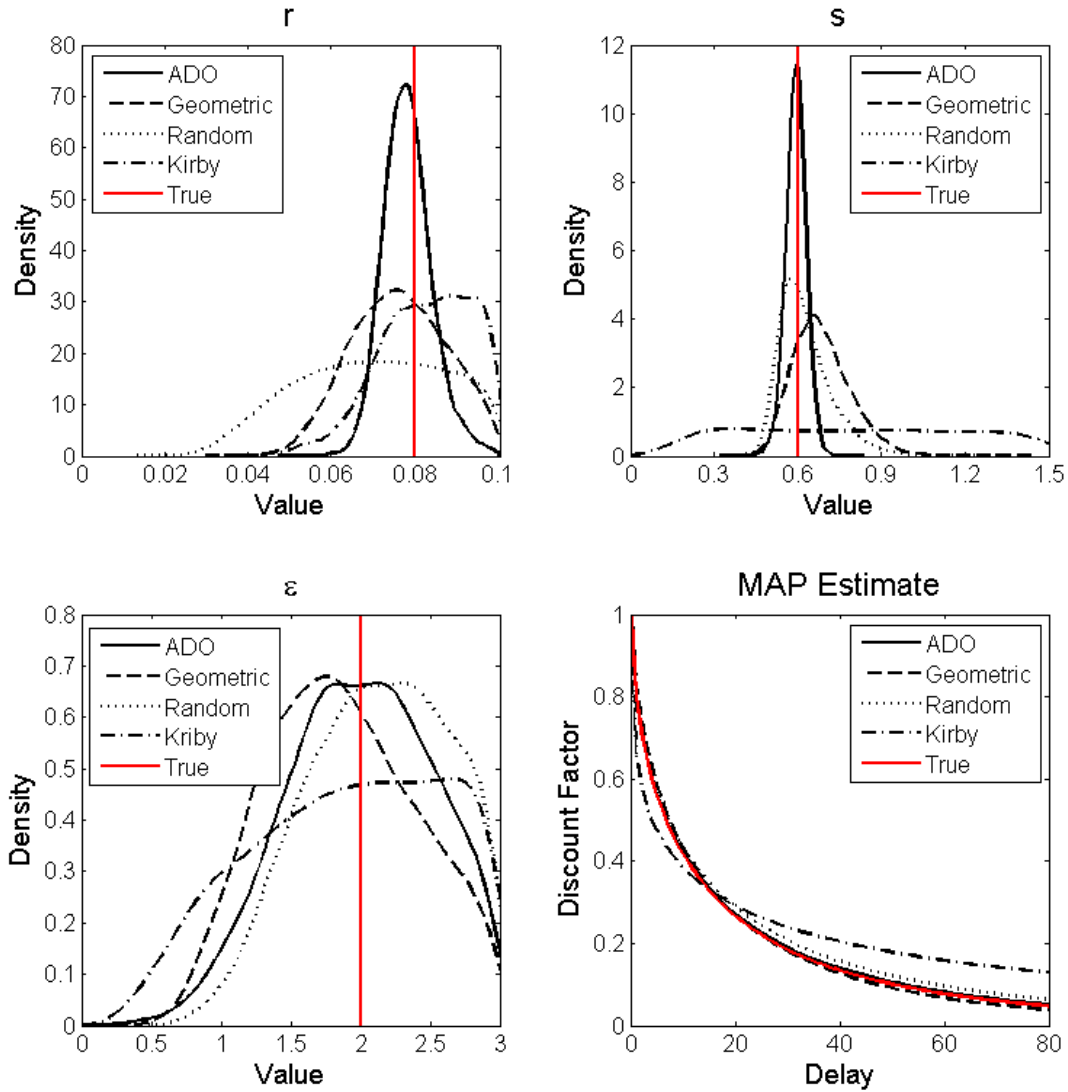


Figure 3: The four overlaid distributions in each of panels r , s and ϵ are the marginal posterior distributions obtained in each model recovery simulation. The vertical line in each graph shows the true value of the parameter, so a tighter distribution around that line indicates a better estimate. In 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.

4 Empirical study

4.1 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 was comprised 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 ($x_S, \Delta t_S$) or larger-later ($x_L, \Delta t_L$) monetary reward simultaneously displayed on either side of the monitor ($x_S \in \{8, 12, 15, 17, 19, 22\}$, $x_L \in \{12, 15, 17, 19, 22, 23\}$, $\Delta t_S \in \{0, 1, 2, 3, 5, 10, 20, 40\}$, and $\Delta t_L \in \{1, 2, 3, 5, 10, 20, 40, 80\}$). 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 trials to familiarize themselves with the task. In order to ensure “incentive compatibility,” one trial was randomly chosen following the experiment and paid to the participant, at the delay specified in the trial, in the form of a (post-dated) check. In addition, all participants received \$5.00 USD for participating in the experiment.

