Estimating Time Preferences from Budget Set Choices Using Optimal Adaptive Design

Taisuke Imai Colin F. Camerer *

September 4, 2018

Abstract

We describe a method for choosing an informationally-optimal sequence of questions in experiments, using subjects' responses to previous questions. The method is applied to induced budget experiments, in which subjects choose allocation of monetary rewards at sooner and later dates, to elicit time preferences. "Ground truth" simulations create artificial choice data based on known parameters and then applies the method, and show how accurately and quickly parameter values can be recovered. Results from online experiments further validate the advantage of our adaptive procedure over the typical designs in which the question sequence is not optimized. The resulting parameter estimates from the adaptive procedure are close to typical values measured in previous studies, but it converges faster than non-optimized designs. The way the adaptive procedure achieves higher accuracy and speed is expressed in subjects' negatively autocorrelated choice patterns, which is a result of the algorithm's active search of informative budget slopes.

Keywords Adaptive designs; Parameter estimation; Model selection; Time preferences

JEL Classifications C90, C52, D90

*Imai: Department of Economics, LMU Munich, 80539, Munich, Germany. E-mail: taisuke.imai@econ.lmu.de. Camerer: Division of the Humanities and Social Sciences, California Institute of Technology, Pasadena, CA 91125, USA. E-mail: camerer@hss.caltech.edu.

Acknowledgement The authors are grateful to Noriko Imai, Gaini Kussainova, and Lasya Venneti, who helped us to develop the software for our adaptive design experiments, and to Jose Apesteguia, Miguel Ballester, John Clithero, Federico Echenique, Ben Gillen, Matthew Shum, Ryan Webb, Romann Weber, Joachim Winter, and Leeat Yariv for helpful discussions and suggestions. The authors are grateful to comments from audience at Economic Science Association 2016 North American Meetings, Society for Judgment and Decision Making 2016 Annual Conference, and Barcelona GSE Summer Forum 2017 (Stochastic Choice session), and 2nd CRC Retreat in Tutzing, and seminar participants at Caltech, University College London, University of Amsterdam, LMU Munich, Osaka University (Institute of Social and Economic Research), Kyoto University (Institute of Economic Research), University of Warwick, Ohio State University, University of Melbourne, and University of Vienna. This research was supported by the Behavioral and Neuroeconomics Discovery Fund at Caltech. Imai acknowledges financial support from Deutsche Forschungsgemeinschaft through CRC TRR 190.

1 Introduction

Efficient accumulation of empirical knowledge about individual's characteristics has become increasingly valuable in many domains of social science, policy design, marketing, and management. In this paper we advance one type of optimal method and apply it to inference about time preference.

In contrast to the approach pursued in this paper, most methods to measure individual-level constructs (like time preference) in experimental social science are developed by intuitive hunches about what types of questions will be precise, easy to implement, understandable to a range of human subjects, and likely to be reproducible. New methods are tried out and adjusted by trial-and-error testing. Then a *de facto* standard method often emerges. Methods become conventional when standardization is useful, because findings produced by a common method can be more easily compared.

In experimental economics, two primary methods have become the conventional ways to measure time preference (Cheung, 2016; Cohen et al., 2016). The older method is asking people to choose between a reward that is smaller but arrives sooner (called SS) and a larger reward but which arrives later (called LL). These choices are typically offered in the form of Multiple Price List (e.g., Andersen et al., 2008; Coller and Williams, 1999; Harrison et al., 2002; Laury et al., 2012; Takeuchi, 2011) or sequential binary choice (used frequently in brain imaging studies, e.g., Kable and Glimcher, 2007, 2010; McClure et al., 2004; Peters and Büchel, 2010).

In the second method, subjects allocate a fixed budget of monetary rewards at each of the two dates (this is called a Convex Time Budget design; Andreoni and Sprenger, 2012; Andreoni et al., 2015; Augenblick et al., 2015). Let us be more precise: Consider two time points t_1 and t_2. A linear budget set of allocations of monetary rewards to be received at those two times is a line connecting two points $(\bar{x}_{t_1}, 0)$ and $(0, \bar{x}_{t_2})$ on a two-dimensional plane. The first point corresponds to an agent receiving a certain amount $\bar{x}_{t_1}$ of reward at time t_1 and nothing at t_2. The second point corresponds to receiving a certain amount $\bar{x}_{t_2}$ at time t_2 and nothing at t_1. Any points on the interior of a budget set represent allocations in which she receives positive rewards on both dates.

In order to identify and estimate parameters of different kinds of time preferences, an experimenter needs to vary the time points (t_1, t_2), the slopes of the budget lines, and the level of the budget lines. Each budget line can be expressed as a set of these numbers.

The overall design question is how to select a *set* of budget lines to best estimate time preferences. In almost all previous studies, the set of budget lines was *predetermined*. Every subject in an experimental treatment thus faced the same budgets, although the orders of presentation could be different across subjects.

This paper uses a different approach, which we call DOSE (an acronym for **D**ynamically **O**ptimized **S**equential **E**xperimentation; the methodology introduced originally by Wang et al., 2010; Chapman et al., 2018), and applies it to estimation of time preferences. [1] In general, the DOSE method requires precise specification of several ingredients: (i) A domain of possible questions (e.g., a set of all possible budget lines); (ii) A set of alternative hypotheses $\mathcal{H}$ (typically, combinations of parameterized theories such as an exponential discounting function δ^t, or a hyperbolic discounting function $1/(1+kt)$, with specific values of parameters); (iii) A prior probability over the set $\mathcal{H}$; and (iv) An information criterion, which is used to measure numerically which question is *expected* to best distinguish the hypotheses in $\mathcal{H}$.

The DOSE algorithm chooses a sequence of budget lines that are optimally informative, as measured by the information criterion. After a subject makes a choice, the posterior probabilities of all hypotheses in $\mathcal{H}$ are updated using Bayes' rule. The posterior is substituted for the prior in step (iii) above and the budget line with the highest information value (computed using (iv) above) is chosen for the next experimental trial or survey item. The algorithm repeats this procedure until it hits a pre-specified stopping criterion such as the maximum number of questions or some function of the posteriors (e.g., when one hypothesis passes a high threshold). Intuitively, when the sequence of choices is customized for each subject in this way, the subjects themselves tell us, through their answers, the "best" (i.e., most informative) question to ask them next.

There are several potential advantages of DOSE approaches in general.

First, the DOSE method maximizes information gained per question. Therefore, they could

[1] Others have developed similar adaptive approaches to estimate preferences; they are described below.

be particularly useful for subject pools who have a high opportunity cost of time, or become bored or habituated quickly. Such groups include highly-trained professionals, subjects in online experiments (such as Amazon's Mechanical Turk) who quit if experiments are too long (creating problems of inference based on attrition), human groups such as lesion patients or children, and animals that typically make long sequences of lab choices.

Second, the posterior distribution of all hypotheses is computed for each subject after each question, since it is a crucial necessary step (ingredient (iv) described above) in finding the most informative budget line for the upcoming trial. Therefore, if the main purpose of the experiment is inferences about preferences, the analysis is already done when the subjects complete the task.

Third, the DOSE method creates an instant statistical parametric assessment of each subject after their experimental session is ended. These portraits can show which subjects seem most impatient, most averse to risk, most reciprocal, most able to learn quickly, most strategic, and so on. These data could then be used to instantly cherry-pick different statistical types of people for the next phase of an experiment. This feature will be particularly useful for experiments with brain imaging using functional magnetic resonance imaging (fMRI) machines. A pre-scanning choice task with DOSE procedure gives researchers sufficient information to individually tailor a set of questions to be presented inside the scanner. [2]

Last, the fact that the DOSE method generates sequences of questions that are provably optimal (given the priors) can sharpen discourse about what different experimental designs are good and bad for. Novel designs which are unconventional should gain credibility if they have desirable informational properties. The DOSE method can be used in pre-experiment simulation to select the best fixed set of questions for survey modules. [3] DOSE methods can also be used to judge the quality of older conventional designs.

We demonstrate the DOSE method in the context of time preference elicitation because those

[2]A similar approach has been taken in several existing brain imaging studies, in which discounting function estimated from SS-LL choices generated by a staircase procedure is used to construct individually-tailored set of questions in later fMRI task (e.g., van den Bos et al., 2014, 2015).

[3]In Falk et al. (2016), questions are selected by identifying the combination of survey items from an extensive battery of alternative survey questions that best predicts choices in incentivized experiments.

estimated preferences are important, often surprisingly different, and may depend systematically on elicitation procedures.

Measures of time preference play a crucial role in many areas of applied economics. Discount rates are likely to influence any choice that reflects valuation of costs and benefits spread over time. Domains include health (food and exercise), education, financial markets, personal and household finance. Reliably estimating individual differences in time preference is useful for explaining variation in choices, development of patience in children, and for creating computational phenotypes of psychiatric disorders.

Furthermore, a huge number of studies show large differences in estimated time preferences (Cohen et al., 2016; Frederick et al., 2002). People are estimated to be more patient for larger magnitudes, for losses compared to gains, and for getting benefits sooner compared to delaying them. There are also substantial differences depending on how attributes of different time-dated rewards are described or emphasized.

We noted earlier that the two most popular methods for measuring time preference are pairwise SS-LL choices, and choosing an allocation from a Convex Time Budget (CTB). A person choosing a point on the budget line is generating more information because they are comparing many different time-reward bundles at a time. An advantage of budget line experiments is that they enable a test of consistency of choices with revealed preference conditions, such as the Generalized Axiom of Revealed Preference (GARP; Afriat, 1967). [4]

Linear budgets experiments have become a popular method for studying individual preferences in laboratory and field. Its first use, to our knowledge, was by Loomes (1991). Linear budgets have been used to study social preferences (e.g., Andreoni and Vesterlund, 2001; Andreoni and

[4] This is a nonparametric test of utility maximization and together with a measure of degree of violation, such as Afriat's (1972) Critical Cost Efficiency Index (CCEI) or the Money Pump Index by Echenique et al. (2011), researchers can quantify the "quality" of decision making of each individual. Virtually all studies show high consistency with GARP (Choi et al., 2007, 2014), including studies with children (Harbaugh et al., 2001) and capuchin monkeys (Chen et al., 2006). For CTB choice data, Echenique et al. (2016b) propose rnonparametric revealed preference tests and measures of degree of violations for several models including exponentially discounted utility model, quasi-hyperbolic discounted utility model, and time-separable utility model.

Miller, 2002; Fisman et al., 2007), risk and ambiguity preferences (e.g., Ahn et al., 2014; Choi et al., 2007, 2014; Loomes, 1991), time preferences (Andreoni and Sprenger, 2012; Augenblick et al., 2015, among others presented in Table A.1 in Appendix A), and general utility maximization with consumer goods and foods as rewards (Burghart et al., 2013; Camille et al., 2011; Harbaugh et al., 2001; Sippel, 1997). In this paper we apply the DOSE to CTB environment, but in principle it is applicable to linear budget experiments with any domains of choices.

Budget line methods are appealing because they generate more information than binary choices (by presenting more choices on each question). However, it is also possible that the complexity of choosing just one point on a line generates different expressed preferences than other methods. [5] The general possibility that two methods produce conflicting results is called "procedure-variance"—i.e., elicited preferences can be sensitive to the procedure used to elicit those preferences. Procedure-variance has been the subject of much research in psychology and behavioral economics (e.g., choice-matching preference reversals; see Tversky et al., 1990), but less in experimental economics.

In any case, the CTB method has caught on quickly. It has been used in more than 40 studies, both in the laboratory, in lab-in-field tests, and in representative surveys (see Table A.1 in Appendix A). [6] However, the earliest estimates of time preference measured using CTB are quite different than measures delivered from binary choice. In Andreoni and Sprenger (2012), for example, there is very little concavity of utility for money and no evident present bias for money. However, rates of time discounting are comparable to those in many other studies (around 30%/year).

[5] Andreoni and Sprenger (2012) did compare the CTB estimates to those from a "double multiple price list" (list of pairwise choices, for both time and risk; Andersen et al., 2008). A focus of many studies, including ours, is specifications in which immediate rewards are weighted by one, and future rewards at time t are weighted by $\beta\delta^t$ (Laibson, 1997; Phelps and Pollak, 1968). The parameter β is a preference for immediacy, or present-bias. The parameter δ is a conventional discount factor. Note that when $\beta = 1$ this quasi-hyperbolic specification reduces to exponential discounting. In the Andreoni and Sprenger's (2012) analysis, the correlation of the inferred discount rates δ in CTB and double multiple price list, within subjects, was 0.42. At the same time, their estimates of β are quite close to one, while most other methods estimate $\beta < 1$.

[6] CTB datasets from some of the published studies are systematically analyzed in Echenique et al. (2016a,b) using a nonparametric revealed preference approach and in Imai et al. (2018) using a meta-analytic approach.

There are also a large majority of allocations chosen as endpoints (also called as corners) of budget lines (i.e., all tokens allocated to rewards at only one date). If endpoint choices are common, more information will be gained by systematically tilting budget line slopes up and down more aggressively than is done in a fixed-sequence design (in order to flip choices from one endpoint to another). The DOSE method applied to CTB will specific exactly how to do that most efficiently.

2 Background

The DOSE method is an innovation in a developing family of adaptive methods used in various fields (though not much in experimental economics publications). The major contribution is a particular measure of information value, called *Equivalence Class Edge Cutting (EC^2)*, which is adaptively submodular, which therefore provably guarantees some useful theoretical and practical properties. The method was introduced in computer science by Golovin et al. (2010), and it is applied here to novel economic questions. See Appendix B for the theoretical background of this information value.

