NBER WORKING PAPER SERIES

LOSS ATTITUDES IN THE U.S. POPULATION: EVIDENCE FROM DYNAMICALLY OPTIMIZED SEQUENTIAL EXPERIMENTATION (DOSE)

Jonathan Chapman
Erik Snowberg
Stephanie Wang
Colin Camerer

Working Paper 25072
http://www.nber.org/papers/w25072

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2018

Thanks to John Beshears, Mark Dean, Kate Johnson, Ian Krajbich, Andreas Krause, Pietro Ortoleva, Deb Ray, Antonio Rangel, Hans-Martin von Gaudecker, Peter Wakker, Nathaniel Wilcox, and the participants of seminars and conferences for their useful comments and suggestions. Judah Okwuobi and Michelle Filiba provided excellent research assistance. Camerer and Snowberg gratefully acknowledge the financial support of NSF Grant SMA-1329195. This paper subsumes the earlier, largely methodological, working paper, "Dynamically Optimized Sequential Experimentation (DOSE) for Estimating Economic Preference Parameters," (2010). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Loss Attitudes in the U.S. Population: Evidence from Dynamically Optimized Sequential Experimentation (DOSE)
Jonathan Chapman, Erik Snowberg, Stephanie Wang, and Colin Camerer
NBER Working Paper No. 25072
September 2018
JEL No. C81,C9,D03,D81,D9

ABSTRACT

We introduce DOSE - Dynamically Optimized Sequential Experimentation - and use it to estimate individual-level loss aversion in a representative sample of the U.S. population (N=2,000). DOSE elicitations are more accurate, more stable across time, and faster to administer than standard methods. We find that around 50% of the U.S. population is loss tolerant. This is counter to earlier findings, which mostly come from lab/student samples, that a strong majority of participants are loss averse. Loss attitudes are correlated with cognitive ability: loss aversion is more prevalent in people with high cognitive ability, and loss tolerance is more common in those with low cognitive ability. We also use DOSE to document facts about risk and time preferences, indicating a high potential for DOSE in future research.

Jonathan Chapman
Division of Social Science
New York University Abu Dhabi
Abu Dhabi
UAE
jchapman@nyu.edu

Erik Snowberg
Division of Humanities and Social Sciences
MC 228-77
California Institute of Technology
Pasadena, CA 91125
and NBER
snowberg@caltech.edu

Stephanie Wang
University of Pittsburgh
Department of Economics
4906 Wesley W. Posvar Hall
230 South Bouquet Street
Pittsburgh, PA 15260
swwang@pitt.edu

Colin Camerer
Department of Economics
California Institute of Technology
Pasadena, CA 91125
camerer@hss.caltech.edu

Screenshots and Design Documents is available at http://www.its.caltech.edu/~snowberg/wep.html
An online appendix is available at http://www.nber.org/data-appendix/w25072/

1 Introduction

An important hypothesis in behavioral economics is that people treat losses and gains differently, resulting in most being *loss averse*: even if they are risk neutral, they tend to shy away from positive expected value gambles with negative payoffs (losses). Loss aversion is used as an explanation for a number of important economic phenomena,[1] and is an essential ingredient in theories of reference-dependent preferences (Kahneman and Tversky, 1979; Kőszegi and Rabin, 2006; O'Donoghue and Sprenger, 2018).

Yet, most evidence of loss aversion comes from economics and psychology labs, usually with university student participants who may have different preferences (Snowberg and Yariv, 2018; Walasek et al., in progress). The hypothesis of differential responses to gains and losses would thus benefit from individual-level assessments in a representative sample. However, as we detail below, current methodologies make such an assessment difficult. To overcome these difficulties, we introduce DOSE—Dynamically Optimized Sequential Experimentation—which estimates preference parameters precisely and quickly by selecting a personalized sequence of simple choices for each participant.

Using DOSE, we find that around 50% of people in the U.S. are *loss tolerant*: even if they are risk neutral, they embrace gambles with negative expected values. Loss aversion is more prevalent in people with high cognitive ability, and loss tolerance is more common in those with low cognitive ability. Moreover, we find that risk aversion over gains and loss attitudes are equally stable (and more stable than previously appreciated), suggesting that both are equally important in understanding risk preferences.

It is important to emphasize that, although surprising, the prevalence of loss tolerance is *not* evidence against the hypothesis of gain-loss differences. Rather, it is evidence of substantial heterogeneity in the asymmetry, with potentially important consequences. In

[1] Examples include the equity premium puzzle (Mehra and Prescott, 1985; Benartzi and Thaler, 1995), asymmetric consumer price elasticities (Hardie et al., 1993), downward sloping labor supply (Dunn, 1996; Camerer et al., 1997; Goette et al., 2004), tax avoidance (Rees-Jones, 2017), opposition to free trade (Tovar, 2009), performance in athletic contests (Pope and Simonsohn, 2011; Allen et al., 2016), and more.

particular, loss aversion can, in theory, reduce the propensity to use financial products that exploit common characteristics like overoptimism and skew-love (Kahneman and Lovallo, 1993; Åstebro et al., 2015). Loss tolerance, on the other hand, makes people more susceptible to exploitation of these characteristics. Moreover, our evidence suggests that loss tolerance is particularly prevalent in precisely the people who might benefit from additional reservations about problematic financial products: those with low income, education, and cognitive ability, and the aged (Kornotis and Kumar, 2010; Chang, 2016).

A new technique is needed to measure loss aversion (and other preferences) in representative populations. In the three studies that elicit loss aversion in a representative population, two lose more than 70% of participants due to non-response and inconsistent choice (Booij and Van de Kuilen, 2009; Booij et al., 2010). A third only recovers population distributions of loss aversion, and is very sensitive to estimation choices. Depending on those choices, it produces population estimates that vary from a large majority being extremely loss averse, to almost everyone being loss tolerant (appendix of von Gaudecker et al., 2011). Individual-level estimates are necessary to study the correlates of heterogeneity of preferences (Harrison et al., 2002), or to calibrate personalized contracts (Andreoni et al., 2016). Moreover, standard techniques for measuring other economic preferences produce unstable estimates (Meier and Sprenger, 2015), and may unintentionally introduce reference points (Sprenger, 2015). DOSE produces accurate (Section 3.2), stable (Section 4.4), and fast (Section 5.2) individual-level estimates of preference parameters that are robust to alternative specifications (Section 5.1). Using DOSE to study risk and time preferences produces similar—often less noisy—results to prior studies using standard elicitation techniques (Section 4.2 and Appendix B).