4.2 Results

As in the simulation studies, the data for each subject were reanalyzed assuming a logistic choice function with a free parameter. 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. In addition to the six discounting models considered in the simulations, an additional “Coin Flip” model was also included in the reanalysis of the experiment data. 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 inclusion of this model is intended to account for the fact that, in the strictest sense, all of the models are wrong. Thus, the coin-flip model serves as a sort of parameter-free null-hypothesis that allows for all of the substantive discounting models to fail.

The within-subject analyses provide the normalized posterior likelihood (i.e., probability) of each of the seven models given the observed data. A rudimentary way to assess which model did the best overall is to count the number of subjects for whom each model is the most probable of the seven. As shown in the first column of Table 3, each model was most probable for at least two subjects. 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.

Table 3: 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

There were also substantial individual differences in the qualitative shape of the discounting curve. A best estimate of the discounting curve for each subject was obtained by selecting the model with the highest posterior probability and plotting it for its maximum a posteriori (MAP) parameter value (i.e., the most probable model at the most probable values of its parameters). Figure 4 plots these best estimates on a common set of axes, color coded by model.

One drawback of simply counting the number of subjects is that it ignores the interval scale information in the posterior probabilities. A more elegant and statistically sound way to summarize the within-subject results is to actually model the data at the group level. One way to do that is to assume that every subject has the same model, so that the joint likelihood of each model can be computed as the product of the within-subject probabilities. The ratio of two such joint likelihoods is referred to as the group Bayes factor (gBF). When there are more than two models under consideration, it is useful to compute the gBF for each model relative to a baseline model. In this case, we will use the Coin Flip model as the baseline so that the gBF for each model m_i for subject j is computed as

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

Where $p_j(m_i)$ and $p_j(m_0)$ are the posterior probabilities of models m_i and m_0 for subject j .

We report the gBF for each model on a log scale in Table 3. 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.

Although Double Exponential is the most likely explanation of the data if every subject