Earlier applications of optimal design methods were made in statistics (Lindley, 1956), decision theory (Howard, 1966), computer-assisted testing (CAT) in psychometrics (e.g., Wainer and Lewis, 1990) and Bayesian experimental design (Chaloner and Verdinelli, 1995).

Adaptive methods extended these approaches to trial-by-trial question choice to optimize information gain. Examples include cognitive psychology (e.g., Myung and Pitt, 2009), adaptive choice-based conjoint measurement in marketing (e.g., Abernethy et al., 2008), and "active learning" methods in computer science (Golovin and Krause, 2010) and machine learning (Nowak, 2009; Dasgupta, 2004). Existing methods created by psychologists and economists to measure parameters such as risk aversion include Cavagnaro et al. (2010, 2011, 2013b,a, 2016), Myung et al. (2009, 2013), Toubia et al. (2013), and Chapman et al. (2018). [7] We compare our method and these

[7] One unpublished paper (Ray et al., 2012) applied the EC^2 criterion in a similar adaptive design framework which they called Bayesian Rapid Optimal Adaptive Design (BROAD), but did not use a clear user interface like ours in the experiments, did not compare BROAD with other sequencing methods, and did not report parameter estimates—which are the numerical results of most interest for economics. Ray et al. (2012) focused on computer-scientific

existing ones in Section 3.5.

Computer scientists have shown that finding an optimal sequence of test choices is not just computationally difficult (NP-hard) but is also difficult to approximate (Chakaravarthy et al., 2007). Several heuristic approaches have been proposed that perform well in some specific applications, but do not have theoretical guarantees (e.g., MacKay, 1992); that is, there are no proofs about how costly the heuristic sequence will be compared to the optimal sequence. (The concept of "costly" in computer science is roughly the number of trials.)

Note that early efforts to introduce static optimal design in experimental economics (El-Gamal et al., 1993; El-Gamal and Palfrey, 1996; Moffatt, 2007, 2016) did not gain traction. The time is now riper for DOSE methods because: Computing power is better than ever; scalable cloud computing services such as Amazon's Elastic Compute Cloud and Microsoft's Azure, are available at a reasonable cost; the new method from computer science (EC^2) applied here provides theoretical guarantees on efficient computability; and there are many new competing theories in behavioral economics which need to be efficiently compared.

In experimental psychology and economics, there are two popular approaches for dynamic selection of question items. In the (binary-choice) *staircase* method, originally developed in psychophysics (Cornsweet, 1962; von Békésy, 1947), one option is fixed while the other option varies from trial to trial, reflecting the subject's response in the previous trial. The method can be used to identify indifference points without Multiple Price List (also called as the *bisection method*; see, e.g., Abdellaoui, 2000; Dimmock et al., 2016; van de Kuilen and Wakker, 2011). In the *iterative Multiple Price List* (e.g., Andersen et al., 2006; Brown and Kim, 2014), subjects complete two lists where the second list has a finer gradation of choices around the option chosen in the first list. The crucial difference between our adaptive procedure and those existing ones is that the latter do not maximize objective measures of informativeness of questions, which DOSE does.

In economic choice applications, there is one possible imperfection in DOSE methods: In theory, subjects might prefer to strategically manipulate their early responses in order to get "better" (more economically valuable) future questions. This is a potential problem because a

aspects of the method and demonstrated advantages of EC^2 over other known algorithms in computer science, including information gain, value of information, and generalized binary search.

strategic earlier choice is likely to be different from the choice they would make if they were making only a single choice.

There are some sensible arguments against why strategizing is unlikely, and several types of evidence that it is not occurring. Since it is easier to understand these arguments and evidence after learning more about the method, and digesting our empirical results, we postpone them to a penultimate section before the conclusion.

3 Adaptive Experimental Design Method

The benchmark model for choices of rewards distributed over time is exponential discounting (Koopmans, 1960; Samuelson, 1937). There is also a huge literature from psychology, behavioral economics, animal behavior, and neuroscience providing evidence that human behavior is often time-inconsistent, and people are willing to forego larger delayed rewards for smaller rewards if they are immediate (Cohen et al., 2016; Frederick et al., 2002). Descriptive models that account for this departure from rationality vary from the one-parameter hyperbolic discounting function (Mazur, 1987), to *present-bias* models, such as quasi-hyperbolic discounting (Laibson, 1997; Phelps and Pollak, 1968), and fixed time cost models that have an additional parameter to account for the observation that people pay a premium to choose options that are immediately available (Benhabib et al., 2010). Models of time preference are useful in decision making in many contexts, including consumer behavior, health (Gafni and Torrance, 1984), savings and consumption (Angeletos et al., 2001), and organizing work (e.g., responses to deadlines O'Donoghue and Rabin, 1999). Given the range of available models, a framework for efficiently comparing time preference models can help choose the most descriptive model quickly.

3.1 Environment

We extend adaptive design methods developed for binary choice experiments to an environment with linear budgets as in Andreoni and Miller (2002), Choi et al. (2007), and Andreoni and Sprenger (2012). This extension is straightforward if the continuous range of possible allocations on the budget line is discretized. [8]

Let $\mathcal{M}$ denote the set of *model classes* and $h \in \mathcal{H}$ denote a *hypothesis*, which is a combination of a model class and a specific parametrization. For example, exponential discounting with discount factor $\delta = 0.98$ can be one hypothesis and quasi-hyperbolic discounting with a pair of present bias and discount factor $(\beta, \delta) = (0.95, 0.99)$ can be another hypothesis. We are endowed with a *prior* μ_0 over $\mathcal{H}$. We assume $\mu_0(h) > 0$ for all $h \in \mathcal{H}$ by pruning zero-prior hypotheses from $\mathcal{H}$ in advance. The subset $\mathcal{H}_m \subseteq \mathcal{H}$ denotes the set of sub-hypotheses (i.e., different parameter specifications) under model $m \in \mathcal{M}$. Note that the sets of sub-hypotheses need not be nested each other (i.e., $\{\mathcal{H}_m\}_{m \in \mathcal{M}}$ is not necessarily a partition of $\mathcal{H}$). The method is able to distinguish among nonnested models (such as exponential discounting model and hyperbolic discounting model).

Let $\mathcal{Q}$ denote the set of all *questions*. A question consists of two options in case of binary choice experiments, while it is a (discrete) budget set in case of liner budget experiments. Let $\mathcal{X}_q$ denote the set of all possible *responses* (or *answers*) to question $q \in \mathcal{Q}$. We can suppress the subscript q by standardizing the representation of responses. For example, $\mathcal{X} = \{0, 1\}$ would represent the set of available options, the left option (0) and the right option (1) in a binary choice question, and $\mathcal{X} = \{0, 1, \ldots, 99, 100\}$ would represent 101 equidistant points on a budget line. [9] We use X to represent a random variable on $\mathcal{X}$.

Let $r \in \mathbf{N}$ represent the *round* in the task. For example, $q_r \in \mathcal{Q}$ indicates that question q_r was presented at round r and $x_r \in \mathcal{X}$ indicates that x_r was selected as a response to that question. A vector $\boldsymbol{q}_r$ represents a sequence of questions presented up to round r, i.e., $\boldsymbol{q}_r = (q_1, q_2, \ldots, q_r)$. Similarly, a vector $\boldsymbol{x}_r = (x_1, x_2, \ldots, x_r)$ represents a sequence of responses up to round r. Combining those, a pair of vectors $\boldsymbol{o}_r = (\boldsymbol{q}_r, \boldsymbol{x}_r)$ summarizes what have been asked and observed so far, which we simply call an *observation*. The set of observations after round r is $\mathcal{O}_r$ and we let $\mathcal{O} = \bigcup_{r \geq 1} \mathcal{O}_r$ denote the set of all possible observations. After every round r,

[8]Discretization is harmless because most subjects choose a very limited set of round numbers, which are approximate multiples of 100/10, 100/4 or 100/3.

[9]One can also view this representation as an allocation of 100 experimental "tokens" into two accounts, each of which is associated with different monetary value as in Andreoni and Sprenger (2012).

Table 1: List of variables.

Variable	Description
$\mathcal{M}$	The set of model classes
$\mathcal{H}$	The set of hypotheses
$\mathcal{H}_m \subseteq \mathcal{H}$	The set of hypotheses under model class $m \in \mathcal{M}$
$\mathcal{Q}$	The set of questions
$\mathcal{X}$	The set of responses to questions
$\boldsymbol{q}_r = (q_1, \ldots, q_r)$	A sequence of questions up to round r
$\boldsymbol{x}_r = (x_1, \ldots, x_r)$	A sequence of responses up to round r
$\boldsymbol{o}_r = (\boldsymbol{q}_r, \boldsymbol{x}_r)$	A sequence of observations (question-response pairs) up to round r
μ_0	A prior belief over $\mathcal{H}$ s.t. $\mu_0(h) > 0$ for all $h \in \mathcal{H}$ and $\sum_{h \in \mathcal{H}} \mu_0(h) = 1$
$\mu_r(\cdot \vert \boldsymbol{o}_r)$	A posterior belief after observing $\boldsymbol{o}_r$

we update our beliefs to $\mu_r(\cdot|\boldsymbol{o}_r)$ by the Bayes' rule. See Table 1 as a reference to those notations and definitions. As usual, $\mathbf{E}$ stands for the expectation operator with respect to an appropriate measure and $\Pr$ is a generic probability measure.

3.2 The Information Value of Questions

Quantifying the information value of questions is the most crucial part of adaptive experimental design methods. In the current study, we consider a particular type of *informativeness function* $\Delta : \mathcal{Q} \times \mathcal{O} \to \mathbf{R}$, the *Equivalence Class Edge Cutting (EC^2)* criterion, proposed originally in Golovin et al. (2010) and later used in an unpublished work (Ray et al., 2012). [10]

Given the sequence of questions and responses $\boldsymbol{o}_r = (\boldsymbol{q}_r, \boldsymbol{x}_r)$, we define the *EC^2 informational*

[10]In the early phase of this project, we also considered another informativeness function based on the Kullback-Leibler (KL) divergence (Kullback and Leibler, 1951), following El-Gamal and Palfrey (1996) and Chapman et al. (2018). In the simulation exercises we found that this informativeness function is significantly slower than EC^2 criterion in preparation of next question. Since computational speed is essential in some applications, we decided not to pursue comparison of EC^2 and KL.

value Δ_{EC^2} of question $q \in \mathcal{Q} \backslash \{q_1, \ldots, q_r\}$ to be asked in round $r+1$ by:

$$\Delta_{\mathrm{EC}^2}(q|\boldsymbol{o}_r) = \left[\sum_{x \in \mathcal{X}_q} \Pr[X_{r+1} = x|\boldsymbol{o}_r] \left(\sum_{h \in \mathcal{H}} \Pr[h|X_{r+1} = x, \boldsymbol{o}_r]^2\right)\right] - \sum_{h \in \mathcal{H}} \mu_r(h|\boldsymbol{o}_r)^2. \quad (1)$$

The first component $\Pr[X_{r+1} = x|\boldsymbol{o}_r]$ is the probability of observing response $x \in \mathcal{X}_q$ conditional on the past observations $\boldsymbol{o}_r$, which is calculated by

$$\Pr[X_{r+1} = x|\boldsymbol{o}_r] = \sum_{h \in \mathcal{H}} \Pr[X_{r+1} = x|h] \mu_r(h|\boldsymbol{o}_r). \quad (2)$$

The second component $\Pr[h|X_{r+1} = x, \boldsymbol{o}_r]$ is the posterior belief of hypothesis $h \in \mathcal{H}$ conditional on the updated observations $((\boldsymbol{q}_r, q), (\boldsymbol{x}_r, x))$. It is calculated using the Bayes' rule:

$$\begin{aligned} \Pr[h|X_{r+1} = x, \boldsymbol{o}_r] &= \frac{\Pr[X_{r+1} = x|h, \boldsymbol{o}_r] \mu_r(h|\boldsymbol{o}_r)}{\sum_{h' \in \mathcal{H}} \Pr[X_{r+1} = x|h', \boldsymbol{o}_r] \mu_r(h'|\boldsymbol{o}_r)} \\ &= \frac{\Pr[X_{r+1} = x|h] \mu_r(h|\boldsymbol{o}_r)}{\sum_{h' \in \mathcal{H}} \Pr[X_{r+1} = x|h'] \mu_r(h'|\boldsymbol{o}_r)}. \end{aligned} \quad (3)$$

The last term, the sum of squared posteriors, is a constant term independent of q. We keep this term for completeness in presentation (see Appendix B for theoretical background), but we can ignore that term in practice.

One can interpret the EC^2 informativeness function as the expected reduction in *Gini impurity* following the observation of X_r. [11] The Gini impurity is commonly used in classification and regression tree (CART) machine learning applications. [12] It is defined by

$$I_{\mathrm{Gini}}(f) = \sum_{j \in J} f_j(1 - f_j) = 1 - \sum_{j \in J} f_j^2,$$

where J is the set of "labels" (or "classes") in the classification problem and f_j is the probability of label $j \in J$. [13]

[11] We thank Romann Weber for pointing out this relationship between the EC^2 informativeness function and Gini impurity.

[12] In a decision tree machine learning problem, the term *purity* refers to the quality of a predictive split within a node of the tree: A split that classifies observations perfectly has no "impurity"; a split which misclassifies is "impure".