DOSE takes the challenges of eliciting loss aversion—the need for multiple choices, and usually a parametric model—and designs around them.[2] DOSE uses the parametric struc-

[2]Estimating an individual index of loss aversion without a parametric structure attributes all differences in the curvature of utility functions over gains and over losses to loss aversion. In principle, a non-parametric approach allows a classification of people into loss averse/neutral/tolerant, but in practice many cannot be classified. For example, using a non-parametric method, Abdellaoui et al. (2007) find that between 29% and 88% of participants cannot be classified, depending on which definition of loss aversion is used.

ture, and rapid computation of Bayesian updating, to dynamically select a personalized sequence from a set of simple choices, as described in Section 2. That is, DOSE starts with a prior over parameters and/or models and, based on that prior, selects a question that will maximize information. DOSE then uses a participant's choice to dynamically update its priors about that participant, and selects the next question in the personalized sequence based on the new distribution over parameters and/or models. The sequence continues this way until either the participant has been asked pre-set number of questions, or the precision of the estimates for that participant is greater than some pre-specified criterion.

We use DOSE to estimate loss aversion using an incentivized, representative survey of the U.S. population (N = 2,000), also described in Section 2. This incentivized survey has several useful features. It is comprehensive, using a wide range of elicitations to measure different preferences. Moreover, it is repeated, meaning the same participants are asked the same questions twice, six months apart. These features allow us to establish a number of facts about loss aversion, as well as evaluate DOSE, in an important practical setting.

Before examining data from the incentivized survey, we first examine, in Section 3, simulations indicating that DOSE is more than twice as accurate as two standard methods for eliciting loss and risk aversion: the multiple price list (MPL; see Andersen et al., 2006, for a review), and the Lottery Menu (Eckel and Grossman, 2002). We show this in two sets of simulations. The first uses the choices from lab/student samples (N = 120; from Sokol-Hessner et al. 2009 and Frydman et al. 2011) in 140 questions similar to those we use in DOSE. We use DOSE to simulate question orderings for these participants, and show that with only 20 (personalized) questions, DOSE captures most of the information from all 140 choices. We then use the distribution of parameters among the 120 laboratory participants to produce 10,000 simulated participants. By generating the choices of the simulated participants in different types of elicitations, we show that DOSE produces estimates that are at least twice as close to the true parameter values as these other methods.

There is a much higher level of loss tolerance in the U.S. population than indicated by prior samples. The loss aversion parameter in Prospect Theory, λ, indicates loss aversion when $\lambda > 1$, and loss tolerance when $\lambda < 1$. In Section 4 we find that 53% of the U.S. population is loss tolerant. This is higher than 13–30% (weighted average 22%, $N = 1{,}023$) in the eight studies we are aware of that investigate heterogeneity in loss aversion (all in lab samples).[3] Moreover, the median level of loss aversion in our study is 0.98, versus 1.5–2.5 in other samples. We show that this difference is not due to DOSE: among 439 lab/student participants in prior studies using DOSE—drawing on the working paper version of this manuscript (Wang et al., 2010)—10% are loss tolerant, with a median value of $\lambda = 1.99$. Moreover, those with greater education and cognitive ability, and lower age, are more likely to be loss averse in our representative sample.[4] These attributes describe the student samples usually used in studies of loss aversion. Indeed, in our data, 23% of those under 35 with a college education ($N = 101$) were loss tolerant, with a median value of $\lambda = 1.75$. Altogether, this suggests that the prevalence of findings of loss aversion, rather than loss tolerance, may be the result of inadvertently selecting highly loss-averse samples.

An important feature of DOSE is that it dynamically estimates, and adjusts for, an individual's level of choice consistency. This produces two more substantive results. First, although we find a correlation between higher cognitive ability and less risk aversion using DOSE, we do not find a statistically significant relationship with an MPL-based measure of risk aversion. However, if we examine only those participants DOSE tells us make consistent choices, we recover a similar relationship using the MPL measure, suggesting that choice inconsistency and resultant measurement error may lead to the mixed results on the relationship between cognitive ability and risk aversion (Dohmen et al., 2018). Second, we show that DOSE estimates of risk and time preferences are more stable across time than MPL-

[3]These studies are Schmidt and Traub (2002); Brooks and Zank (2005); Abdellaoui et al. (2007, 2008); Sokol-Hessner et al. (2009); Abdellaoui et al. (2011); Sprenger (2015); Goette et al. (2018). The figure for Sprenger (2015) is reported in Goette et al. (2018), Footnote 8.

[4]Most studies of the relationship between cognitive ability and risk preferences have focused on lotteries over gains (see Andersson et al., 2016b, Table E1 for a summary). The few studies with questions involving losses have found that lower cognitive ability is associated with fewer expected value maximizing choices on those lotteries—consistent with our results—although differences in design and data reporting make it difficult to ascertain the degree of agreement.

based measures. DOSE estimates are equally stable regardless of choice consistency, while MPL-based estimates are less stable for less consistent participants, suggesting measurement error in current methodologies is responsible for low levels of estimated preference stability (Gillen et al., Forthcoming). Additionally, loss aversion is nearly as stable as risk aversion, indicating that both are similarly important in describing a participant's risk preferences.

Our results from DOSE are robust to a number of factors, such as misspecification and removing participants most likely to not be paying attention, as shown in Section 5. Allowing for different specifications of the utility function still results in much lower estimates of loss aversion, and much higher estimates of loss tolerance, than prior studies on student/lab populations. Moreover, we show that DOSE is equally fast to complete for people of all cognitive ability levels. In contrast, MPLs take longer for everyone, but especially for those with lower cognitive abilities. Removing participants that may be "rushing through" DOSE, or our entire study, has minimal effects on the distribution of DOSE-estimated parameters. Additional robustness checks are conducted in Appendices D, E, F, and G.

The paper concludes with a discussion, in Section 6, of the potential for DOSE to be used more widely, and of research settings in which the procedure is likely to be particularly valuable. Of particular interest are further questions about loss aversion that were not covered by our design. We finish with a description of how our results fit with the broader work in psychology and neuroscience on the processes underlying the gain-loss hypothesis.

1.1 Related Literature

Our work is related to three broad literatures: optimal experimental design, measuring economic preferences and their correlates in broad populations, and loss aversion. We review these literatures here: relationships between specific factual findings in this paper and others are included when we discuss those specific findings, and in Appendix B.

There is a large literature on optimal experimental design in computer science and statistics, but there is surprisingly little development of applications for economics.[5] Those few studies that exist focus on static, rather than dynamic, experiments (El-Gamal et al., 1993; El-Gamal and Palfrey, 1996). DOSE extends these ideas by implementing a dynamic design, in which questions are selected sequentially based on a participant's answers. This allows for the identification of models and parameters at the individual level, in contrast to prior designs which could only discriminate between models, or measure the distribution parameters, at a population level. Taking advantage of recent advances in computing power, we are also able to account for a much larger range of parameters in designing an experiment.