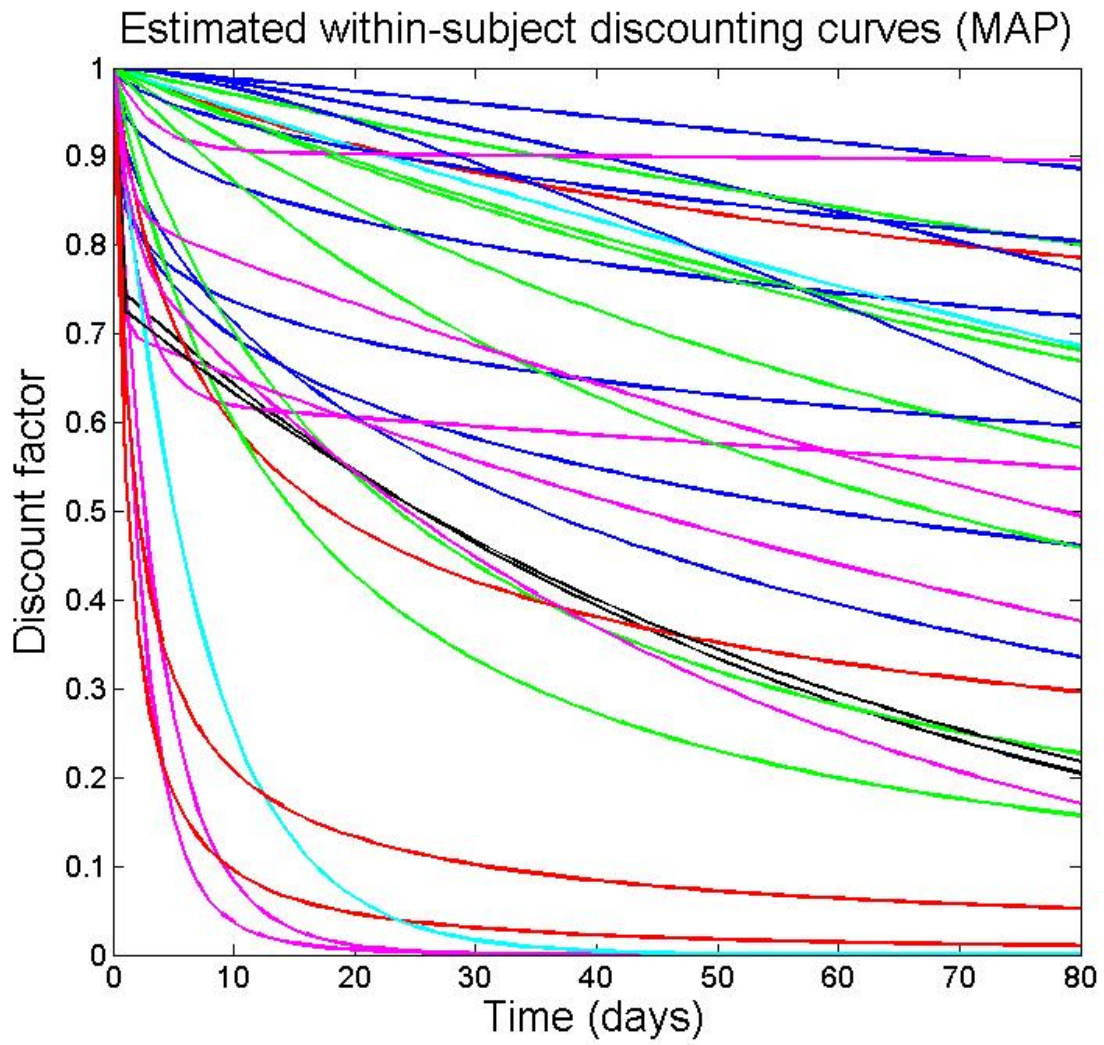


Figure 4: MAP estimate of the most probable discounting curve for each participant, color coded by model

has the same model (but different parameters), there were many subjects for whom the posterior probability of Double Exponential was quite low. Model discrimination at the subject level was often quite decisive thanks to the ADO design, and as a result, the probability of Double Exponential was even lower than that of Coin Flip for nearly half (18 out of 37) of the subjects. This suggests that the assumption that every subject has the same model is not a good one. Therefore, it would be better to model the group data in a way that allows every subject to have their own subject-level model and parameters. This can be done formally using a hierarchical model for the group that treats the subject-level model as a random variable, which allows the distribution of subject-level models to be estimated from the data (citations to justify this as a standard approach). This distribution can be approximated from the within-subject model probabilities that we already have by using a 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 $\alpha_i = 1 + \sum_j p_j(m_i)$. Thus, assuming k candidate models, the expected probability of obtaining the i^{th} model for any randomly selected subject is $\frac{\alpha_i}{\alpha_1 + \dots + \alpha_k}$.

The expected probability of each model is reported in the third column of Table 3. 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.

4.2.1 AIC results

An alternative measure frequently utilized in model comparison analyses is the Akaike Information Criterion (AIC; Akaike, 1976). The AIC is 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. Thus, like posterior model probability, the AIC trades off goodness-of-fit and complexity. However, rather than being interpreted as a likelihood of the model generating the data, the AIC is interpreted as an information-theoretic difference between the true model and the fitted model (Myung, 2000). Moreover, since it is based on the maximized likelihood instead of an averaged likelihood (weighted by a prior), the AIC is less dependent on prior assumptions than posterior model probability. If the two measures (AIC and posterior probability) select the same model then we can be reasonably confident that the decision was not overly dependent on the assumed prior (Liu and Aitkin, 2008).

The number of subjects for whom each model was best according to the AIC is given in the first column of Table 4. Although the counts are similar to the analogous counts based on posterior model probability (1st column of Table 3), the model selected by the AIC matched the model selected by posterior probability for just 18 out of 37 subjects. The discrepancy is mostly due to the bounds imposed on the parameters of some of the models in their prior distributions,

which were used to compute the posterior probabilities. In short, the best fitting parameters of some of the models for each subject were much more extreme than anticipated. In fact, they were so extreme that the maximum likelihood estimates were frequently outside the prior range. When this happens for a given model, no prior weight is given to the parameter values on which it provides the best fit, so its posterior probability is much lower than it should be. This was the case for 25 subjects for the generalized Hyperbolic model, 17 for the Constant Sensitivity model, 13 for the Hyperbolic model, and 10 for the Exponential model. As a result, those models performed much better according to the AIC than according to the posterior probability, relative to the Coin Flip model. The prior bounds on the parameters for the Beta-Delta and Double Exponential models are natural (i.e., implied by theory), hence we did not consider parameter values outside those ranges in computing the maximum likelihood. Nevertheless, the Double Exponential model also performed better according to the AIC, which could be attributed to a difference in the way that the AIC penalizes for complexity relative to the Bayes factor.

Figure 5 plots the maximum likelihood estimate of the best model (lowest AIC) for each participant. The difference between this figure and Figure 4 is that the former restricted the parameters to the support of the prior, while the latter allowed placed only theoretically meaningful constraints on the parameters. The between-subjects variability is dramatic.

The counts of how often each model had the lowest AIC provide some insight, but analyzing the group data in this way is just as rudimentary as counting the number of participants with the highest posterior probability. Therefore, to provide AIC-based counterparts to the model-based, group-level results described earlier, we used the AIC to fit and compare the same group-level models as before. This can be done in a straightforward way using Akaike weights (e.g., Burnham and Anderson, 2004). The 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, and is interpreted as the probability that the model is “best” in the sense of minimizing the information-theoretic distance (Wagenmakers and Farrell, 2004). Therefore, we can use Akaike weights to fit and compare the same group-level models as before by simply inserting them in place of posterior probabilities in the calculations. For example, the Akaike evidence ratio (analogous to the Bayes factor) can be computed as the ratio of two AIC weights, and so we compute what we call the *group AIC factor* (gAF, analogous to the gBF) for a given model m_i versus the Coin Flip model as

$$gAF_i = \prod_j \frac{w_i(AIC)}{w_0(AIC)}$$

where $w_i(AIC)$ and $w_0(AIC)$ are the Akaike weights for model m_i and the Coin Flip model, respectively.