[13] An *impurity function* is a function defined on a $(K-1)$-dimensional simplex $\{(f_1, \ldots, f_K) : f_k \geq 0, k = 1, \ldots K, \sum_{k=1}^{K} f_k = 1\}$ such that: (i) it is maximized only at $(1/K, \ldots, 1/K)$; (ii) it achieves its minimum at the vertices of the simplex (where all probability is placed on one hypothesis, $f_j = 1$ for some j); and (iii) it is a symmetric function (i.e., permutation of does not change the value of the function).

Then, $I_{\text{Gini}}(f)$ gives the expected rate of incorrect labeling if the classification was decided according to the label distribution f. Replacing the label set J with the hypothesis set $\mathcal{H}$ and the label distribution with the posterior belief μ_r, we obtain the equivalence between our EC^2 informativeness function and the expected reduction in Gini impurity:

$$\Delta_{\text{EC}^2}(q|\boldsymbol{o}_r) = I_{\text{Gini}}(\mu_r(\cdot|\boldsymbol{o}_r)) - \mathbf{E}[I_{\text{Gini}}(\mu_{r+1}(\cdot|\boldsymbol{o}_r, (q, x)))],$$

where the expectation in the second term is taken with respect to $\Pr[X_{r+1}|\boldsymbol{o}_r]$.

Computation of Δ_{EC^2}. The necessary ingredients for calculation of $\Delta_{\text{EC}^2}(q|\boldsymbol{o}_r)$ are (conditional) choice probabilities $\Pr[X_{r+1}|h]$ and the posterior beliefs $\mu_r(h|\boldsymbol{o}_r)$ for $h \in \mathcal{H}$.

The posterior beliefs $\mu_r(\cdot|\boldsymbol{o}_r)$ are calculated using Bayes' rule. The belief formation process starts with an initial prior μ_0. As new observations are accumulated, posterior beliefs are updated using equation (3) based on the actual response x_{r+1}. For the conditional choice probability $\Pr[X|h]$, we need to impose some behavioral assumption that maps hypothesized preference to observed choice, and preferably includes a reasonable type of noise in responses. In the current study, we mainly consider a stochastic choice model in the form of multinomial logit (also called as sotfmax choice model; in the context of CTB choices see Harrison et al., 2013; Janssens et al., 2017):

$$\Pr[X = x|h] = \frac{\exp(U_h(x)/\lambda)}{\sum_{x' \in \mathcal{X}} \exp(U_h(x')/\lambda)}, \tag{4}$$

where U_h is a parametrized utility function under hypothesis $h \in \mathcal{H}$. The *"temperature"* (or response sensitivity) parameter $\lambda \geq 0$ controls the sensitivity of choice probabilities to the underlying utility values. [14] The choice probability approaches to a uniform distribution as $\lambda \to \infty$ while it approaches to a degenerate probability distribution assigning all mass at the utility-maximizing option as $\lambda \to 0$.

In general, possible values of λ can be incorporated as part of the hypothesis space $\mathcal{H}$ to capture individual heterogeneity of noisiness or to distinguish optimally between different models

[14] Apesteguia and Ballester (2018) point out that the use of a random utility model (RUM) in estimation of risk and time preferences has an identification problem while the random parameter model (RPM) does not. Extension of the DOSE methods and systematic comparison between RUM and RPM are left for future works.

of noise (e.g., Bardsley et al., 2009; Wilcox, 2008). In the application of the DOSE methods in our online experiments, we set λ at an exogenously fixed value, since identifying the temperature parameter at the same time as identifying other core preference parameters have proved to be challenging in a simpler choice domains than ours (Chapman et al., 2018). We complement the analysis by running additional simulations which simultaneously estimate preference parameters and noise (Section 4).

3.3 Selecting the Most Informative Next Question

Given an informativeness function Δ_{EC^2}, a question is selected to be asked in round $r+1$ by

$$q_{r+1} \in \underset{q \in \mathcal{Q} \setminus \{q_1, \ldots, q_r\}}{\operatorname{argmax}} \Delta_{\mathrm{EC}^2}(q|\boldsymbol{o}_r). \tag{5}$$

In the rare case of multiple maximizers of $\Delta_{\mathrm{EC}^2}(q|\boldsymbol{o}_r)$— equally informative questions— the algorithm selects one randomly. Notice that the question selection rule (5) is *myopic*—we are not taking the effect of response x_{r+1} to the potential future question selection into account. We discuss this limiting feature briefly in the concluding Section 7.

3.4 Prior Beliefs

In order to initiate the adaptive question selection procedure, we have to specify a Bayesian prior μ_0 over hypotheses. The easiest way to specify a prior is to assume that each model class $m \in \mathcal{M}$ has equal probability, which is then spread uniformly across all hypothesis $h \in \mathcal{H}_m$ in that model class. A useful alternative is a data-driven prior which uses distributions of parameters obtained from existing studies. For example, Wang et al. (2010) suggest the following procedure. First, estimates of each parameter are binned into n equiprobable bins. Second, the midpoints of those bins are used as discrete mass points, each of which is assumed to have $1/n$ probabilities. One can also add "extreme" parameters to capture possibilities of outliers. Assuming that three parameters are independently distributed, we obtain the prior $\mu_0(h)$ by the product of the Bayesian priors over the parameters. After running experiments and obtain more data, we go back to the first point and refine our beliefs.

3.5 Comparison to Other Adaptive Design Approaches

DOSE is different than other existing methods such as Dynamic Experiments for Estimating Preferences (DEEP; Toubia et al., 2013) and Adaptive Design Optimization (ADO; Cavagnaro et al., 2010, 2011, 2013b,a, 2016; Myung et al., 2009, 2013). Essentially, the main difference across methodologies lies in the formulation of the informativeness function measuring the value of next questions, and some computational details.

In the DEEP method, the question that maximizes the expected norm of the Hessian of the posterior distribution at its mode, also called as the maximum a posteriori estimate (MAP estimate; DeGroot, 1970), is selected for next round. The authors used the absolute value of the determinant as the norm of the Hessian. This choice of informativeness function was motivated by the fact that the asymptotic covariance matrix of the maximum likelihood estimator (MLE) is equal to the inverse of the Hessian of the log-likelihood function at the MLE. In ADO method, on the other hand, the informativeness of a question is measured in terms of Shannon mutual information (Cover and Thomas, 1991).

In addition to the formulation of the informativeness function, there is another key difference that distinguishes those existing approaches and the one we take here—DOSE requires discretization of the parameter space while DEEP and ADO deal with continuous spaces. This feature can be a disadvantage of our methodology, but at the same time it is inevitable given that the space of choice alternatives in our linear budget environment is much larger than simple binary choice environment in those previous studies. Comparing DOSE against DEEP and ADO is beyond the scope of the current study is left for future work.

The closest paper to ours is Chapman et al. (2018) (a revision of Wang et al., 2010). They use a Kullback-Liebler divergence and emphasize surprising results about loss-aversion from a survey, whereas we apply the EC^2 method to Convex Time Budget protocol (a method which could be easily extended to other budget line experiments for risk and social preferences, such as: Choi et al., 2007; Fisman et al., 2007).

4 Simulation Exercises

To evaluate the performance of our adaptive design approach, we conduct several simulation exercises (called "ground truth" in computer science). In a CTB experiment, every round a subject is asked to allocate experimental budget between two time periods t and $t+k$. Date t is called the "sooner" payment date and $t+k$ is called the "later" payment date; the gap between them is the delay length k. Choices were made by allocating 100 tokens between two payment dates. There are token exchange rates (a_t, a_{t+k}) that convert tokens to money. The slope of the budget line is thus determined by the gross interest rate over k periods, $1+\rho = a_{t+k}/a_t$. By choosing sets of (t, k, a_t, a_{t+k}), the researcher can identify preference parameters both at the aggregate and the individual level.

Let $\mathbf{D} = (\mathbf{D}(t), \mathbf{D}(k), \mathbf{D}(a_t), \mathbf{D}(a_{t+k}))$ denote the *design space*, a collection of vectors specifying the spaces of parameters. For example, $\mathbf{D}(t) = (0, 7, 35)$ and $\mathbf{D}(a_{t+k}) = (0.20, 0.25)$ specify an environment in which sooner payment dates can be today, a week from today, or five weeks from today, and one token allocated to sooner or later dates is worth \$0.20 and \$0.25, respectively. The set of questions $\mathcal{Q}$ is all possible combinations of the numbers in the vectors $\mathbf{D}$, denoted $\mathcal{Q}(\mathbf{D})$.

In the following simulation exercises, as a first step, our interest is in parameter estimation within one fixed model class. Since we focus the CTB method applied to estimate parameters in quasi-hyperbolic discounting (QHD) model (Laibson, 1997; Phelps and Pollak, 1968). [15] Assuming a QHD with power utility function, a consumption (c_t, c_{t+k}) is evaluated (at time 0) as:

$$U(c_t, c_{t+k}) = \frac{1}{\alpha}c_t^{\alpha} + \beta^{\mathbf{1}\{t=0\}}\delta^k \frac{1}{\alpha}c_{t+k}^{\alpha}, \tag{6}$$

where δ is the per-period discount factor, β is the present bias, and α is the utility curvature parameter.

We report results from several model recovery exercises (also known as a "ground truth" analysis) below. Each simulation assumes a "true" underlying preference $h^0 \in \mathcal{H}$ and generates

[15] See Andreoni and Sprenger (2012), Andreoni et al. (2015), Augenblick et al. (2015), Balakrishnan et al. (2015), Bousquet (2016), Brocas et al. (2016), Janssens et al. (2017), Kuhn et al. (2017), Sawada and Kuroishi (2015), Sun and Potters (2016).

choices according to that model. Questions are prepared either by an adaptive procedure or by a random selection from $\mathcal{Q}$ (without replacement). We measured how fast and precisely the adaptive design can recover the true model.

4.1 Simulation Parameters

We construct two hypothesis spaces using the data-driven approach (see Appendix D for details). The first hypothesis space $\mathbf{H}_1$ (top panel in Table 2) focuses on QHD parameters (α, δ, β) while fixing the softmax temperature $\lambda \in \{0.04, 0.18\}$ to check the effects of noisiness in stochastic choices. There are 175 unique combinations of parameters in this hypothesis space. The second hypothesis space $\mathbf{H}_2$ (bottom panel in Table 2) is used to study performance when jointly estimating preference parameters and the noise parameter. There are 135 unique combinations of parameters in this hypothesis space.

We combine these hypothesis spaces with a common question design space:

$$\mathbf{D} = \begin{bmatrix} t: & (0, 7, 28) \\ k: & (21, 35, 42, 56) \\ a_t: & (0.14, 0.15, 0.16, 0.17, 0.18) \\ a_{t+k}: & (0.17, 0.18, 0.19, 0.20, 0.21) \end{bmatrix}.$$

The numbers in this design space are chosen so that the reward magnitudes are comparable to those in Andreoni and Sprenger's (2012), Andreoni et al. (2015), and Augenblick et al. (2015), upon which our data-driven priors are based. The total number of questions in $\mathcal{Q}$ is 300.

4.2 Procedure

Every simulation consists of $|\mathcal{H}|$ "subsimulations," in which: (i) One hypothesis $h \in \mathcal{H}$ is fixed as the "true model"; (ii) 45 questions are generated by three selection rules: EC^2, "Fixed," and "Random"; and (iii) Choices are generated with stochastic choice model (4) together with assumed parameter values h. [16] We repeat this procedure 100 times for each h.

[16] We simulate choices following a procedure described in Meier and Sprenger (2015). The stochastic choice (4) gives a cumulative distribution function $F(x) = \sum_{y \in \{0,\ldots,100\}}^{x} \Pr[X = y|h]$. We then draw a number ξ from a

Table 2: Data-driven prior: parameter values and their initial probabilities. Top panel: $\mathbf{H}_1$ focusing on QHD parameters. Bottom panel: $\mathbf{H}_2$ including noise parameter.

$\mathbf{H}_1$	α	0.7675	0.9016	0.9519	0.9719	0.9833		
	$\mu_0(\alpha)$	0.2	0.2	0.2	0.2	0.2		
	δ	0.9945	0.9972	0.9986	0.9992	1.0012		
	$\mu_0(\delta)$	0.2	0.2	0.2	0.2	0.2		
	β	0.8839	0.9401	0.9786	1.0000	1.0233	1.0623	1.1237
	$\mu_0(\beta)$	0.1	0.1	0.2	0.2	0.2	0.1	0.1
$\mathbf{H}_2$	α	0.7946	0.9480	0.9827				
	$\mu_0(\alpha)$	0.33	0.33	0.33				
	δ	0.9950	0.9982	1.0011				
	$\mu_0(\delta)$	0.33	0.33	0.33				
	β	0.9010	0.9786	1.0000	1.0233	1.1079		
	$\mu_0(\beta)$	0.2	0.2	0.2	0.2	0.2		
	λ	0.72	1.08	1.8				
	$\mu_0(\lambda)$	0.33	0.33	0.33				

The Random rule selects questions purely randomly (without replacement) from $\mathcal{Q}$. The Fixed rule pre-specifies an order of 45 questions to capture how a CTB design is most often used in existing studies. A typical CTB design "blocks" questions based on a time frame (t, k), so that subjects complete several questions under the same time frame before moving to a new time frame. Within each time frame, subjects often see questions that are ordered by the gross interest rates (see Table F.1 in Appendix F).