Two papers have examined dynamic experimental procedures, drawing on an earlier working paper version of this manuscript (Wang et al., 2010). Toubia et al. (2013) use a very similar method to study risk and time preferences. Imai and Camerer (2018) use DOSE to evaluate time preferences, but focus on model selection.[6] The latter paper uses a different information criterion—EC^2 rather than the Kullback-Leibler divergence we use—and a different model of inconsistencies in decision-making. The Kullback-Leibler criterion is particularly well suited to efficient parameter estimation (Ryan et al., 2016) but may not be as efficient in model selection. Thus, the main contribution of that paper is to illustrate how the novel criterion EC^2 is used, and apply it to distinguish different models of time preferences (in participants recruited from MTurk) more rapidly and precisely than earlier research. In contrast, the primary contributions here are to establish general performance characteristics of the DOSE method, compare them to extant methods, and document new facts about loss aversion in a representative population.

[5]The idea of optimal experimental design appears to originate most clearly in Peirce (1879), who described an "economic" theory of experimentation and applied it to the study of gravity. The idea of dynamic designs begins with Wald (1950). Chaloner and Verdinelli (1995) provides a useful review of applications in statistics. Although little used in economics, optimal designs have been used in many applied fields including in neurophysiology (Lewi et al., 2009), psychophysics (Kujala and Lukka, 2006; Lesmes et al., 2006), marketing (Toubia et al., 2004; Abernethy et al., 2008), and medicine (Müller et al., 2007). See also Aigner (1979) for an early survey in economics, and Moffatt (2007) for a discussion of potential applications of optimal design to parameter estimation, including the elicitation of risk preferences.

[6] Cavagnaro et al. (2010, 2013a,b, 2016) independently develop an adaptive framework for model discrimination. Their implementations use many more questions than DOSE—for example, 80 in Cavagnaro et al. (2016), 101 in Cavagnaro et al. (2013a)—making it difficult to use with a representative sample.

Our paper also contributes to the recent literature studying the correlates of economic preferences in broad populations. Many of these studies focus on the role of cognitive ability in economic preferences, generally concluding that higher cognitive ability is associated with greater normative rationality (Frederick, 2005; Burks et al., 2009; Oechssler et al., 2009; Dohmen et al., 2010; Benjamin et al., 2013). We add to this literature in two ways. First, we examine the relationship between loss aversion and cognitive ability, and find that both low and high cognitive ability people tend to depart from normative rationality, but in different ways. Second, we show that the mixed results on the relationship between risk aversion and cognitive ability (Dohmen et al., 2018) are likely due to measurement error and imprecision in the elicitation techniques used in prior studies.

Finally, our paper relates to a larger literature interested in understanding loss aversion. Most studies focus on lab/student populations.[7] von Gaudecker et al. (2011) is the most similar to our work. As noted above, they focus on population distributions, and their estimates are sensitive to estimation choices. Depending on those choices, they report estimates of the median λ ranging from 0.12 to 4.47 (in their appendix).[8] As we show in Section 5.1 and Appendix F, our results are relatively stable with respect to different specifications. It is worth noting, however, that their results are not inconsistent with ours: the shape of the loss aversion distribution we find (in Figure 4) is very similar to theirs. Moreover, some of their specifications produce results much closer to ours than the prior literature.

2 The DOSE Procedure

This section introduces the DOSE procedure and our incentivized survey. We start with an abstract overview of DOSE, focusing on the choices experimenters can make to tune it to

[7] See Table 1 of Booij et al. (2010) and Table S4 of Sokol-Hessner et al. (2009) for estimates from lab studies. We are aware of four field studies that measure loss aversion in non-representative populations, but only report first moments. These studies feature samples of customers at a car manufacturer (Gächter et al., 2007), Vietnamese villagers (Tanaka et al., 2010), Mechanical Turk workers (Toubia et al., 2013), and U.S. mortgage holders (Atlas et al., 2017). Reported first moments of loss aversion are similar to lab studies.

[8] Their estimation strategy also does not allow them to use the S-shaped utility function suggested by Prospect Theory (Kahneman and Tversky, 1979).

their application, before describing the specific design choices we made to estimate risk and time preferences in a representative sample of the U.S. population.

2.1 DOSE in the Abstract

DOSE asks each participant a personalized set of questions. Questions are selected sequentially, using a participant's previous answers to identify the most informative question at that point in time. When selecting each successive question, DOSE accounts for the possibility that the participant may have made mistakes in his or her previous choices. Altogether this leads to accurate parameter estimates after only a few questions.

The procedure starts with a prior over a set of parameter values, and then optimally (according to some pre-defined criterion) selects questions to pinpoint a participant's preferences. The experimenter can choose a different prior, based on observables, for each participant. DOSE selects the optimal question given the prior and the optimality criterion. After a participant answers the first question, DOSE updates beliefs using Bayes's law, optimally selects the next question, and so on. The process continues for as many questions as the experimenter wants, or until posterior beliefs are more precise than some pre-set criterion.

DOSE can elicit more accurate parameter estimates than other common dynamic experimental designs because it allows for the possibility that participants make mistakes, as we illustrate by comparing DOSE with a simple partitioning method, in Figure 1. Partitioning techniques include the iterative MPL (see, for example, Andersen et al., 2006; von Gaudecker et al., 2011) and the staircase method (Falk et al., 2018).[9] In the example in Figure 1, both methods start with a uniform prior and offer participants a binary choice. In the first round, each participant faces the same question (Q_1 or q_1). Beliefs are then updated depending on the answer they provide, and the next question is picked optimally given the new beliefs. The key difference between the two procedures is that a partitioning method successively

[9] The iterative MPL presents participants with an initial MPL, and then offers them a refined set of options in another MPL. For example, if the choice on the first MPL implied a participant's certainty equivalent for a lottery lay between $X and $Y, the next MPL would have options in [$X,$Y].

eliminates ranges of parameter values after each question. DOSE, in contrast, allows for the possibility that any choice may have been a mistake, and hence places a positive probability on all parameter values regardless of previous answers.

In a partitioning method, a single incorrect choice causes considerable inaccuracy. Consider a participant with true parameter θ_0, displayed in the bottom panel of Figure 1. This participant should choose B in both questions of a partition method. If, however, he or she incorrectly chooses A in the first question, his or her estimated parameter value is constrained to be less than the median value—regardless of the number of rounds of questions. Errors early in the procedure thus lead to considerable measurement error. Any error makes it particularly hard to identify parameter values at the extremes of the distribution.

DOSE, in contrast, can elicit accurate parameter values even after a participant makes a mistake. Even after an initial incorrect choice of a, the posterior distribution places a positive probability on the true parameter value θ_0. As a result, with enough correct answers in future rounds, an accurate parameter estimate will still be obtained. Further, the procedure keeps track of the extent of inconsistent choice, which, as we demonstrate empirically in Section 4.3, provides a valuable measure of participant behavior.

The precise way in which DOSE selects questions, or accounts for possible mistakes, can easily be adapted to meet the needs of a particular research question. Researchers can modify the optimality criterion used to pick the personalized question sequence, and can choose any parametric model to capture the way in which participants make mistakes. We explain the design choices we use in this paper in the following subsection.