The gAF for each model is reported in the second column of Table 4. A higher gAF means

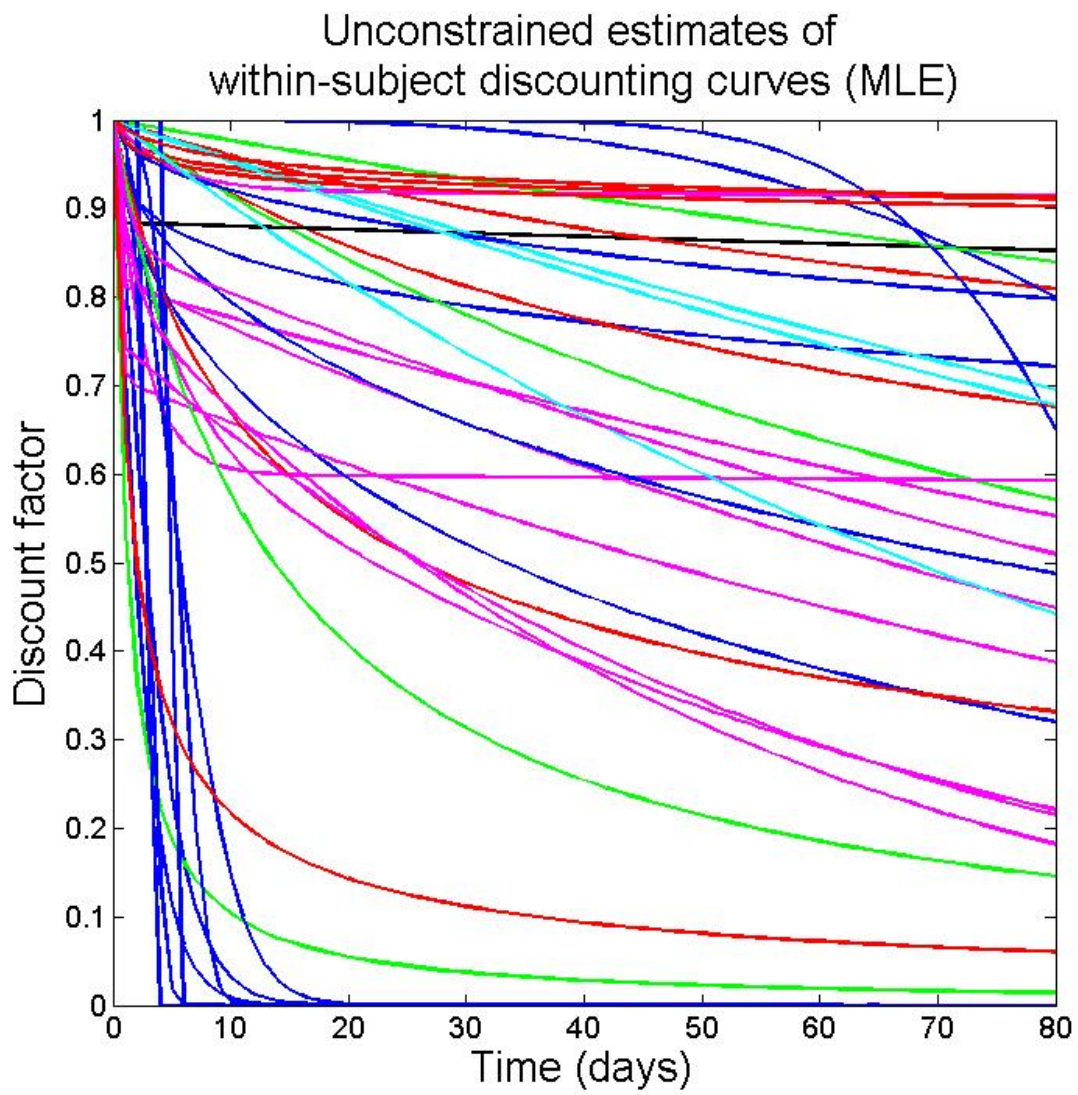


Figure 5: MLE of the model with the lowest AIC for each participant

stronger evidence in favor of that model, unlike the raw AIC for which lower values are better. Interestingly, all of the gAF’s are large and positive, indicating that they are overwhelming more likely than the Coin Flip model. The vast improvement in the performance of these models over what was seen with the gBFs can be attributed in part the AIC penalizing less for complexity. However, most of the improvement is more likely due to the fact that the maximum likelihood parameters were so frequently outside the bounds of the priors for the gBFs. The added flexibility afforded to each model by not restricting their parameters improved their overall performance substantially.