We note, however, that in many CTB experiments subjects see several questions presented simultaneously on the same sheet of paper or on the computer screen. Therefore, the order at

uniform distribution on $[0, 1]$. We assign a choice x^* if $F(x^* - 1) \leq \xi < F(x^*)$ with $F(-1) = 0$.

which subjects answer questions may not necessarily coincide with the order of presentation, which typically has monotonic structure as described above. Even with this caveat in mind, the "monotonic" Fixed rule will be a useful benchmark to compare against EC^2 algorithm.

Some technical details about implementation of EC^2 algorithm is presented in Appendix C.

4.3 Results

The primary variables of interests are: (i) speed of underlying parameter recovery, (ii) frequency of correct parameter recovery, and the effects of noise level in choices and reward magnitudes.

Now, we introduce some notation that becomes useful later. Suppose we run S *iterations* of R questions under true model h^0. Each iteration consists of the following steps: Let $\mu_r^s(h|h^0)$ denote the posterior belief of a hypothesis h in round r of iteration s, when h^0 is the true model; let $\bar{\mu}_r(h^0|h^0) = \sum_{s=1}^{S} \mu_r^s(h^0|h^0)/S$ denote the posterior belief of the true model averaged over all iterations; let $h_s^{\text{MAP}} = \text{argmax} \sum_{r=R-n+1}^{R} \mu_r^s(h|h^o)/n$ denote the *maximum a posteriori (MAP) estimate* given by average beliefs of the last n rounds in iteration s; and let $\text{hit}_s(h^0) = \mathbf{1}\{h^{\text{MAP}} = h^0\} \in \{0, 1\}$ be an indicator for whether the MAP hypothesis matches the true model in iteration s. Finally, the *hit rate* is given by $\sum_{s=1}^{S} \text{hit}_s(h^0)/S$.

Accuracy of parameter recovery. The EC^2 algorithm recovers the underlying preference parameters *more accurately*, and *more quickly*, compared to two benchmark cases. Figure 1 compares hit rates of EC^2 and Fixed, using the MAP estimates given by the average posteriors from the final five rounds. Columns A to C in each row represent the same information, but are color-coded based on the parameter values of the underlying hypotheses. Since we take hit rates from EC^2 algorithm on the y-axis, data points appearing above the 45-degree line, as in this figure, indicate that EC^2 algorithm is more accurate (at the end of the simulation), compared to Fixed question design. We also find better performance of EC^2 compared to Random, but no visible difference between Fixed and Random (Figures E.1 and E.2 in Appendix E).

Comparing distributions of hit rates between panels (within each row) or between rows further reveals the following. First, whether or not the algorithm can achieve higher performance

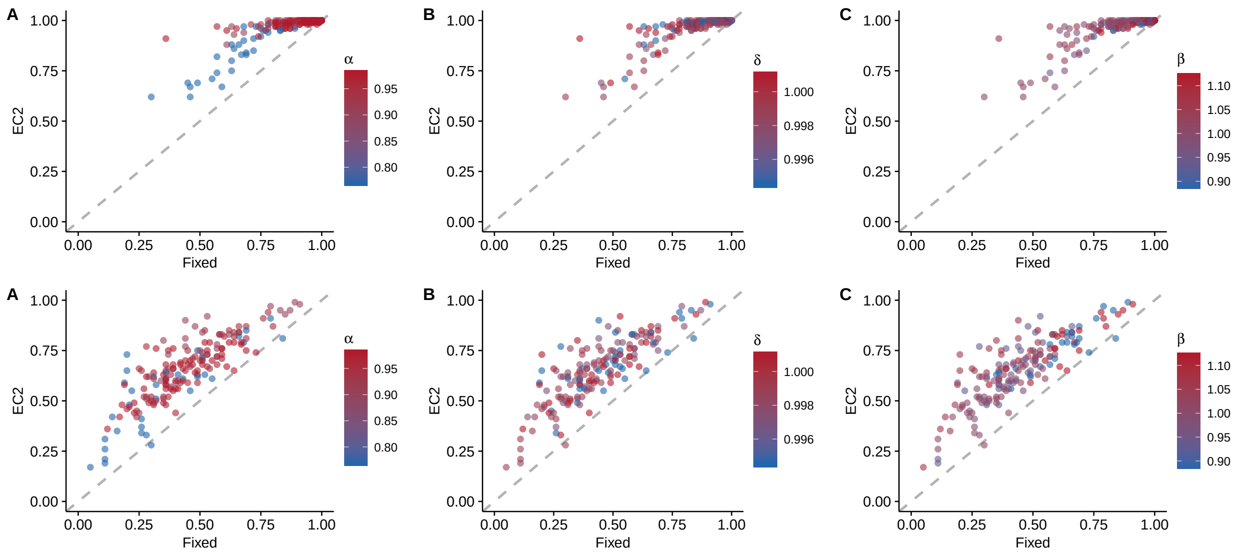

FIGURE 1: "Hit rates" comparison between EC^2 and Fixed. The first row is from simulation with $\mathbf{H}_1$ and $\lambda = 0.04$, and the second row is with $\mathbf{H}_1$ and $\lambda = 0.18$. Each panel is color-coded by the parameter value, and it shows fundamental difficulty in recovering smaller α's.

depends on the underlying parameter values, especially α (see panel A in each row of the figure). Regardless of the algorithm, there is a fundamental difficulty in accurately recovering utility functions which are "sufficiently concave." The other two parameters, on the other hand, do not have such clear effects on accuracy.

Second, as expected, noisier choices reduce overall performance of the algorithms (compare the first and second rows). Even the EC^2 algorithm sometimes produce hit rates less than 0.5.

Speed of parameter recovery. Next, we show that the EC^2 algorithm works faster than the benchmarks. Figure 2 shows the time series of posterior standard deviation of each parameter (α in panel A, δ in panel B, and β in panel C). The solid lines represent the dynamics of the median of posterior standard deviations after round r response, across all simulations $s = 1, \ldots, S$ and all true underlying models $h^0 \in \mathcal{H}$. The shaded bands represent inter-quartile range at each point in time.

We observe: (i) Both EC^2 and Random algorithms reduce a lot of uncertainty by the 10-th question; (ii) All three question selection rules perform comparably in identification of α; (iii)

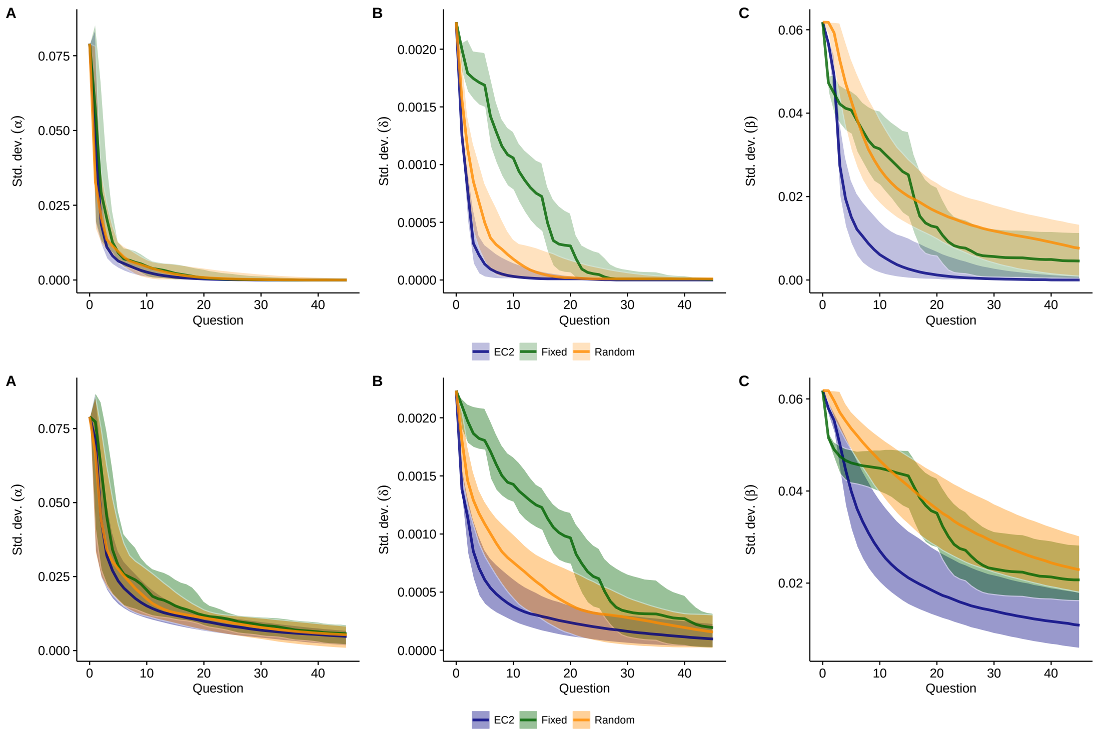

FIGURE 2: Posterior standard deviation over time. The first row is from simulation with $\mathbf{H}_1$ and $\lambda = 0.04$, and the second row is with $\mathbf{H}_1$ and $\lambda = 0.18$. The 25-percentile, median, 75-percentile for each algorithm are presented.

Fixed design performs worst, especially in identification of δ (the lines look like "steps" because time frames change after every five questions); (iv) higher degree of noise in choices stretch the inter-quartile range of standard deviations (compare the bottom row, with higher noise, to the top). Overall, the figure shows that EC^2 is faster than the benchmarks, especially the fixed design.

The dynamics of posteriors over the true model is another measure of speed with which we can compare different question selection rules. Figure E.3 in Appendix E presents $\bar{\mu}_r(h^0|h^0)$, $r = 1, \ldots, 45$, for several combinations of (α, δ, β). The EC^2 algorithm always gives higher posterior beliefs compared to other two benchmarks, but the speed of updating and the final level of the posterior depend crucially on the underlying true model h^0. For example, it suffers to identify parameters when utility function has large curvature ($\alpha = 0.7675$; top right panel in Figure E.3 in Appendix E).

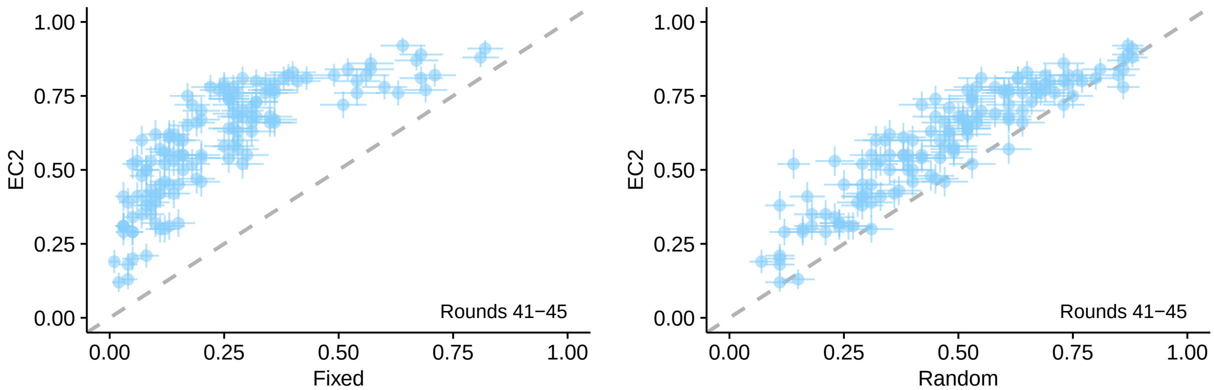

FIGURE 3: "Hit rates" comparison between EC^2 and Fixed (left) and EC^2 and Random (right).

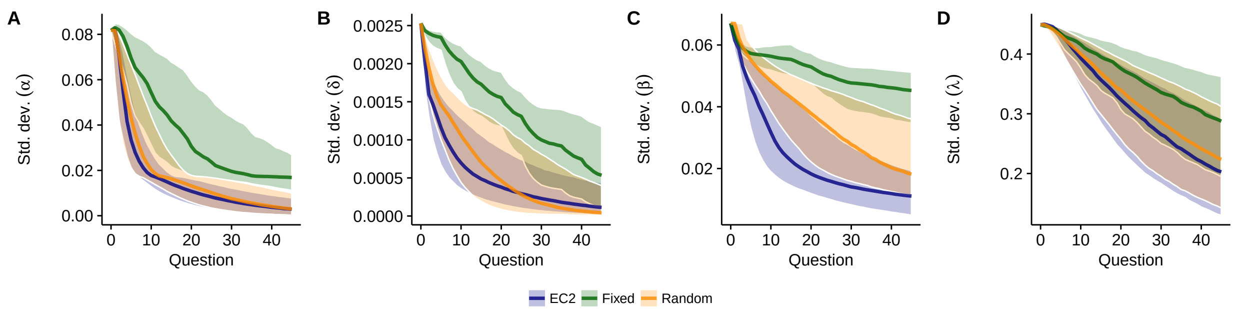

FIGURE 4: Posterior standard deviation over time.

Joint estimation of preference parameters and noise. Participants in real experiments and surveys are expected to exhibit heterogeneity not only in the core preference parameters but also in their tendency to make mistakes. The second set of simulations with hypothesis space $\mathbf{H}_2$, which includes softmax temperature λ, will assess the performance of the DOSE method in such an environment.

Figures 3 and 4 demonstrate that EC^2 achieves higher hit rates and removes uncertainty faster compared to both Fixed and Random designs. These observations validate the relative performance advantage of the DOSE method over standard designs even in the presence of heterogeneous noise.