2.2 DOSE Procedure to Estimate Risk and Loss Aversion

DOSE can be customized for particular research questions. The main objects of choice for a researcher are the parametric specification(s), the prior distribution over parameters or models, the set of choices to present to participants, how parameters map to choices—that is, the structure of possible mistakes—and the information criterion used to select the next

Figure 1: DOSE improves estimate accuracy by allowing for choice inconsistency.

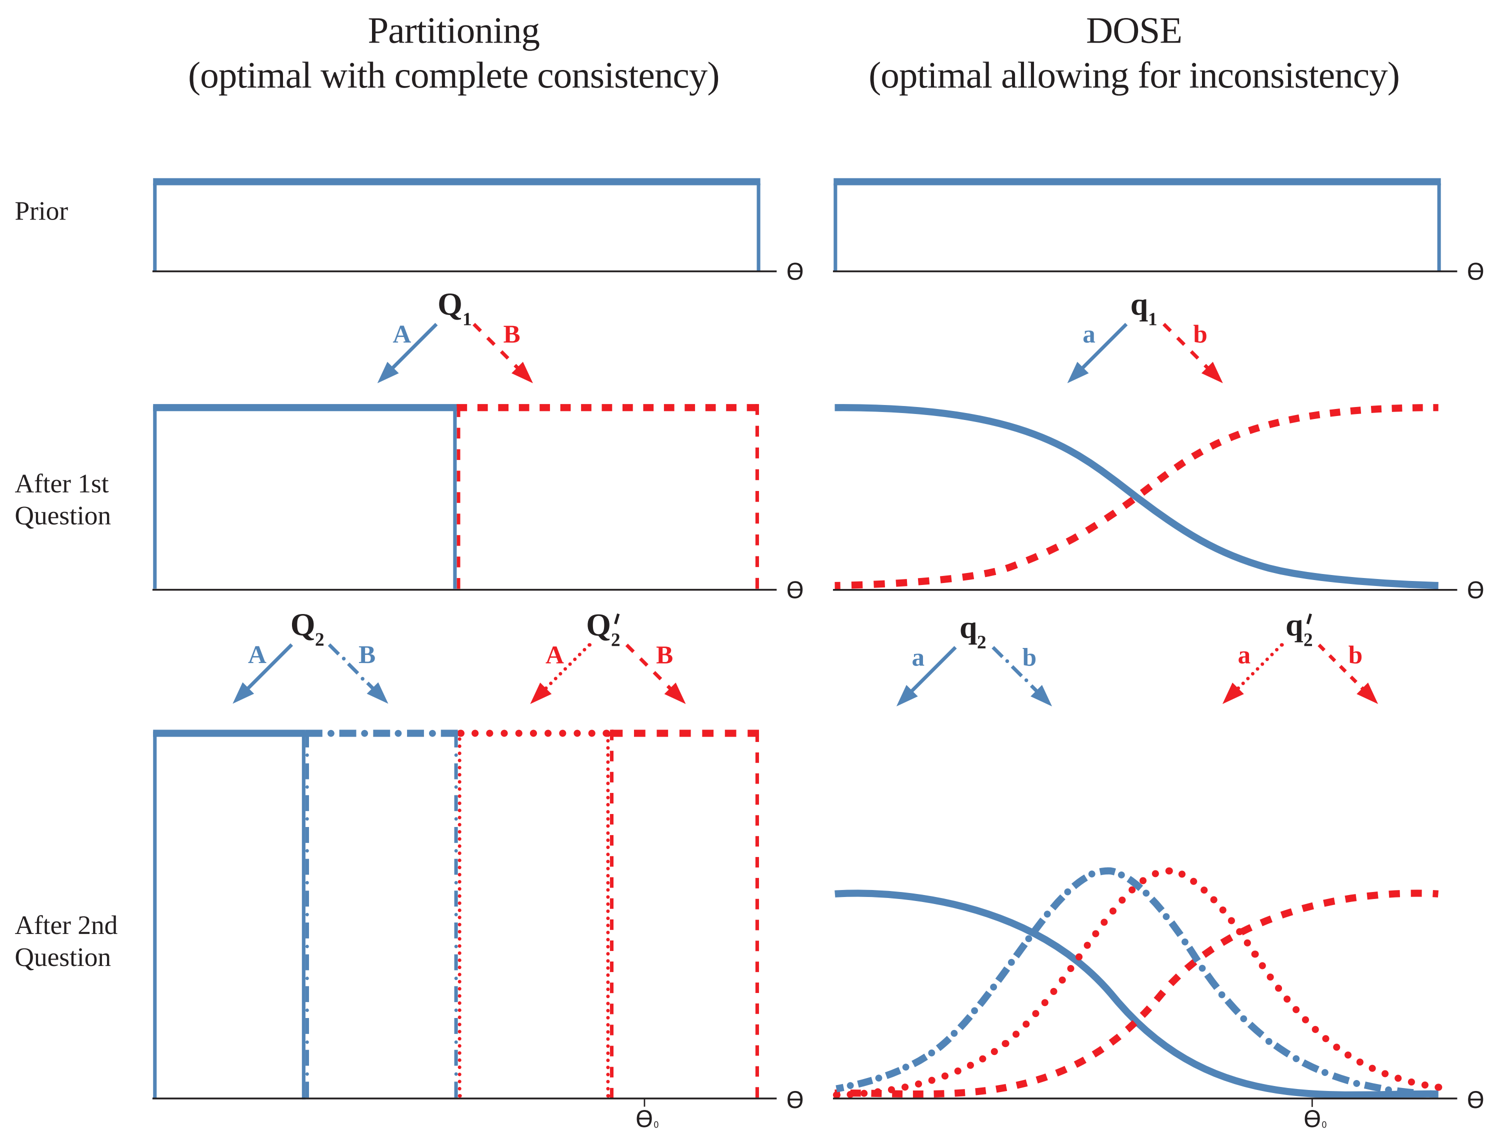

question based on current beliefs. This subsection details the choices we made to elicit loss aversion from a representative population.

2.2.1 Utility Function and Priors over Parameters

We elicit risk and loss aversion using a Prospect Theory utility function with power utility (Kahneman and Tversky, 1979). This utility function assumes that participants value payments relative to a reference point, which we assume is zero. The standard S-shaped utility function in Prospect Theory implies that, for common parameter values, participants are risk averse over positive payments (gains), and risk loving over negative payments (losses). A kink in the utility function at zero represents loss aversion. Formally:

$$v(x, \rho_i, \lambda_i) = \begin{cases} x^{\rho_i} & \text{for } x \geq 0 \\ -\lambda_i(-x)^{\rho_i} & \text{for } x < 0, \end{cases} \tag{1}$$

in which λ_i parameterizes loss aversion, ρ_i parameterizes risk aversion, and $x \in \mathbb{R}$ is a monetary outcome relative to the reference point. If $\lambda_i > 1$, then the participant is loss averse. If $\lambda_i < 1$, then the participant is loss tolerant. An individual with $\rho_i < 1$ demonstrates risk aversion over gains and risk love over losses. So that higher numbers indicate greater risk aversion, we use the *coefficient of relative risk aversion*: $1 - \rho_i$, in tables and figures.