The most improved model in the AIC-based analysis is Constant Sensitivity, which is also the best group-level model overall, with a gAIC of 106.290. That is to say, if we must select one model to describe every subject (allowing parameters to vary between subjects), it should be Constant Sensitivity. The Double Exponential, which was the best group-level model according to the gBF, is the next best with a gAIC of 90.294, followed closely by the generalized Hyperbolic with a gAIC of 83.002. The gAICs of the remaining models were substantially lower.

Table 4: 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

The wAIC values can also be used to compute the expected probability of each model being the best for a randomly selected subject. This is done using the same hierarchical model as before, with a Dirichlet hyper-prior over multinomial distributions, using the wAIC in place of the posterior probability to approximate the concentration parameters of the posterior distribution. That is, each parameter α_i of the Dirichlet posterior is computed as $\alpha_i = 1 + \sum_j w_j(AIC)$.

The expected probability of each subject-level model is reported in the third column of Table 4. 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.

4.2.2 The role of *increasing* impatience and the extended present

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. Recall that, in terms of the discounting curve, increasing impatience means concavity (i.e., the discounting curve gets steeper as the delay to reward increases). Although a few recent studies have found increasing impatience (e.g., Abdellaoui et al., 2013), most commonly used discounting functions were developed to accommodate *decreasing* impatience (i.e., convexity) which is the prevailing empirical finding in intertemporal choice, particularly at the aggregate level. Among the models under consideration here, only Constant Sensitivity can accommodate increasing impatience. In particular, when $s > 1$, the Constant Sensitivity discounting curve is convex for $t \leq 0 \leq \frac{1}{a}$. Therefore, we can identify subjects who exhibit increasing impatience by examining the MLEs of s in the Constant Sensitivity model.

In our sample, the median estimate of the s parameter of the Constant Sensitivity model (i.e., the median of the within-subject MLEs) was 0.6225, indicating decreasing impatience on average. However, among the 11 subjects for whom Constant Sensitivity was best (according to the AIC), the median MLE of s was 1.261, indicating increasing impatience on average for that group. 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.

To better illustrate the qualitative differences between the models that allowed them to succeed for different subjects, the right panel of Figure 6 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 6 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.

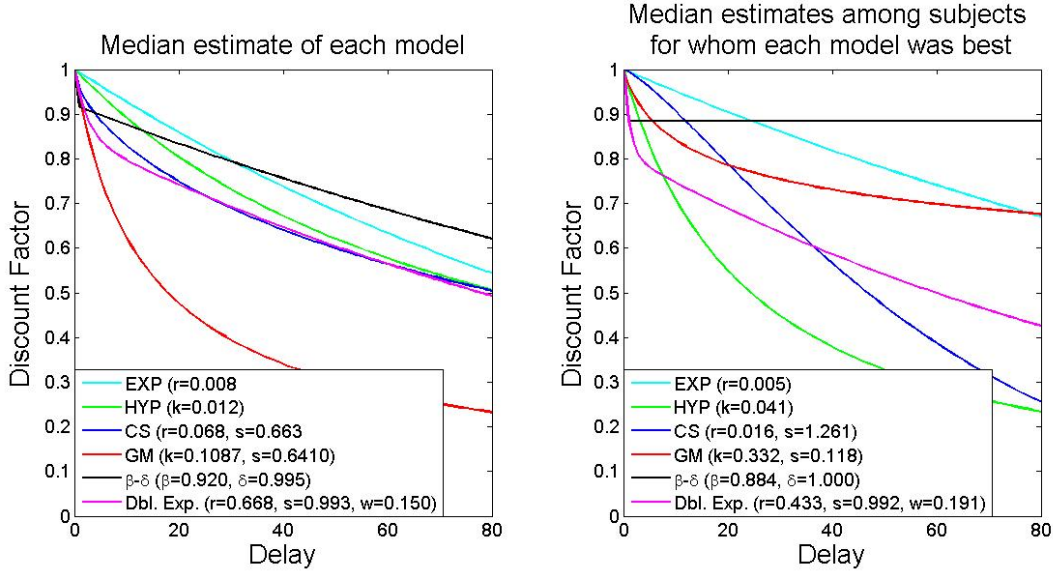


Figure 6: Comparison of median estimates of each model in the entire sample versus the subgroups in which each model was best.

The Exponential, Hyperbolic, and generalized Hyperbolic models work best when the discounting behavior is somewhere in between these extremes.