Joint estimation of preference parameters and noise itself is not a major challenge in the DOSE method. However, in applying the method, the researcher has to trade-off benefit and cost of allowing heterogenous noise. As we see in Figure 4D, even EC^2 requires many questions to identify the true λ, which also affects the speed of identifying other preference parameters.

Computation speed of EC^2. In order for the DOSE method to be a useful adaptive design algorithm, it has to calculate the informational value of questions and present the next question to the subjects instantly. Under the sizes of the design space and the hypothesis space used in the simulations (300 and 175, respectively), it takes less than 10 seconds to initialize the set of all questions $\mathcal{Q}$ and hypotheses $\mathcal{H}$, and takes about 50 to 70 milliseconds to prepare the next question between each round. Therefore, in experiments of this size the subjects will not have to wait unnaturally long between questions.

5 Experimental Design

Simulation exercises presented in the previous section establish the power of our application of the adaptive question selection mechanism. We now examine usefulness of this new design in empirical applications, using online experiments.

5.1 Design and Implementation

The experiment was conducted using Amazon's Mechanical Turk platform (hereafter AMT). The platform has become popular in any domains of experimental social sciences. Detailed explanations are presented in Goodman et al. (2013), Horton et al. (2011), Mason and Suri (2012), and Paolacci et al. (2010).

We conducted experiments with *hypothetical* choices. One may argue that hypothetical choice tasks conducted on AMT would deliver quite different results from incentivized laboratory experiments. However, available evidence show that this is not the case—time preference estimates from Montiel Olea and Strzalecki (2014), Ericson and Noor (2015), and Hardisty et al. (2013) are all comparable to what is usually observed in incentivized experiments. Other studies, such as Bickel et al. (2009), Johnson and Bickel (2002), Madden et al. (2003, 2004), and Ubfal (2016), also found no effects of incentives.

We used the hypothesis space $\mathbf{H} = \mathbf{H}_1$ used in the first set of simulations presented above and design space

$$\mathbf{D} = \begin{bmatrix} t: & (0, 14, 28) \\ k: & (14, 21, 28, 35) \\ a_t: & (0.91, 0.94, 0.97, 1.00, 1.03) \\ a_{t+k}: & (1.00, 1.03, 1.06, 1.09, 1.12) \end{bmatrix}.$$

We set the number of questions to 20. For the Fixed rule, we use the sequence of questions presented in Table F.2 in Appendix F. In this version of the experiment, the temperature parameter λ is fixed at the same level (0.18) across treatments and subjects. We plan to address this noise issue in future research.

We conducted three treatments, EC^2, Fixed, and Random, with 45 workers in each. Each worker received a $3 participation fee after completing all 20 questions and an exit survey. Since the entire experiment took about 15 to 20 minutes, the hourly wages for those workers were around $10, which is quite high by AMT standards.

5.2 Results

Preliminary data analysis. Out of the $N = 135$ AMT workers participated in our study, two dropped out in the middle of the experiment, four had no variability in their allocation decisions, and five exhibited a strong "anchoring effect" identified by significant linear correlation between randomly selected initial slider locations and final allocation decisions (t-test on Pearson's correlation coefficient, $p < 0.001$). We exclude those 11 subjects from our data analysis.

One disadvantage of AMT is that we cannot monitor workers while they are performing the task and check if they are paying attention to it. However, the response time data reveal that the subjects in our experiments did not inattentively click "proceed" button to finish the task as quickly as possible: The median subject spent about five to six minutes in reading the instructions, about seven to nine seconds in each CTB question, and about four to five minutes in completing all 20 CTB questions.

Preference parameter estimates. The first data are the estimated values, and precision, of the preference parameters (α, δ, and β). We first show values computed from the subject-specific Bayesian posterior distributions in the last four trials.

Consider β values (present-biasedness) as a specific example. For each subject and round r the EC^2 procedure derives posterior probabilities of the seven discretized β values in the hypothesis space $\mathbf{H}$. The mean of the posterior distribution and its standard deviation represent the subject's estimate and accompanying precision. For *pairs* of parameters these data can be represented in a scatter plot. Plotting the pairs gives evidence about whether there is any correlation, across subjects, between parameter values (e.g., do those who are present-biased, as evidenced by low β, discount the future more or less, as evidenced by the value of δ?).

Figure 5 plots parameter estimates (averaged over last four trials) along with bivariate confidence bars. Figure 6 plots empirical CDFs of these estimates. Several features of the estimates can be seen in these Figures:

1. Many estimates, particularly curvature α and discount factor δ, are on the boundary near the maximum or minimum of support of the data-driven priors. This is generally a sign that the Bayesian priors need to be stretched out further to better fit subjects who have unusually high or low parameter values. Keep in mind that after the data are collected, they can be reanalyzed using any Bayesian priors. The particular data-driven priors that we used only constrained the sequence of budget lines that each subject faced.

2. Most estimates of δ (95.5%), and most estimates of β (66.2%) are below one. The percentages for subjects who have posteriors larger than 0.9 on values $\delta < 1$ and $\beta < 1$ are 94.7% and 35.2%, respectively.

3. Pairs of parameters are not very correlated across subjects. While the procedure is not optimized to estimate cross-parameter correlation, these data suggest that the constructs are rather separate.

4. The values we estimate are comparable to those in previous CTB experiments. Recall that the method starts with a data-driven prior constructed from estimates from Andreoni and

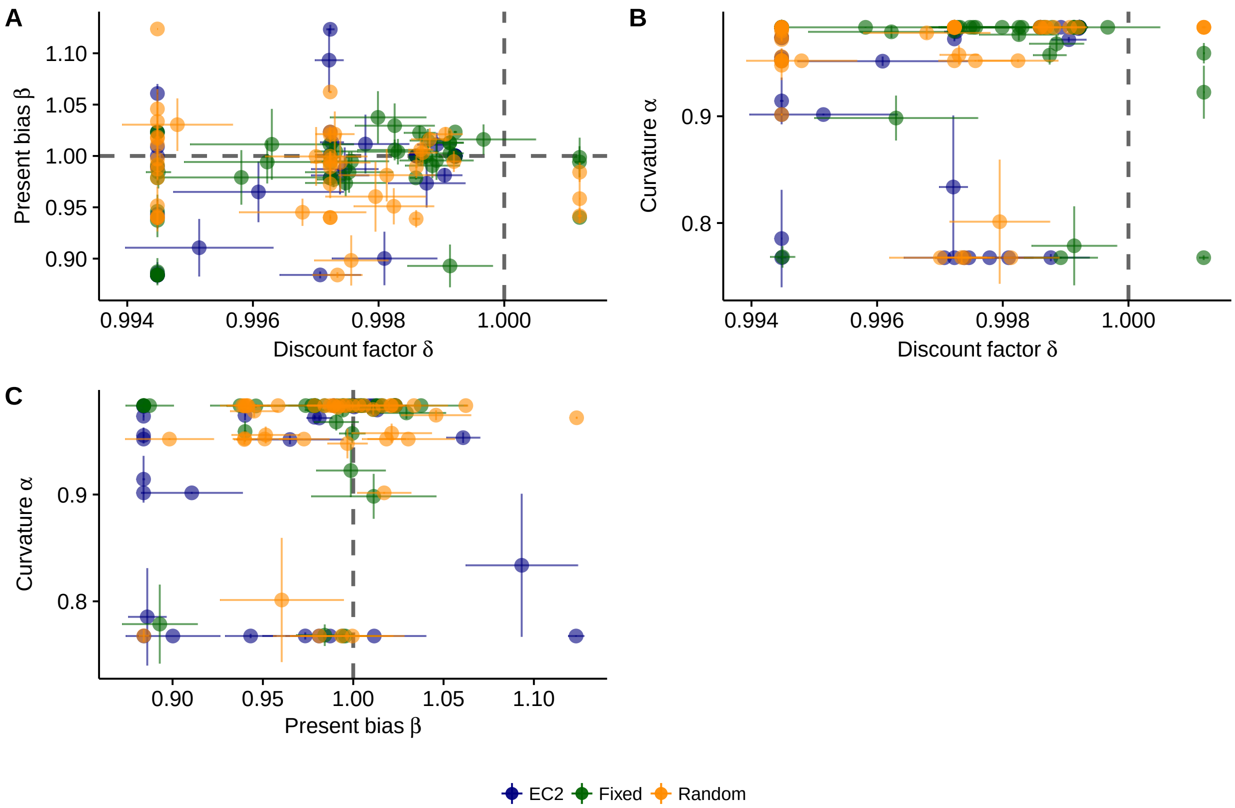

FIGURE 5: Scatterplots of estimated parameters. Each dot represents single subject's mean and lines represent standard deviation, both from posterior belief averaged across the last four questions.

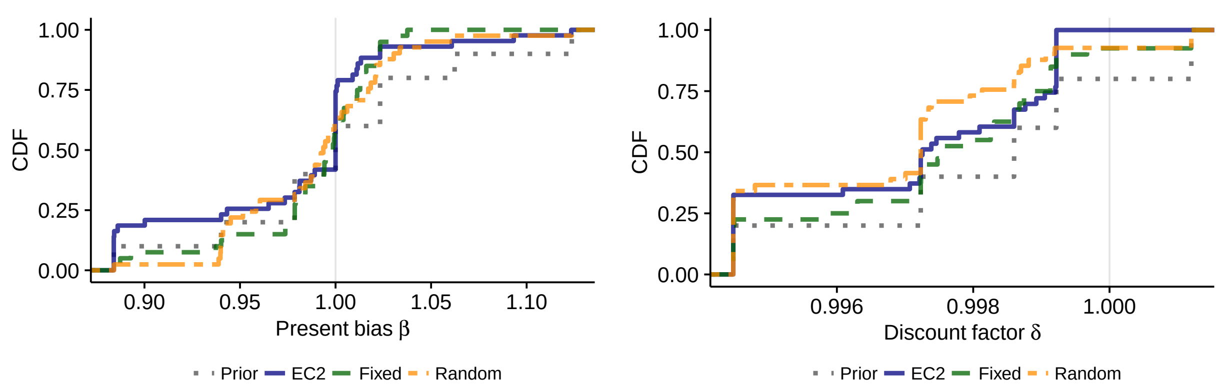

FIGURE 6: Empirical CDFs of estimated parameters.

Sprenger (2012), Andreoni et al. (2015), and Augenblick et al. (2015). The distributions of the means of posterior parameter distributions do not move much away from the prior means.

Our median β estimate is close to 1, which is higher than some of the recent studies finding significant present bias (e.g., Balakrishnan et al., 2015; Bousquet, 2016; Sun and Potters, 2016). This result could be due to the hypothetical procedure, which may not generate as strong a desire for the immediate reward as would be present with actual payments.

5. The distribution of the final parameter estimates are indistinguishable across treatments (Figure 6). It suggests that presenting a sequence of questions in an adaptive design fashion does not induce biases in estimated final parameters (which is a requirement for an adaptive method to be practically useful).

Speed of convergence. Next we present some statistics illustrating how rapidly the different sequencing methods achieve precision. Figure 7 shows "survival" curves (based on the actual choices, and averaged across subjects). These curves count how many hypothesized parameter configurations have posterior probability above a particular cutoff (in these figures, the cutoff is 0.01). A good method will reduce the set of surviving hypotheses rapidly, which will be evident visually as a steeply plunging curve. For example, the procedure starts with 175 different three-parameter (α, δ, β) hypotheses, each with prior probability of either 0.004 or 0.008 (see Table 2). After five questions, on average 15, 28, and 28 hypotheses survive using EC^2, Fixed, and Random procedures (panel A of Figure 7). The results for 35 different two-parameter (δ, β) hypotheses are similar, although the advantage of EC^2 is a bit less pronounced (panel B of Figure 7). Another way to measure the advantage is to fix the number of surviving hypotheses after five questions, and compute how many questions are needed, using Fixed or Random sequences, to achieve the number of surviving hypotheses. The answers are 10 in both cases. So regardless of how the speedup advantage is measured, the EC^2 procedure is about twice as good. [17]

Another measure of quality is how precisely parameters are estimated partway through an experiment. To illustrate, we compare parameter estimates (i) after 10 questions had been asked

[17] Note also that the Fixed method is slightly inferior to Random. Intuitively, in Fixed-sequence designs the design may get stuck using questions which are not providing information which is useful for estimating parameters. Because the design is fixed it persistently asks "uninteresting" questions. The Random design does not get stuck in this way.