The specification in (1) focuses on accurately estimating loss aversion with as few questions as possible. It does not allow for other common features of Prospect Theory: probability weighting or differential curvature of the utility function over losses and gains. The lotteries in our questions are further designed to minimize probability distortions, as all have 50/50 probabilities of two outcomes. Moreover, as most studies have found limited difference in curvature across the two domains (see Booij et al., 2010, Table 1), we impose the same utility curvature for both gains and losses.[10] This improves the accuracy of the estimates

[10] Assuming the same curvature across gains and losses also avoids an issue with power utility: different curvatures mean that estimates of loss aversion depend on scaling (Köbberling and Wakker, 2005). Moreover, there will always be an amount x s.t. $U(x) > U(-x)$ (Wakker, 2010). Our results are similar using

of λ_i as questions are not selected to separately identify the curvature in the loss and gain domain. However, it is quite simple to allow for different curvatures even after conducting the experiment—in Section 5.1 we show our results are robust to estimating this alternative specification, and that the assumption of equal curvature is supported by the data.

Time discounting is modeled with a standard monthly discount factor and the power utility function in (1).[11] Utility from the perspective of the survey date is given by $u(x_t, \rho_i, \delta_i) = \delta_i^t x_t^{\rho_i}$, where δ_i is a discount factor and ρ_i captures the curvature of the utility function from (1), t is the time from the survey date in months, and x_t is a payment at time t.

While researchers must choose a parametric specification and prior distribution for data collection, an alternative prior or specification can be used to calculate parameter values ex post. The experimenter's initial choices are used only to generate the personalized question sequence: once participant choices are recorded, any other prior distribution or parametric specification can be used to derive new estimates from the data. We use a joint uniform prior over preference parameters to both select questions and estimate parameters. The range of the prior distribution is chosen to cover the individual estimates obtained in Section 3.1.[12] As shown in Section 3.1 this results in near-optimal question selection.

Re-analyzing the choices obtained using questions selected by DOSE with a different prior or parametric specification allows researchers to obtain accurate parameter estimates even when the initial choices are misspecified, as shown in Sections 3.1 and 5.1. Despite initial misspecification, the questions asked still home in on a participant's preferences, even though the question sequence is not optimal in the sense defined by the question selection criterion. As a result, DOSE still provides a great deal of information about individual preferences, and precise estimates can be recovered ex post even with some initial misspecification.

---

the exponential (CARA) utility function Köbberling and Wakker (2005) suggest to avoid these issues (see Section 5.1).

[11] The specification used for question selection in the time preference module also allowed for present bias. In practice, however, we found very little evidence of present bias either in the DOSE module or the time MPLs—possibly due to the fact that payment was, in general, not instantly convertible into consumption. As such, the specification used to obtain estimates did not include a present bias parameter.

[12] In particular, the prior ranges are $\lambda \in [0, 4.6]$, $\rho \in [0.2, 1.7]$, $\mu \in [0, 8]$, and $\delta \in [0.2, 1]$.

2.2.2 Mistakes and Choice Consistency

As described above, an important advantage of DOSE is that, when selecting the personalized sequence of questions, it takes into account the possibility that participants make mistakes. The process by which mistakes are made must be parametrically modeled for DOSE to account for it. We model the mapping between utility and choices using the logit function, which has been widely used in both economics and psychophysics due to its connection with the random utility model.[13] For any choice between options o_1 and o_2 with $V(o_1) > V(o_2)$:

$$\text{Prob}[o_1] = \frac{1}{1 + e^{-\mu_i(V(o_1)-V(o_2))}}. \tag{2}$$

The logit function depends both on the utility difference between options o_1 and o_2 and the choice *consistency* parameter $\mu_i \in \mathbb{R}^+$. The probability of making a mistake—that is, not choosing the value-maxmizing option—is $1 - \text{Prob}[o_1]$. This is decreasing in the value difference between o_1 and o_2. This decrease is more rapid when μ_i is larger, so higher values of this parameter represent greater consistency in choices.

The set of questions can be designed to reduce the likelihood of mistakes. Our options were constructed to make expected value comparisons as simple as possible. All of our questions include only two options, with only one of these being a non-degenerate lottery. Each lottery has only two possible prizes. One of three payoffs (the sure payoff, and the two lottery payoffs) is always zero. Choices are thus between either a lottery with a zero payoff and some gain, versus some (possibly negative) sure amount; or between a lottery with a gain and loss, versus a sure amount of zero. In the former case, the expected value can be found by dividing by two. In the latter, one can ascertain if the expected value of the lottery is greater or less than zero by comparing the size of the positive and negative payoffs.

We believe these questions are also unlikely to produce inadvertent reference points, as

[13] Specifically, choice probabilities will be logit if the errors in the random utility model have an Extreme Value Type I distribution. See McFadden (2001) for a broader discussion of the history of the logit specification and its properties. DOSE can easily be implemented with multi-answer question using a multinomial logit or alternative probabilistic choice function.

MPLs have been shown to do (Sprenger, 2015; Chapman et al., 2017). However, this is also a testable prediction: we can try to fit specifications with alternative reference points, such as incorporating the endowment of points given at the beginning of the DOSE module. As shown in Appendix F.2, this model produces a much worse fit: it only predicts 48% of choices correctly, whereas our main specification predicts 88% correctly.

2.2.3 Information Criterion

In our implementation, DOSE selects each question to maximize the expected Kullback-Leibler (KL) divergence between the prior and possible posteriors associated with each answer. That is, the question that is picked at each point is the one with the highest expected information gain given the initial prior and previous answers. The KL criterion has been used widely in the optimal design literature in statistics due to its conceptual simplicity and grounding in information theory (see Ryan et al., 2016, for a discussion and examples). Further, the information maximization approach leads to consistent and efficient parameter estimates under weak modeling conditions (Paninski, 2005). However, DOSE is easily modified to incorporate alternative information criteria—for example, Imai and Camerer (2018) use DOSE with the EC^2 criterion to discriminate between models of time preferences.