5 Discussion and conclusions

The thoroughly uncontroversial 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 conclusively demonstrated 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 with human participants, 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. Though largely overlooked, increasing impatience has been highlighted in 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 6). In these cases, the discounting curve was essentially flat initially before dropping to zero after $\Delta t = \frac{1}{r}$ 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 present a starting point for 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, A. E., and Bleichrodt, H. (2010). Intertemporal tradeoffs for gains and losses: An experimental measurement of discounted utility*. *The Economic Journal*, 120(545):845–866.
- Abdellaoui, M., Bleichrodt, H., and lHaridon, O. (2013). Sign-dependence in intertemporal choice. *Journal of Risk and Uncertainty*, 47(3):225–253.
- Ainslie, G. (1975). Specious reward: a behavioral theory of impulsiveness and impulse control. *Psychological bulletin*, 82(4):463.
- Ainslie, G. and Herrnstein, R. J. (1981). Preference reversal and delayed reinforcement. *Animal Learning & Behavior*, 9(4):476–482.
- Akaike, H. (1976). Canonical correlation analysis of time series and the use of an information criterion. *Mathematics in Science and Engineering*, 126:27–96.
- Attema, A. E., Bleichrodt, H., Rohde, K. I., and Wakker, P. P. (2010). Time-tradeoff sequences for analyzing discounting and time inconsistency. *Management Science*, 56(11):2015–2030.
- Benhabib, J., Bisin, A., and Schotter, A. (2010). Present-bias, quasi-hyperbolic discounting, and fixed costs. *Games and Economic Behavior*, 69(2):205–223.

- Berns, G. S., Laibson, D., and Loewenstein, G. (2007). Intertemporal choice—toward an integrative framework. *Trends in cognitive sciences*, 11(11):482–488.
- Bickel, W. K., Yi, R., Landes, R. D., Hill, P. F., and Baxter, C. (2011). Remember the future: working memory training decreases delay discounting among stimulant addicts. *Biological psychiatry*, 69(3):260–265.
- Bleichrodt, H., Rohde, K. I., and Wakker, P. P. (2009). Non-hyperbolic time inconsistency. *Games and Economic Behavior*, 66(1):27–38.
- Broomell, S. B. and Bhatia, S. (2014). Parameter recovery for decision modeling using choice data. *Decision*, 1(4):252.
- Burnham, K. P. and Anderson, D. R. (2004). Multimodel inference understanding aic and bic in model selection. *Sociological methods & research*, 33(2):261–304.
- Camerer, C., Loewenstein, G., and Prelec, D. (2005). Neuroeconomics: How neuroscience can inform economics. *Journal of economic Literature*, pages 9–64.
- Carter, R. M., Meyer, J. R., and Huettel, S. A. (2010). Functional neuroimaging of intertemporal choice models. *Journal of Neuroscience, Psychology, and Economics*, 3(1):27–45.
- Cavagnaro, D. R., Gonzalez, R., Myung, J. I., and Pitt, M. A. (2013a). Optimal decision stimuli for risky choice experiments: An adaptive approach. *Management science*, 59(2):358–375.
- Cavagnaro, D. R., Myung, J. I., Pitt, M. A., and Kujala, J. V. (2010). Adaptive design optimization: A mutual information-based approach to model discrimination in cognitive science. *Neural computation*, 22(4):887–905.
- Cavagnaro, D. R., Pitt, M. A., Gonzalez, R., and Myung, J. I. (2013b). Discriminating among probability weighting functions using adaptive design optimization. *Journal of risk and uncertainty*, 47(3):255–289.
- Cavagnaro, D. R., Pitt, M. A., and Myung, J. I. (2011). Model discrimination through adaptive experimentation. *Psychonomic bulletin & review*, 18(1):204–210.
- Chaloner, K. and Verdinelli, I. (1995). Bayesian experimental design: A review. *Statistical Science*, pages 273–304.
- Dallery, J. and Raiff, B. R. (2007). Delay discounting predicts cigarette smoking in a laboratory model of abstinence reinforcement. *Psychopharmacology*, 190(4):485–496.
- Dasgupta, P. (2008). Discounting climate change. *Journal of risk and uncertainty*, 37(2-3):141–169.

- Ebert, J. E. and Prelec, D. (2007). The fragility of time: Time-insensitivity and valuation of the near and far future. *Management Science*, 53(9):1423–1438.
- Frederick, S., Loewenstein, G., and O’donoghue, T. (2002). Time discounting and time preference: A critical review. *Journal of economic literature*, pages 351–401.
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., and Rubin, D. B. (2013). *Bayesian data analysis*. CRC press.
- Giordano, L. A., Bickel, W. K., Loewenstein, G., Jacobs, E. A., Marsch, L., and Badger, G. J. (2002). Mild opioid deprivation increases the degree that opioid-dependent outpatients discount delayed heroin and money. *Psychopharmacology*, 163(2):174–182.
- Green, L., Fisher, E., Perlow, S., and Sherman, L. (1981). Preference reversal and self control: Choice as a function of reward amount and delay. *Behaviour Analysis Letters*.
- Green, L. and Myerson, J. (2004). A discounting framework for choice with delayed and probabilistic rewards. *Psychological bulletin*, 130(5):769.
- Johnson, M. W. and Bickel, W. K. (2002). Within-subject comparison of real and hypothetical money rewards in delay discounting. *Journal of the experimental analysis of behavior*, 77(2):129–146.
- Kable, J. W. and Glimcher, P. W. (2007). The neural correlates of subjective value during intertemporal choice. *Nature neuroscience*, 10(12):1625–1633.
- Kass, R. E. and Raftery, A. E. (1995). Bayes factors. *Journal of the american statistical association*, 90(430):773–795.
- Kim, W., Pitt, M. A., Lu, Z.-L., Steyvers, M., and Myung, J. I. (2014). A hierarchical adaptive approach to optimal experimental design. *Neural Computation*, (26):2463–2492.
- Kirby, K. N. (1997). Bidding on the future: evidence against normative discounting of delayed rewards. *Journal of Experimental Psychology: General*, 126(1):54.
- Kirby, K. N., Petry, N. M., and Bickel, W. K. (1999). Heroin addicts have higher discount rates for delayed rewards than non-drug-using controls. *Journal of Experimental Psychology: General*, 128(1):78.
- Koffarnus, M. N., Jarmolowicz, D. P., Mueller, E. T., and Bickel, W. K. (2013). Changing delay discounting in the light of the competing neurobehavioral decision systems theory: a review. *Journal of the experimental analysis of behavior*, 99(1):32–57.