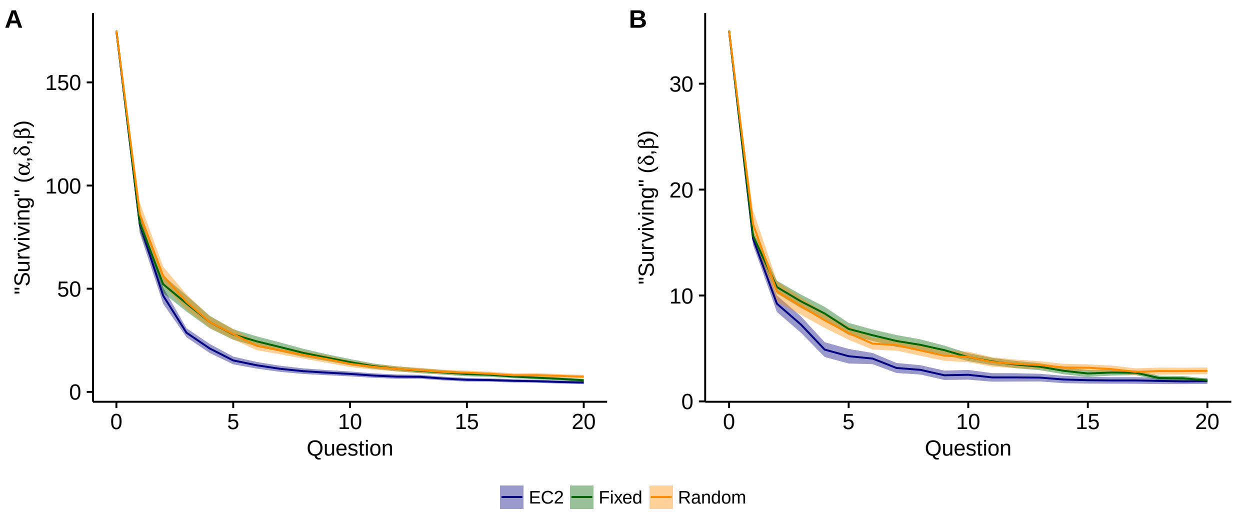

FIGURE 7: Speed of achieving parameter precision. (A) Survival curves for three-parameter (α, δ, β) hypotheses for the three methods EC^2 (purple), Fixed sequence (green) and Random (orange). (B) Survival curves for two-parameter (δ, β) hypotheses.

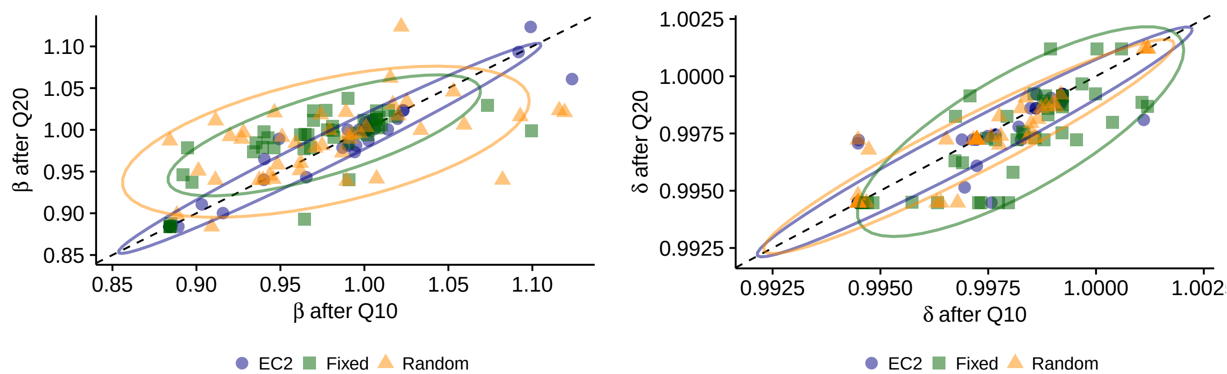

FIGURE 8: Parameter estimates after 10-th question and the final question. 95% confidence ellipses are overlaid.

and (ii) after all 20 questions had been asked (Figure 8). Estimated β in EC^2 treatment is more narrowly scattered around the 45-degree line compared to other two treatments, indicating that EC^2 learned (and became confident) true underlying parameter and did not revise this estimate much afterwards. The performance of EC^2 and Random are similar in estimating δ.

Finally, it is notable that nearly half the responses are choices of either 0 or 100 tokens allocated to the later reward date. The high frequency of these extreme "corner" allocations has been observed in many studies using CTB.

Corner choices are not unreasonable. But suppose a person is consistently choosing, say, 100 token allocations to the later reward, and 0 to the sooner reward, for several different budget lines in a row. Such a pattern of persistent choices of 100 implies that the budget lines which were chosen are not efficiently determining the *strength of preference* for allocations to the earlier reward. An efficient method would quickly locate a budget line for which some tokens are allocated to the sooner reward.

More generally, in a good method allocations should be *negatively autocorrelated* across trials (e.g., subjects who are choosing corners should flip back and forth between allocating 0 and 100 on consecutive trials quite often). As an illustration, Figure 9 take three representative subjects from the EC^2, Fixed, and Random conditions (from top to bottom) and plots the dynamics of sooner allocation percentages (panel A) and a scatterplot between "% sooner in question $r+1$" and "% sooner in question r" (panel B). The subject in the Fixed condition changed his/her sooner allocation monotonically, which makes sense by design of the sequence (asking four questions in the same time frame, from low gross interest rate to high, and then move on to next five with different time frame). The subject in the Random condition chose corners frequently, but s/he sometimes stuck to one corner (between questions 11 and 16, for example). Unlike those two, the subject from EC^2 condition flipped back and forth between two corners with high frequency—s/he never stopped at one corner for more than three questions in a row.

Figure 10 generalizes this idea and plots the cumulative distribution functions of consecutive-trial autocorrelations for the three sequencing methods, across subjects (where a separate autocorrelation is computed for each subject). The fixed sequence generates hardly any significantly negative autocorrelations. For the EC^2 method most autocorrelations ($30/43 = 0.70$) are negative with a mass around -0.50. Even though we cannot reject the null hypothesis of equal distribution between EC^2 and Random using the two-sample Kolmogorov-Smirnov test ($p = 0.23$), there is a qualitative difference between those two distributions. Among the 30 subjects who have negative autocorrelation, 11 (37%) of those values are significant at 5% level (t-test) in EC^2. In Random, on

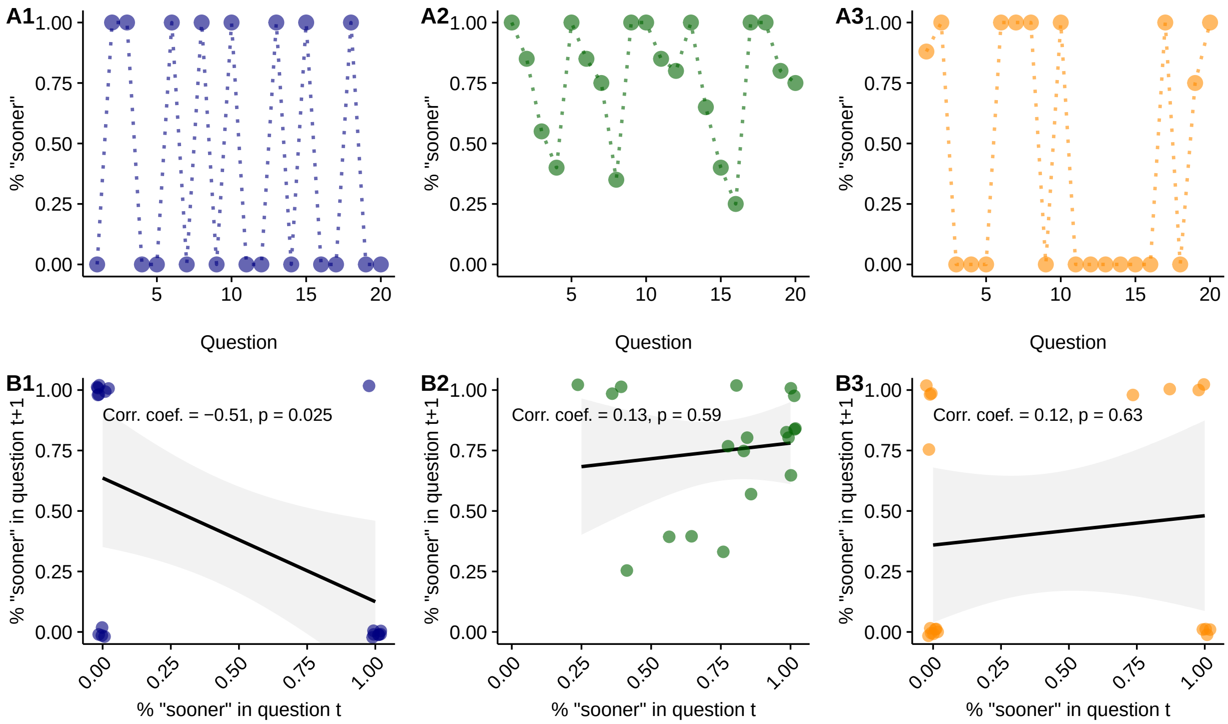

Figure 9: (A) Dynamics of percentages of tokens allocated to sooner payment. (B) Lagged scatterplot between sooner allocation percentages between two consecutive time periods. From left to right: EC^2, Fixed, Random.

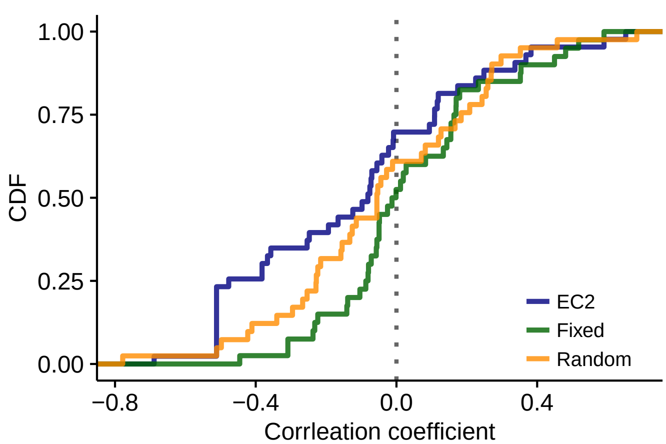

Figure 10: Empirical CDFs of correlation coefficient.

the other hand, there are only three significantly negative autocorrelations out of 25.

6 Possible Strategic Manipulation

Experimental economists have found it prudent to treat our subjects as (possibly) intelligent enough to think very carefully about how they should behave in an experiment, in order to earn more money.

This concern for how conniving subjects might be, while perhaps a bit paranoid, can help to expose weaknesses in design that could jeopardize inference, and which are often easily repaired. (It is like worrying in advance about black hat cyberattacks when designing cyber security.)

In the case of adaptive experimental design, the obvious concern is that subjects could 'game' or strategize by making choices in early trials which increase the quality of choices that are available to them in future trials.

In adaptive designs, subjects are likely to misrepresent their true preferences in some choices *if* all of the following chain of conditions hold: (i) they believe that future test choices depend on previous responses; (ii) they can compute *how* to misrepresent preferences in early choices to create better future choices (as evaluated by their own preferences); and (iii) the value of misrepresentation is high enough to be worthwhile. We present arguments and evidence that misrepresentation resulting from the chain of conditions (i)-(iii) is likely to be rare. And if misrepresentation does occur, it could be easily detected and is not likely to lead to wrong conclusions about revealed preferences which cannot be undone.

- Does strategizing pay? To partially answer this question, it is helpful to establish an upper bound on the maximum gain from strategizing, for a particular design and player type. The upper bound on the marginal gain is likely to be low. Here's why: In later periods, it does not pay to strategize since doing makes suboptimal immediate choices. And in early periods, strategizing is immediately costly for the same reason. So there is a natural tradeoff between the cost of strategizing in a period—the utility losses from deliberately making the wrong choices—and the future gains from improved choice sets. It could be that in a 10-period experiment, for example, strategizing is only beneficial in the first three periods. If so, the posterior probabilities computed after 10 periods might be close to the correct posteriors because they include 7 periods of non-strategizing choice data after three periods of misleading data. It is also possible that when ranking different subjects by their risk-aversion, for example, we can recover an approximately correct ranking across subjects even if manipulation leads to biased estimates of their means.

- Can strategizing be detected? Strategizing will typically leave clear fingerprints in the data from choices across a sequence of questions. In typical cases without strategizing, the posterior probability of the most likely hypothesis—as judged from final round result– goes up across the trials. In contrast, a strategizing respondent will appear to be one hypothesis type in early trials, and then revert to their true type in later trials (as the future benefit of strategizing shrinks). This will leave a telltale pattern of posterior probabilities veering from one type to another, from earlier to later trials. This is not evident in our data.

- How can strategizing be limited? There are several possible ways strategizing could be limited, presuming one budget line will be chosen at random at the end of the experiment as a basis for actual payment (the norm in experimental economics). The best remedy is ingenious and simple: Choose randomly from the *entire design space* of possible budget lines. [18] *Do not* choose from the set of lines that were presented. (Note that if the chosen budget line is one that was not presented during the adaptive question selection, the subject has to make a fresh choice.) The key to this method is that strategizing does not pay because it does not improve the quality of the budget lines that will be used to eventually determine the payment. Each of the entire set of budget lines is equally likely, regardless of what the subject chooses. (The only flaw in this method is that it lowers the probability that any of the actual choices that are made during the experiment will determine actual payment.) An alternative method is to tell the subjects that all their choices will be used to estimate their preferences, and the estimated preferences will be used to choose an allocation from

[18]This idea was suggested by Cathleen Johnson. The Prince (acronym summarizing principles that define the method: priority, instructions to experimenter, concreteness, entirety) method of Johnson et al. (2015), begins with a *real choice situation (RCS)*, which is randomly selected from a set of all possible questions. The RCS is written on a sheet of paper and put in a sealed envelope. The experimenter asks subjects to give "instructions" about the real choice to be implemented. At the end of the experiment, the experimenter opens the envelope and selects the subject's desired option using the instruction provided by the subject.

a different budget line (Krajbich et al., 2017). [19] In this method, the subjects are essentially "training" an algorithm, much as choices of Amazon books are training a recommender system.

7 Conclusion

In this paper we described and applied a method, called DOSE, for choosing an informationally-optimal sequence of questions in experiments. This method should be useful to the many economics experimenters who are currently using those methods in lab and field experimenters, and in surveys, and would value halving the time it takes to estimate parameters.