Formally, consider a finite set of possible parameter vectors θ_k for $k = 1, ..., K$, where each $\theta_k = (\rho_k, \lambda_k, \delta_k, \mu_k)$ is a combination of possible values of the parameters of interest.[14] Each θ_k has an associated probability p_k of being the correct parameters. In the first question, these probabilities are the priors chosen by the experimenter; they are then updated in each round according to the participant's answers. The expected Kullback-Leibler divergence between the prior and the posterior when asking question Q_j is:

$$KL(Q_j) = \sum_{k \leq K} \sum_{a \in A} \log \left(\frac{l_k(a; Q_j)}{\sum_{j \in \mathcal{K}} p_j l_j(a; Q_j)} \right) p_k l_k(a; Q_j) \tag{3}$$

[14] We assume a finite space of parameters for computational ease. The KL divergence here is slightly different from that in El-Gamal and Palfrey (1996). Their variant maximizes the distance between posteriors (information) obtained under different models, whereas ours maximizes the information about parameters.

where $a \in A$ are the possible answers to the question, and $l_k(a; Q_j)$ is the likelihood of answer a given θ_k—in our implementation this is determined by the logit funciton in (2). DOSE selects the question that maximizes $KL(Q)$, the participant answers it, model posteriors are updated, the question Q_j that now maximizes $KL(Q)$ is selected, and so on.[15]

2.3 DOSE in a Representative Survey

We now turn to the practical details of implementing DOSE in two waves of a large, representative, incentivized survey of the U.S. population. The survey includes two DOSE modules—one relating to risk and one to time preferences—as well as other behavioral elicitations, and cognitive and sociodemographic questions.[16]

2.3.1 Survey Implementation

The two waves of the incentivized survey used the same questions and the same people, about six months apart. The first wave of the survey collected responses from 2,000 U.S. adults and was conducted online by YouGov between March 27 and April 3, 2015. A second wave recontacted the same population and received 1,465 responses between September 21 and November 23, 2015. We use data from the first wave for most analyses. Results are similar when using the second wave data, as shown in Appendix D.3.

Participants in the survey were drawn from a panel of respondents maintained by YouGov. YouGov continually recruits new people to the panel, especially from hard-to-reach and low-socioeconomic-status groups. To generate a representative sample, it randomly draws people from various Census Bureau products, and matches them on observables to members of their panel. Differential response rates lead to the over- and under-representation of certain populations and so YouGov provides sample weights to recover estimates that would be

[15] We restricted the procedure to only consider questions that had not yet been asked of that participant. In order to improve the estimate of μ, the procedure would eventually ask the same question multiple times.

[16] For specific details of the implementation of these, and other, questions see Appendix C and Chapman et al. (2017), or screenshots and design documents at hss.caltech.edu/∼snowberg/wep.html.

obtained from a fully representative sample. We use these weights throughout the paper.[17]

The behavioral measures in this paper were all incentivized: at the end of the survey, two survey modules were selected for payment at random.[18] All outcomes were expressed in YouGov points, an internal YouGov currency used to pay panel members, which can be converted to U.S. dollars using the approximate rate of $0.001 per point.[19] To enhance the credibility of these incentives, we took advantage of YouGov's relationship with its panel, and restricted the sample to those who had already been paid (in cash or prizes) for their participation in surveys. The average payment to respondents (including the show-up fee) was $9 (9,000 points), which is approximately three times the average for YouGov surveys. For comparability, we convert points to dollars, using the exchange rate, in our analyses.

2.3.2 DOSE Modules

All respondents were asked two, ten-question, DOSE modules.

Risk Preferences: The first DOSE module elicited risk and loss aversion. Participants were given 10,000 points and offered a sequence of ten binary choices between a 50:50 lottery and a sure amount. Two types of lottery were used. The first had a 50% chance of 0 points, and a 50% chance of winning a (varying) positive amount of points (of up to 10,000). The

[17]The attrition rate of ≈25% is lower than most online surveys. This is due, in part, to YouGov's panel management, and in part to the large incentives we offered. According to Pew Research, YouGov's sampling and weighting procedure yields better representative samples than traditional probability sampling methods with non-uniform response rates, including Pew's own probability sample (Pew Research Center, 2016, YouGov is Sample I).

[18]We chose to pay two randomly selected questions to increase the stakes while making fewer participants upset about their payoffs. Paying for two questions instead of one may theoretically induce some wealth effects, but these are known to be negligible, especially in an experiment such as ours (Charness et al., 2016). Paying for randomly selected questions is incentive compatible under Expected Utility, but not necessarily under more general risk preferences, where it is known that no such mechanism may exist (Karni and Safra, 1987; Azrieli et al., 2018). An old and still growing literature suggests this theoretical concern may not be empirically important (Beattie and Loomes, 1997; Cubitt et al., 1998; Hey and Lee, 2005; Kurata et al., 2009), but there are some exceptions (Freeman et al., 2015). Dynamic designs are generally not incentive compatible, however in practice this is of little concern—see Appendix A.

[19]The conversion from points to awards can only be done at specific point values, which leads to a slightly convex payoff schedule. This is of little concern here as these cash-out amounts are further apart than the maximum payoff from the survey.

second had a 50% chance of winning an amount up to 10,000 points, and a 50% chance of a loss of up to 10,000 points. In the latter case, the sure amount was always 0 points.

Time Preferences: The second module elicited discount factors and refined estimates of the curvature of the utility function. Participants were offered a sequence of ten binary choices between a lower amount of points at an earlier date (either the day of the survey, or in the future) or a higher amount at a later date (up to 90 days in the future). The maximum payment in each question was 10,000 points.

2.3.3 Additional Measures

The survey also contained more standard ways of eliciting risk and time preferences. These serve as a useful comparison for DOSE.

Risk Aversion MPLs: Two MPLs asked participants to choose between a fixed 50/50 lottery and a series of ascending sure amounts. The row in which the participant first chose the sure amount identified a range of possible certainty equivalents for the lottery—we use the midpoint of this range. There were two MPLs of this type: the first had a 50/50 lottery over 0 and 10,000 points, the second, a 50/50 lottery over 2,000 and 8,000 points.[20]

Time Preference MPLs: In addition to the DOSE module, the survey included two MPLs to elicit time preferences. The first time MPL elicited the amount of points that the participant valued the same as 6,000 points 45 days later. The second MPL elicited the amount of points in 45 days that the participant valued the same as 6,000 points in 90 days. This measure is used primarily in Section 4.4.

[20]See Appendix D.1 for more details on these measures, and additional analyses using them.

3 Performance of DOSE versus Current Methods

We first demonstrate that DOSE works in simulated environments, before turning to our results on loss aversion in a representative sample in the next section. We do so in two ways. First, we simulate DOSE question selection using the choices of participants in two laboratory experiments that used questions similar to those in our DOSE implementation. We show that a 20-question DOSE procedure obtains parameter estimates that are close to (within 15% of) parameter estimates after 140 questions. Second, we use simulated participants to show that, in the presence of mistakes or inconsistency, DOSE recovers estimates that are at least twice as accurate as standard elicitations used to measure risk and loss aversion.

3.1 Simulating DOSE with Laboratory Data

The first simulation exercise demonstrates the benefits of DOSE's personalized question sequence. We use data from 120 student participants in two prior laboratory experiments.[21] In each experiment, participants were asked the same set of 140 binary choices from the same two types of questions described in the prior section. The order of these questions in the experiments was somewhat random. In our simulation, we optimally order these question for each participant using DOSE. After DOSE selects a question, we provide it with the answer the participant gave in the experiment. The procedure then updates the probability distribution over parameters, selects the next question, and so on. This allows us to compare, question by question, the *inaccuracy*—the absolute distance from the true parameter value as a percentage of the true value—of DOSE's estimates with those elicited by a random question ordering. As we do not have access to true parameter values, we substitute the values one would obtain using the choices in all 140 questions.