- Kontsevich, L. L. and Tyler, C. W. (1999). Bayesian adaptive estimation of psychometric slope and threshold. *Vision research*, 39(16):2729–2737.
- Kujala, J. V. and Lukka, T. J. (2006). Bayesian adaptive estimation: The next dimension. *Journal of Mathematical Psychology*, 50(4):369–389.
- Laibson, D. (1997). Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics*, pages 443–477.
- Li, X. (2008). The effects of appetitive stimuli on out-of-domain consumption impatience. *Journal of Consumer Research*, 34(5):649–656.
- Lindley, D. V. (1956). On a measure of the information provided by an experiment. *The Annals of Mathematical Statistics*, pages 986–1005.
- Liu, C. C. and Aitkin, M. (2008). Bayes factors: Prior sensitivity and model generalizability. *Journal of Mathematical Psychology*, 52(6):362–375.
- Loewenstein, G. (1996). Out of control: Visceral influences on behavior. *Organizational behavior and human decision processes*, 65(3):272–292.
- Loewenstein, G. and Prelec, D. (1992). Anomalies in intertemporal choice: Evidence and an interpretation. *The Quarterly Journal of Economics*, pages 573–597.
- Loewenstein, George and Weber, R., Flory, J., Manuck, S., and Muldoon, M. (2001). Dimensions of time discounting. Paper Presented at Conference on Survey Research on Household Expectations and Preferences, Ann Arbor, Nov. 2-3.
- MacKillop, J. and Kahler, C. W. (2009). Delayed reward discounting predicts treatment response for heavy drinkers receiving smoking cessation treatment. *Drug and alcohol dependence*, 104(3):197–203.
- Madden, G. J. and Bickel, W. K. (2010). *Impulsivity: The behavioral and neurological science of discounting*. American Psychological Association.
- Mazur, J. E. (1987). An adjusting procedure for studying delayed reinforcement. *Quantitative analyses of behavior*, 5:55–73.
- McClure, S. M. and Bickel, W. K. (2014). A dual-systems perspective on addiction: contributions from neuroimaging and cognitive training. *Annals of the New York Academy of Sciences*, 1327(1):62–78.
- McClure, S. M., Ericson, K. M., Laibson, D. I., Loewenstein, G., and Cohen, J. D. (2007). Time discounting for primary rewards. *The Journal of Neuroscience*, 27(21):5796–5804.

- McClure, S. M., Laibson, D. I., Loewenstein, G., and Cohen, J. D. (2004). Separate neural systems value immediate and delayed monetary rewards. *Science*, 306(5695):503–507.
- McKerchar, T. L., Green, L., Myerson, J., Pickford, T. S., Hill, J. C., and Stout, S. C. (2009). A comparison of four models of delay discounting in humans. *Behavioural Processes*, 81(2):256–259.
- Metcalfe, J. and Mischel, W. (1999). A hot/cool-system analysis of delay of gratification: dynamics of willpower. *Psychological review*, 106(1):3.
- Moore, S. C. and Cusens, B. (2010). Delay discounting predicts increase in blood alcohol level in social drinkers. *Psychiatry research*, 179(3):324–327.
- Myerson, J. and Green, L. (1995). Discounting of delayed rewards: Models of individual choice. *Journal of the experimental analysis of behavior*, 64(3):263–276.
- Myung, I. J. (2000). The importance of complexity in model selection. *Journal of Mathematical Psychology*, 44(1):190–204.
- Myung, J. I., Cavagnaro, D. R., and Pitt, M. A. (2013). A tutorial on adaptive design optimization. *Journal of mathematical psychology*, 57(3):53–67.
- Myung, J. I., Cavagnaro, D. R., and Pitt, M. A. (in press). Model evaluation and selection. In Batchelder, W. H., Colonius, H., Dzhafarov, E., and Myung, J. I., editors, *New Handbook of Mathematical Psychology*. Cambridge University Press, London.
- Myung, J. I. and Pitt, M. A. (2009). Optimal experimental design for model discrimination. *Psychological review*, 116(3):499.
- Peters, J. and Büchel, C. (2011). The neural mechanisms of inter-temporal decision-making: understanding variability. *Trends in cognitive sciences*, 15(5):227–239.
- Peters, J., Miedl, S. F., and Büchel, C. (2012). Formal comparison of dual-parameter temporal discounting models in controls and pathological gamblers. *PloS one*, 7(11):e47225.
- Pine, A., Seymour, B., Roiser, J. P., Bossaerts, P., Friston, K. J., Curran, H. V., and Dolan, R. J. (2009). Encoding of marginal utility across time in the human brain. *The Journal of Neuroscience*, 29(30):9575–9581.
- Pine, A., Shiner, T., Seymour, B., and Dolan, R. J. (2010). Dopamine, time, and impulsivity in humans. *The Journal of Neuroscience*, 30(26):8888–8896.
- Rachlin, H. (2006). Notes on discounting. *Journal of the experimental analysis of behavior*, 85(3):425–435.