The first empirical finding is that the distributions of estimated β, δ, and α parameter are similar to those observed earlier. The second, novel, finding is that the EC^2 method is able to estimate parameters much more precisely during the middle part of an experiment– about twice as fast.

If one accepts the value of optimal adaptive design, there is a lot of interesting work to do. Here is a short to-do list:

1. *Other choice domains*: There are many areas of behavioral economics in which multiple theories or parametric frameworks are used to explain the same stylized facts. As noted in the introduction, adaptive optimal design is one way to make progress when there are multiple well-specified theories, and some intuitions (or evidence, as implemented here) about a prior probability distribution of parameters. These methods could be applied to distinguish theories about: Risky choice; social preferences and fairness; non-equilibrium

[19] In their application of DOSE method in a binary-choice risk preference elicitation task, subjects were told that: (i) subjects' responses during the task were hypothetical and would not count for the final payment; (ii) those hypothetical choices would be used to determine their risk preferences; (iii) a new question that had not been asked during the task would be drawn at random, and a computer algorithm would make a choice for the subject based on the hypothetical answers. Since every decision made during the task would influence how the computer algorithm would decide in a new question that determines the payment, the proposed mechanism would mute the subjects' incentive to misreport.

choices; and learning in games. [20]

2. *Multiple (non myopic) question selection*: Our implementation chooses one question at a time. It is possible that choosing sequences of two or more questions would be a substantial improvement, at the cost of more computation. For example, many people have an intuition that when choosing questions to estimate β and δ, say, it could be better to use a two-stage procedure like the following: Choose questions to estimate δ first (by imposing a front-end delay for the earlier reward, so all valuations depend on β), then transition to estimate β in the second stage. The myopic implementation cannot do this automatically because it cannot select a "package" of multiple questions to capture sequential information complementarities. That is, a δ-focussed question in trial 4 might be informationally valuable only if it is followed by two more δ-focussed questions. Our myopic procedure will not include this complementarity. However, the method can be easily adapted to see if selecting sequences of trials non-myopically is a large improvement.

3. *Handling noise*: There are two ways to handle noise in the DOSE framework. In the first approach, the researcher can jointly estimate preference parameters and the stochastic choice parameter by preparing a rich enough hypothesis space. The downside of this approach is that the algorithm may need to ask many questions to be confident. In the second approach, the researcher can ask small set of questions just to learn each subject's tendency to make mistakes, and then plug this subject-specific noise parameter in the DOSE algorithm to ask main questions. Designing and comparing these two approaches are left for future works.

4. *Optimal stopping*: Part of experimental design is when to stop asking questions. It is easy to compute an optimal stopping rule in theory: Quit asking questions when the marginal cost begins to exceed the expected marginal information benefit (or some loss function summarizing the expected possible benefits of learning more). However, in practice these cost and benefit numbers are not always easy to compute.

[20]These methods could also be applied to identify individual specific "boundaries" of context effects, such as compromise and asymmetrically dominated effects (Huber et al., 1982; Simonson, 1989). The method would allow researchers (and marketers) to quickly identify the best placement of decoy options.

5. *Using non-choice data*: The procedure uses only observed choices. In our experiments, however, we also observed response times (RTs) and the positions of a slider bar over time. These non-choice data could contain information that would help diagnose what theories describe behavior and what parameter values are. One potential example exploits the common correlation between how close in value two choices are, and how long a decision takes. Typically, "difficult" decisions—when objects are close in value—are slower and have longer RTs (see Clithero, 2016). Suppose there are two hypotheses about possible behavior, C and F. Also suppose that for a particular budget line under hypothesis C the allocations on the line are close in value and under the other hypothesis F they are far apart in value. A slow RT is more consistent with hypothesis C than with hypothesis F, and could be used to update probabilities much as observed choices are. [21]

Finally, we think optimally adaptive design is relevant to the recent growth of interest in scientific reproducibility (to which we have also contributed; see Camerer et al., 2016, 2018). Concern about reproducibility is partly about weak statistical power, partly about publication bias and snowballing of attention to weak results, and partly about incentives of career-concerned scientists, journal editors and referees, funding agencies, science journalists, and others. All of these elements are important and will probably be improved upon, but let's consider only statistical power for now.

Statistical power obviously depends on sample size, variability in responses, and the type of statistical tests that are used to analyze data. Experimental design also matters. What we have shown in this paper is that for one type of choice experiment which is widely used in experimental economic, there is a sweet spot for short experiments—about 5-10 trials—in which about twice as much information is generated by an adaptive design. This innovation is not that difficult to implement, and will immediately improve the quality of inference and therefore improve reproducibility.

[21] Many previous studies have made this point and used non-choice data. Some recent papers include Clithero (2018), Franco-Watkins et al. (2016), Frydman and Krajbich (2016), Konovalov and Krajbich (2016).

References

Abdellaoui, M. (2000): "Parameter-Free Elicitation of Utility and Probability Weighting Functions," *Management Science*, 46, 1497–1512.

Abernethy, J., T. Evgeniou, O. Toubia, and J.-P. Vert (2008): "Eliciting Consumer Preferences Using Robust Adaptive Choice Questionnaires," in *IEEE Transactions on Knowledge and Data Engineering*, 145–155.

Afriat, S. N. (1967): "The Construction of Utility Functions from Expenditure Data," *International Economic Review*, 8, 67–77.

——— (1972): "Efficiency Estimation of Production Functions," *International Economic Review*, 13, 568–598.

Ahn, D. S., S. Choi, D. Gale, and S. Kariv (2014): "Estimating Ambiguity Aversion in a Portfolio Choice Experiment," *Quantitative Economics*, 5, 195–223.

Andersen, S., G. W. Harrison, M. I. Lau, and E. E. Rutström (2006): "Elicitation Using Multiple Price List Formats," *Experimental Economics*, 9, 383–405.

——— (2008): "Eliciting Risk and Time Preferences," *Econometrica*, 76, 583–618.

Andreoni, J., M. A. Kuhn, and C. Sprenger (2015): "Measuring Time Preferences: A Comparison of Experimental Methods," *Journal o f Economic Behavior & Organization*, 116, 451–464.

Andreoni, J. and J. Miller (2002): "Giving According to GARP: An Experimental Test of the Consistency of Preferences for Altruism," *Econometrica*, 70, 737–753.

Andreoni, J. and C. Sprenger (2012): "Estimating Time Preferences from Convex Budgets," *American Economic Review*, 102, 3333–3356.

Andreoni, J. and L. Vesterlund (2001): "Which is the Fair Sex? Gender Differences in Altruism," *Quarterly Journal of Economics*, 116, 293–312.

Angeletos, G.-M., D. Laibson, A. Repetto, J. Tobacman, and S. Weinberg (2001): "The Hyberbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation," *Journal of Economic Perspectives*, 15, 47–68.

Apesteguia, J. and M. A. Ballester (2018): "Monotone Stochastic Choice Models: The Case of Risk and Time Preferences," *Journal of Political Economy*, 126, 74–106.

Augenblick, N., M. Niederle, and C. Sprenger (2015): "Working Over Time: Dynamic Inconsistency in Real Effort Tasks," *Quarterly Journal of Economics*, 130, 1067–1115.

Balakrishnan, U., J. Haushofer, and P. Jakiela (2015): "How Soon Is Now? Evidence of Present Bias from Convex Time Budget Experiments," Unpublished manuscript.

Bardsley, N., R. Cubitt, G. Loomes, P. Moffatt, C. Starmer, and R. Sugden (2009): *Experimental Economics: Rethinking the Rules*, Princeton, NJ: Princeton University Press.

Benhabib, J., A. Bisin, and A. Schotter (2010): "Present-Bias, Quasi-Hyperbolic Discounting, and Fixed Costs," *Games and Economic Behavior*, 69, 205–223.

Bickel, W. K., J. A. Pitcock, R. Yi, and E. J. C. Angtuaco (2009): "Congruence of BOLD Response across Intertemporal Choice Conditions: Fictive and Real Money Gains and Losses," *Journal of Neuroscience*, 29, 8839–8846.

Bousquet, L. (2016): "Measuring Time Preference and Anticipation: A Lab Experiment," Unpublished manuscript.

Brocas, I., J. D. Carrillo, and J. Tarrasó (2016): "How Long is a Minute?" Unpublished manuscript.

Brown, A. L. and H. Kim (2014): "Do Individuals Have Preferences Used in Macro-Finance Models? An Experimental Investigation," *Management Science*, 60, 939–958.

Burghart, D. R., P. W. Glimcher, and S. C. Lazzaro (2013): "An Expected Utility Maximizer Walks into a Bar..." *Journal of Risk and Uncertainty*, 46, 215–246.

Camerer, C. F., A. Dreber, E. Forsell, T.-H. Ho, J. Huber, M. Johannesson, M. Kirchler, J. Almenberg, A. Altmejd, T. Chan, E. Heikensten, F. Holzmeister, T. Imai, S. Isaksson, G. Nave, T. Pfeiffer, M. Razen, and H. Wu (2016): "Evaluating Replicability of Laboratory Experiments in Economics," *Science*, 351, 1433–1436.

Camerer, C. F., A. Dreber, F. Holzmeister, T.-H. Ho, J. Huber, M. Johannesson, M. Kirchler, G. Nave, B. A. Nosek, T. Pfeiffer, A. Altmejd, N. Buttrick, T. Chan, Y. Chen, E. Forsell, A. Gampa, E. Heikensten, L. Hummer, T. Imai, S. Isaksson, D. Manfredi, J. Rose, E.-J. Wagenmakers, and H. Wu (2018): "Evaluating the Replicability of Social Science Experiments in *Nature* and *Science* between 2010 and 2015," *Nature Human Behavior*, 2, 637–644.

Camille, N., C. A. Griffiths, K. Vo, L. K. Fellows, and J. W. Kable (2011): "Ventromedial Frontal Lobe Damage Disrupts Value Maximization in Humans," *Journal of Neuroscience*, 31, 7527–7532.

Cavagnaro, D. R., G. J. Aranovich, S. M. McClure, M. A. Pitt, and J. I. Myung (2016): "On the Functional Form of Temporal Discounting: An Optimized Adaptive Test," *Journal of Risk and Uncertainty*, 52, 233–254.

Cavagnaro, D. R., R. Gonzalez, J. I. Myung, and M. A. Pitt (2013a): "Optimal Decision Stimuli for Risky Choice Experiments: An Adaptive Approach," *Management Science*, 59, 358–375.

Cavagnaro, D. R., J. I. Myung, M. A. Pitt, and J. V. Kujala (2010): "Adaptive Design Optimization: A Mutual Information-Based Approach to Model Discrimination in Cognitive Science," *Neural Computation*, 22, 887–905.

Cavagnaro, D. R., M. A. Pitt, R. Gonzalez, and J. I. Myung (2013b): "Discriminating among Probability Weighting Functions Using Adaptive Design Optimization," *Journal of Risk and Uncertainty*, 47, 255–289.

Cavagnaro, D. R., M. A. Pitt, and J. I. Myung (2011): "Model Discrimination Through Adaptive Experimentation," *Psychonomic Bulletin & Review*, 18, 204–210.

Chakaravarthy, V. T., V. Pandit, S. Roy, P. Awasthi, and M. Mohania (2007): "Decision Trees for Entity Identification: Approximation Algorithms and Hardness Results," in *Proceedings of the twenty-sixth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*, 53–62.

Chaloner, K. and I. Verdinelli (1995): "Bayesian Experimental Design: A Review," *Statistical Science*, 10, 273–304.

Chapman, J., E. Snowberg, S. W. Wang, and C. F. Camerer (2018): "Dynamically Optimized Sequential Experimentation (DOSE) for Estimating Economic Preference Parameters," Unpublished manuscript.

Chen, M. K., V. Lakshminarayanan, and L. R. Santos (2006): "How Basic are Behavioral Biases? Evidence from Capuchin Monkey Trading Behavior," *Journal of Political Economy*, 114, 517–537.

Cheung, S. L. (2016): "Recent Developments in the Experimental Elicitation of Time Preference," *Journal of Behavioral and Experimental Finance*, 11, 1–8.

Choi, S., R. Fisman, D. Gale, and S. Kariv (2007): "Consistency and Heterogeneity of Individual Behavior under Uncertainty," *American Economic Review*, 97, 1921–1938.

Choi, S., S. Kariv, W. Müller, and D. Silverman (2014): "Who is (More) Rational," *American Economic Review*, 104, 1518–1550.

Clithero, J. A. (2016): "Response Times in Economics: Looking Through the Lens of Sequential Sampling Models," Unpublished manuscript.

——— (2018): "Improving Out-of-Sample Predictions Using Response Times and a Model of the Decision Process," *Journal of Economic Behavior & Organization*, 148, 344–375.

Cohen, J., K. Ericson, D. Laibson, and J. White (2016): "Measuring Time Preferences," Unpublished manuscript.

Coller, M. and M. B. Williams (1999): "Eliciting Individual Discount Rates," *Experimental Economics*, 2, 107–127.

Cornsweet, T. N. (1962): "The Staircase-Method in Psychophysics," *American Journal of Psychology*, 75, 485–491.

Cover, T. M. and J. A. Thomas (1991): *Elements of Information Theory*, New York: Wiley.

Dasgupta, S. (2004): "Analysis of a Greedy Active Learning Strategy," in *Advances in Neural Information Processing Systems*, 337–344.