A 20-question DOSE sequence provides a similar amount of information as about 50

[21] 90 participants come from Frydman et al. (2011) and 30 from Sokol-Hessner et al. (2009). We attempted to compare the performance of DOSE using Maximum Likelihood Estimation (MLE), the method in Sokol-Hessner et al. (2009) and Frydman et al. (2011). However, as reported in Appendix G, we were unable to obtain MLE estimates for a large portion of the sample. The MLE estimates that were obtained were less accurate—relative to the estimate after 140 questions—than those obtained from Bayesian estimation.

Figure 2: Optimal question selection rapidly leads to accurate estimates.

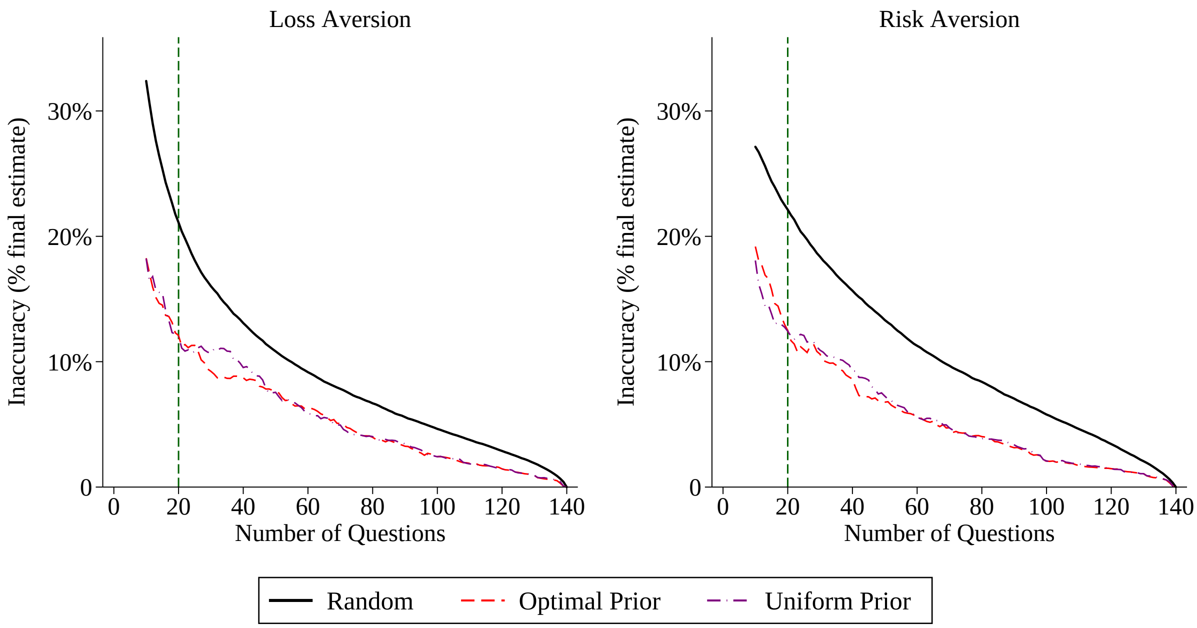

Notes: Based on data from Sokol-Hessner et al. (2009) and Frydman et al. (2011). Each line shows the inaccuracy of Bayesian estimates (with uniform initial prior) obtained after each question, starting at question 10, under different orders. "Optimal Prior" and "Uniform Prior" refer to DOSE question selection using corresponding priors. "Random" orders questions randomly, averaging over 100 different random orderings.

randomly ordered questions, as shown in Figure 2.[22] The DOSE estimates of both risk and loss aversion are consistently closer to the final parameter estimate, indicating—under the assumption that the final estimate closely approximates an individual's true parameters—that the procedure provides accurate estimates considerably faster than selecting questions at random.[23] After 20 questions, the DOSE estimates are almost twice as close to the final estimate as those under a random question ordering (12% vs. 21–22%). The DOSE estimates are also more highly correlated with the final estimates (shown in Appendix Figure G.1), an important feature when seeking to identify correlations between preferences and other population characteristics.[24]

[22]For loss aversion, 45 randomly-ordered questions are needed to be as close to the final estimate as 20 DOSE questions. For risk aversion, 55 questions are required.

[23]Supporting this assumption, the next subsection finds that DOSE achieves similar levels of accuracy in a simulation where we know the true parameter values.

[24]Further, the average benefits we estimate are not limited to the particular distribution of preferences we observe in the laboratory. As we show in Appendix G, the DOSE estimates converge rapidly to the final estimate for the entire range of λ and ρ.

These simulations also show that using the uniform prior is close to optimal for question selection. To do so, we compare the performance of DOSE question selection using a uniform prior to that using an *optimal prior* constructed from the distribution of the estimates after 140 questions. To focus on the question selection impacts of the prior, we estimate the parameter values using a uniform prior in both cases. As shown in Figure 2, the accuracy is similar whether using the optimal or uniform prior.

3.2 Parameter Recovery Study

When participants make mistakes DOSE produces estimates that are about twice as accurate as traditional risk and loss aversion elicitation mechanisms. We demonstrate this with a parameter recovery study (or Monte Carlo simulation). This is conducted with an entirely simulated dataset that allows us to both know and control the true parameters governing (simulated) participant behavior. Appendix E provides full details.

We evaluate the relative (in)accuracy of DOSE, and two other common risk elicitation procedures, using 10,000 simulated participants with power utility given in (1), who make binary choices probabilistically according to (2). For each of these participants, a set of parameter values (ρ_i, λ_i, and μ_i) are drawn from the posterior distribution obtained from the 120 laboratory participants in the previous subsection. We obtain DOSE estimates by running 10- and 20-question DOSE procedures for each simulated participant, determining the answer to each question according to her parameters.

As a benchmark for DOSE, we also allow our simulated participants to make choices in two other common risk elicitation methods: the Lottery Menu (Eckel and Grossman, 2002), and a double MPL (Andersen et al., 2008a; Andreoni and Sprenger, 2012b). In the Lottery Menu, participants choose from a list of six 50/50 lotteries over gains. In the double MPL, participants complete two MPLs, each offering a choice between a fixed 50/50 lottery and a series of ascending sure amounts. The first—which identifies risk aversion—offers a lottery over gains ($0 and $10), while the second—identifying loss aversion—offers a lottery

Table 1: DOSE produces more accurate estimates.