- Rangel, A., Camerer, C., and Montague, P. R. (2008). A framework for studying the neurobiology of value-based decision making. *Nature Reviews Neuroscience*, 9(7):545–556.
- Ray, D., Golovin, D., Krause, A., and Camerer, C. (2012). Bayesian rapid optimal adaptive design (broad): Method and application distinguishing models of risky choice. *California Institute of Technology working paper*.
- Reynolds, B. (2006). A review of delay-discounting research with humans: relations to drug use and gambling. *Behavioural pharmacology*, 17(8):651–667.
- Samuelson, P. A. (1937). A note on measurement of utility. *The Review of Economic Studies*, 4(2):155–161.
- Scholten, M. and Read, D. (2006). Beyond discounting: the tradeoff model of intertemporal choice. *Operations research working papers*.
- Scholten, M. and Read, D. (2010). The psychology of intertemporal tradeoffs. *Psychological Review*, 117(3):925.
- Sharp, C., Barr, G., Ross, D., Bhimani, R., Ha, C., and Vuchinich, R. (2012). Social discounting and externalizing behavior problems in boys. *Journal of Behavioral Decision Making*, 25(3):239–247.
- Shefrin, H. M. and Thaler, R. H. (1988). The behavioral life-cycle hypothesis. *Economic inquiry*, 26(4):609–643.
- Stephan, K. E., Penny, W. D., Daunizeau, J., Moran, R. J., and Friston, K. J. (2009). Bayesian model selection for group studies. *Neuroimage*, 46(4):1004–1017.
- Story, G. W., Vlaev, I., Seymour, B., Darzi, A., and Dolan, R. J. (2014). Does temporal discounting explain unhealthy behavior? a systematic review and reinforcement learning perspective. *Frontiers in behavioral neuroscience*, 8.
- Takahashi, T., Oono, H., and Radford, M. H. (2008). Psychophysics of time perception and intertemporal choice models. *Physica A: Statistical Mechanics and its Applications*, 387(8):2066–2074.
- Thaler, R. (1981). Some empirical evidence on dynamic inconsistency. *Economics Letters*, 8(3):201–207.
- Toubia, O., Johnson, E., Evgeniou, T., and Delquié, P. (2013). Dynamic experiments for estimating preferences: An adaptive method of eliciting time and risk parameters. *Management Science*, 59(3):613–640.

- Van den Bergh, B., Dewitte, S., and Warlop, L. (2008). Bikinis instigate generalized impatience in intertemporal choice. *Journal of Consumer Research*, 35(1):85–97.
- van den Bos, W. and McClure, S. M. (2013). Towards a general model of temporal discounting. *Journal of the experimental analysis of behavior*, 99(1):58–73.
- Wagenmakers, E.-J. and Farrell, S. (2004). Aic model selection using akaike weights. *Psychonomic bulletin & review*, 11(1):192–196.
- Wang, S. W., Filiba, M., and Camerer, C. F. (2010). Dynamically optimized sequential experimentation (dose) for estimating economic preference parameters. *California Institute of Technology working paper*.
- Wang, X. and Dvorak, R. D. (2010). Sweet future fluctuating blood glucose levels affect future discounting. *Psychological Science*, 21(2):183–188.
- Weiss, D. J. and Kingsbury, G. (1984). Application of computerized adaptive testing to educational problems. *Journal of Educational Measurement*, 21(4):361–375.
- Wilcox, N. T. (2008). Stochastic models for binary discrete choice under risk: A critical primer and econometric comparison. *Research in experimental economics*, 12:197–292.
- Wilson, M. and Daly, M. (2004). Do pretty women inspire men to discount the future? *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 271(Suppl 4):S177–S179.
- Zauberman, G., Kim, B. K., Malkoc, S. A., and Bettman, J. R. (2009). Discounting time and time discounting: Subjective time perception and intertemporal preferences. *Journal of Marketing Research*, 46(4):543–556.