DeGroot, M. H. (1970): *Optimal Statistical Decisions*, New York: McGraw-Hill.

Dimmock, S. G., R. Kouwenberg, and P. P. Wakker (2016): "Ambiguity Attitudes in a Large Representative Sample," *Management Science*, 62, 1363–1380.

Echenique, F., T. Imai, and K. Saito (2016a): "A Cross-Study Analysis of Behavior in the Convex Time Budget Experiments," Unpublished manuscript.

——— (2016b): "Testable Implications of Models of Intertemporal Choice: Exponential Discounting and Its Generalizations," Caltech Social Science Working Paper #1388.

Echenique, F., S. Lee, and M. Shum (2011): "The Money Pump as a Measure of Revealed Preference Violations," *Journal of Political Economy*, 119, 1201–1223.

El-Gamal, M. A., R. D. McKelvey, and T. R. Palfrey (1993): "A Bayesian Sequential Experimental Study of Learning in Games," *Journal of the American Statistical Association*, 88, 428–435.

El-Gamal, M. A. and T. R. Palfrey (1996): "Economical Experiments: Bayesian Efficient Experimental Design," *International Journal of Game Theory*, 25, 495–517.

Ericson, K. M. and J. Noor (2015): "Delay Functions as the Foundation of Time Preference: Testing for Separable Discounted Utility," NBER Working Paper No. 21095.

Falk, A., A. Becker, T. J. Dohmen, D. Huffman, and U. Sunde (2016): "The Preference Survey Module: A Validated Instrument for Measuring Risk, Time, and Social Preferences," Unpublished manuscript.

Fisman, R., S. Kariv, and D. Markovits (2007): "Individual Preferences for Giving," *American Economic Review*, 97, 1858–1876.

Franco-Watkins, A. M., R. E. Mattson, and M. D. Jackson (2016): "Now or Later? Attentional Processing and Intertemporal Choice," *Journal of Behavioral Decision Making*, 29, 206–217.

Frederick, S., G. Loewenstein, and T. O'donoghue (2002): "Time Discounting and Time Preference: A Critical Review," *Journal of Economic Literature*, 40, 351–401.

Frydman, C. and I. Krajbich (2016): "Using Response Times to Infer Others' Beliefs: An Application to Information Cascades," Unpublished manuscript.

Gafni, A. and G. W. Torrance (1984): "Risk Attitude and Time Preference in Health," *Management Science*, 30, 440–451.

Golovin, D. and A. Krause (2010): "Adaptive Submodularity: A New Approach to Active Learning and Stochastic Optimization," in *Proceedings of Annual Conference on Learning Theory (COLT)*.

Golovin, D., A. Krause, and D. Ray (2010): "Near-Optimal Bayesian Active Learning with Noisy Observations," in *Advances in Neural Information Processing Systems (NIPS)*.

Goodman, J. K., C. E. Cryder, and A. Cheema (2013): "Data Collection in a Flat World: The Strengths and Weaknesses of Mechanical Turk Samples," *Journal of Behavioral Decision Making*, 26, 213–224.

Harbaugh, W. T., K. Krause, and T. R. Berry (2001): "GARP for Kids: On the Development of Rational Choice Behavior," *American Economic Review*, 91, 1539–1545.

Hardisty, D. J., K. F. Thompson, D. H. Krantz, and E. U. Weber (2013): "How to Measure Time Preferences: An Experimental Comparison of Three Methods," *Judgment and Decision Making*, 8, 236–249.

Harrison, G. W., M. Lau, and E. E. Ruström (2013): "Identifying Time Preferences with Experiments: Comment," Unpublished manuscript.

Harrison, G. W., M. I. Lau, and M. B. Williams (2002): "Estimating Individual Discount Rates in Denmark: A Field Experiment," *American Economic Review*, 92, 1606–1617.

Horton, J. J., D. G. Rand, and R. J. Zeckhauser (2011): "The Online Laboratory: Conducting Experiments in a Real Labor Market," *Experimental Economics*, 14, 399–425.

Howard, R. A. (1966): "Information Value Theory," in *IEEE Transactions on Systems Science and Cybernetics*, 22–26.

Huber, J., J. W. Payne, and C. Puto (1982): "Adding Asymmetrically Dominated Alternatives: Violations of Regularity and the Similarity Hypothesis," *Journal of Consumer Research*, 9, 90–98.

Imai, T., T. Rutter, and C. F. Camerer (2018): "Meta-Analysis of Estimation of Present Bias from Convex Time Budget Experiments," Unpublished manuscript.

Janssens, W., B. Kramer, and L. Swart (2017): "Be Patient When Measuring Hyperbolic Discounting: Stationarity, Time Consistency and Time Invariance in a Field Experiment," *Journal of Development Economics*, 126, 77–90.

Johnson, C., A. Baillon, H. Bleichrodt, Z. Li, D. van Dolder, and P. P. Wakker (2015): "Prince: An Improved Method for Measuring Incentivized Preferences," Unpublished manuscript.

Johnson, M. W. and W. K. Bickel (2002): "Within-Subject Comparison of Real and Hypothetical Money Rewards in Delay Discounting," *Journal of the Experimental Analysis of Behavior*, 77, 129–146.

Kable, J. W. and P. W. Glimcher (2007): "The Neural Correlates of Subjective Value during Intertemporal Choice," *Nature Neuroscience*, 10, 1625–1633.

——— (2010): "An "As Soon As Possible? Effect in Human Intertemporal Decision Making: Behavioral Evidence and Neural Mechanisms," *Journal of Neurophysiology*, 103, 2513–2531.

Konovalov, A. and I. Krajbich (2016): "Revealed Indifference: Using Response Times to Infer Preferences," Unpublished manuscript.

Koopmans, T. C. (1960): "Stationary Ordinal Utility and Impatience," *Econometrica*, 28, 287–309.

Krajbich, I., C. Camerer, and A. Rangel (2017): "Exploring the Scope of Neurometrically Informed Mechanism Design," *Games and Economic Behavior*, 101, 49–62.

Kuhn, M., P. Kuhn, and M. C. Villeval (2017): "Decision-Environment Effects on Intertemporal Financial Choices: How Relevant are Resource-Depletion Models?" *Journal of Economic Behavior & Organization*, 137, 72–89.

Kullback, S. and R. A. Leibler (1951): "On Information and Sufficiency," *Annals of Mathematical Statistics*, 22, 79–86.

Laibson, D. (1997): "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, 112, 443–478.

Laury, S. K., M. M. McInnes, and J. T. Swarthout (2012): "Avoiding the Curves: Direct Elicitation of Time Preferences," *Journal of Risk and Uncertainty*, 44, 181–217.

Lindley, D. V. (1956): "On a Measure of the Information Provided by an Experiment," *Annals of Mathematical Statistics*, 27, 986–1005.

Loomes, G. (1991): "Evidence of a New Violation of the Independence Axiom," *Journal of Risk and Uncertainty*, 4, 91–108.

MacKay, D. J. C. (1992): "Information-Based Objective Functions for Active Data Selection," *Neural Computation*, 4, 590–604.

Madden, G. J., A. M. Begotka, B. R. Raiff, and L. L. Kastern (2003): "Delay Discounting of Real and Hypothetical Rewards," *Experimental and Clinical Psychopharmacology*, 11, 139–145.

Madden, G. J., B. R. Raiff, C. H. Lagorio, A. M. Begotka, A. M. Mueller, D. J. Hehli, and A. A. Wegener (2004): "Delay Discounting of Potentially Real and Hypothetical Rewards: II. Between-and Within-Subject Comparisons," *Experimental and Clinical Psychopharmacology*, 12, 251–261.

Mason, W. and S. Suri (2012): "Conducting Behavioral Research on Amazon's Mechanical Turk," *Behavior Research Methods*, 44, 1–23.

Mazur, J. E. (1987): "An Adjusting Procedure for Studying Delayed Reinforcement," in *The Effect of Delay and Intervening Events on Reinforcement Value*, ed. by M. L. Commons, J. E. Mazur, J. A. Nevin, and H. Rachlin, Hillsdale, NJ: Erlbaum, 55–73.

McClure, S. M., D. I. Laibson, G. Loewenstein, and J. D. Cohen (2004): "Separate Neural Systems Value Immediate and Delayed Monetary Rewards," *Science*, 306, 503–507.

Meier, S. and C. D. Sprenger (2015): "Temporal Stability of Time Preferences," *Review of Economics and Statistics*, 97, 273–286.

Moffatt, P. G. (2007): "Optimal Experimental Design in Models of Decision and Choice," in *Measurement in Economics*, ed. by M. Boumans, London: Academic Press, 357–375.

——— (2016): *Experimetrics: Econometrics for Experimental Economics*, London: Palgrave Macmillan.

Montiel Olea, J. L. and T. Strzalecki (2014): "Axiomatization and Measurement of Quasi-Hyperbolic Discounting," *Quarterly Journal of Economics*, 129, 1449–1499.

Myung, J. I., D. R. Cavagnaro, and M. A. Pitt (2013): "A Tutorial on Adaptive Design Optimization," *Journal of Mathematical Psychology*, 57, 53–67.

Myung, J. I. and M. A. Pitt (2009): "Optimal Experimental Design for Model Discrimination," *Psychological Review*, 116, 499–518.

Myung, J. I., M. A. Pitt, Y. Tang, and D. R. Cavagnaro (2009): "Bayesian Adaptive Optimal Design of Psychology Experiments," in *Proceedings of the 2nd International Workshop in Sequential Methodologies (IWSM 2009)*.

Nowak, R. (2009): "Noisy Generalized Binary Search," in *Advances in Neural Information Processing Systems*, 1366–1374.

O'Donoghue, T. and M. Rabin (1999): "Incentives for Procrastinators," *Quarterly Journal of Economics*, 114, 769–816.

Paolacci, G., J. Chandler, and P. G. Ipeirotis (2010): "Running Experiments on Amazon Mechanical Turk," *Judgment and Decision Making*, 5, 411–419.

Peters, J. and C. Büchel (2010): "Episodic Future Thinking Reduces Reward Delay Discounting through an Enhancement of Prefrontal-Mediotemporal Interactions," *Neuron*, 66, 138–148.

Phelps, E. S. and R. A. Pollak (1968): "On Second-Best National Saving and Game-Equilibrium Growth," *Review of Economic Studies*, 35, 185–199.

Ray, D., D. Golovin, A. Krause, and C. Camerer (2012): "Bayesian Rapid Optimal Adaptive Design (BROAD): Method and Application Distinguishing Models of Risky Choice," Unpublished manuscript.

Samuelson, P. A. (1937): "A Note on Measurement of Utility," *Review of Economic Studies*, 4, 155–161.

Sawada, Y. and Y. Kuroishi (2015): "How Does a Natural Disaster Change People's Preference? The Case of a Large Scale Flood in the Philippines Using the Convex Time Budget Experiments," in *Disaster Risks, Social Preferences, and Policy Effects: Field Experiments in Selected ASEAN and East Asian Countries*, ed. by Y. Sawada and S. Oum, ERIA Research Project Report FY2013, No. 34, 27–56.

Simonson, I. (1989): "Choice Based on Reasons: The Case of Attraction and Compromise Effects," *Journal of Consumer Research*, 16, 158–174.

Sippel, R. (1997): "An Experiment on the Pure Theory of Consumer's Behaviour," *Economic Journal*, 107, 1431–1444.

Sun, C. and J. Potters (2016): "Magnitude Effect in Intertemporal Allocation Tasks," Unpublished manuscript.

Takeuchi, K. (2011): "Non-parametric Test of Time Consistency: Present Bias and Future Bias," *Games and Economic Behavior*, 71, 456–478.

Toubia, O., E. Johnson, T. Evgeniou, and P. Delequié (2013): "Dynamic Experiments for Estimating Preferences: An Adaptive Method of Eliciting Time and Risk Parameters," *Management Science*, 59, 613–640.

Tversky, A., P. Slovic, and D. Kahneman (1990): "The Causes of Preference Reversal," *American Economic Review*, 80, 204–217.

Ubfal, D. (2016): "How General are Time Preferences? Eliciting Good-Specific Discount Rates," *Journal of Development Economics*, 118, 150–170.

van de Kuilen, G. and P. P. Wakker (2011): "The Midweight Method to Measure Attitudes Toward Risk and Ambiguity," *Management Science*, 57, 582–598.

van den Bos, W., C. A. Rodriguez, J. B. Schweitzer, and S. M. McClure (2014): "Connectivity Strength of Dissociable Striatal Tracts Predict Individual Differences in Temporal Discounting," *Journal of Neuroscience*, 34, 10298–10310.

——— (2015): "Adolescent Impatience Decreases with Increased Frontostriatal Connectivity," *Proceedings of the National Academy of Sciences*, 112, E3765–E3774.

von Békésy, G. (1947): "A New Audiometer," *Acta Oto-Laryngologica*, 35, 411–422.

Wainer, H. and C. Lewis (1990): "Toward a Psychometrics for Testlets," *Journal of Educational Measurement*, 27, 1–14.

Wang, S. W., M. Filiba, and C. F. Camerer (2010): "Dynamically Optimized Sequential Experimentation (DOSE) for Estimating Economic Preference," Unpublished manuscript.

Wilcox, N. T. (2008): "Stochastic Models for Binary Discrete Choice Under Risk: A Critical Primer and Econometric Comparison," in *Research in Experimental Economics: Risk Aversion in Experiments*, ed. by C. Cox, James and G. W. Harrison, Emerald Group Publishing, vol. 12, 197–292.