	Average Inaccuracy	Spearman Rank Correlation with True Value
Loss Aversion		
DOSE 10 question	21%	0.86
DOSE 20 question	15%	0.91
Multiple Price List	36%	0.65
Risk Aversion		
DOSE 10 question	21%	0.66
DOSE 20 question	15%	0.79
Multiple Price List	37%	0.45
Lottery Menu	35%	0.28

Notes: Inaccuracy is the absolute distance from the true parameter value as a percentage of the true value.

between a gain and a loss, both of $10. In both cases, we define a probability distribution over the possible choices using sequential pairwise comparison of the options with the same logit choice function in (2). This probability distribution is used to calculate the expected inaccuracy of the parameter estimate for each simulated participant.

The estimates of risk and loss aversion from DOSE are approximately twice as accurate as those from the other elicitation procedures, as shown in Table 1. After 20 questions, DOSE obtains estimates of risk and loss aversion that are, on average, within 15% of the true parameter value. The average inaccuracy of the MPL and Lottery Menu procedures, in contrast, is at least 35%—much higher than even a 10-question DOSE procedure.[25]

DOSE produces more accurate estimates than both other procedures regardless of participants' level of choice consistency (μ), as shown in Figure 3. This figure repeats the parameter recovery analysis above, but assigns all simulated participants the same level of choice consistency, μ. We then vary μ across percentiles of the population distribution. The 20-question DOSE procedure always provides the most accurate estimates. Even the 10-

[25]The improvement in accuracy from DOSE is similar when the utility function used in the question selection procedure is misspecified—see Section 5.1 and Appendix F.

Figure 3: DOSE is more accurate than other methods at all levels of choice consistency.

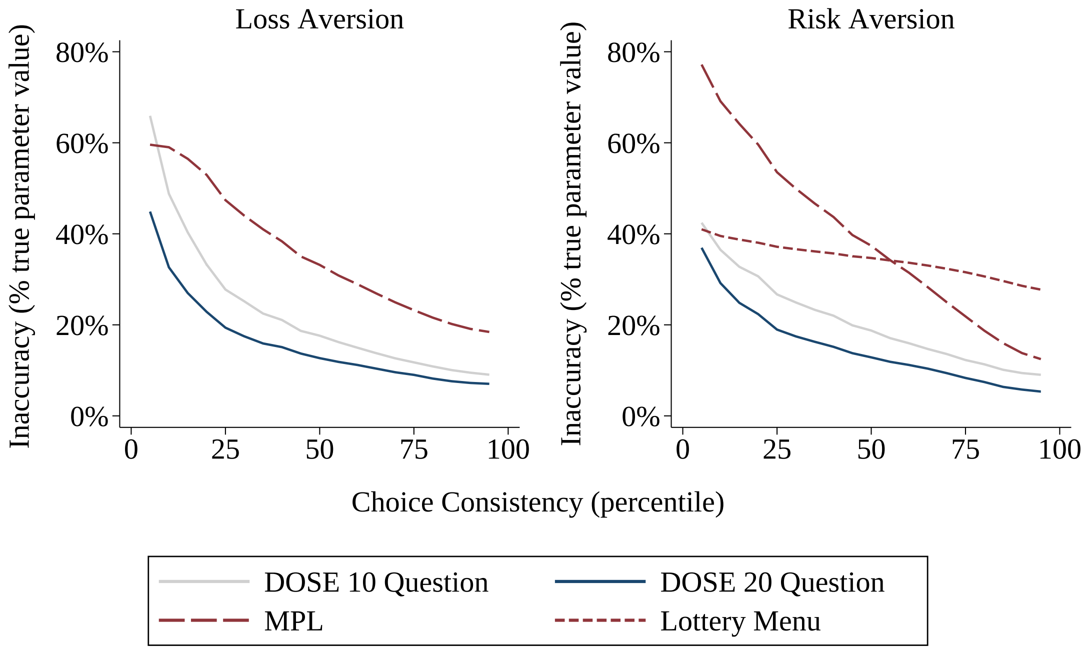

Notes: Estimates obtained using simulation procedure described in Appendix E, with all simulated participants at a point in the graph having the same value of μ.

question procedure performs better than either the MPL or Lottery Menu, except amongst extremely inconsistent participants.[26]

The high accuracy of the DOSE estimates also leads to higher correlations with the true parameter values than the other two procedures (column 2 of Table 1). Thus, DOSE is less likely to miss associations between economic preferences and other characteristics through attenuation bias. The correlation between the true risk aversion parameter and the DOSE estimate is 0.79, compared to 0.45 with the MPL estimates and 0.28 for the Lottery Menu. For loss aversion, the DOSE procedure produces correlations above 0.85 with the true values, even after a 10-question procedure. This is reflected in our survey results, see Section 4.3.

Unlike the MPL, DOSE is able to elicit loss aversion estimates even when participants' choices violate First Order Stochastic Dominance (FOSD), although this is not an important factor in the simulation results of Table 1. Because DOSE accounts for the possibility that

[26] Although the Lottery Menu procedure appears to perform better than the MPL for inconsistent participants, this advantage is not robust to alternative simulation assumptions, which can drive the average inaccuracy for low consistency participants as high as 139%. See Appendix E.1 for further details.

a participant's choice is a mistake, the procedure can always recover parameter estimates. In the double MPL, on the other hand, participants may erroneously make choices on the second MPL (used to elicit loss aversion) that are First Order Stochastically Dominated given their choices on the first MPL (used to elicit risk aversion). This prevents estimation of the loss aversion parameter. In our simulation, the MPL could not recover estimates for 11% of participants—increasing to more than 50% of participants with low choice consistency.

In practice, the double MPL procedure is unable to elicit loss aversion for a significant proportion of the population, which may lead to biased conclusions about loss aversion. In particular, the double MPL used by Chapman et al. (2018) is unable to recover estimates for 37% of their participants, as measurement error and other factors lead to these participants appearing to make FOSD choices. The prevalence of these choice patterns was not random: loss aversion could only be computed for 50% of participants in the bottom quartile of cognitive ability, compared to 70% in the upper quartile. The results in this paper (see Section 4.2) indicate that high cognitive ability is associated with more loss aversion. This pattern of missing observations may thus lead to a biased over-estimate of loss aversion.[27]

In practice, choice data in the MPL is likely noisier than our simulations assume. To estimate the relative amount of noise in the survey, we compare simulated and real responses for three additional MPLs—the double MPL procedure (but with different payoffs), and a second risk aversion MPL as implemented in Chapman et al. (2018).[28] The proportion of FOSD responses in the loss aversion MPL is much lower in this simulation than the real data: 20% rather than 37%. Further, the correlation between the certainty equivalents in the risk aversion MPLs—which is higher in the presence of less measurement error (Gillen et al., Forthcoming)—is higher in the simulation: 0.73 vs. 0.69.[29]

[27]The survey in Chapman et al. (2018) is similar to the one in this paper, although it did not utilize DOSE. Note this pattern of responses is consistent with results in the next section, as loss-tolerant individuals will wish to make choices that are close to violating FOSD, and measurement error can push them over the threshold. Low cognitive ability participants are more loss tolerant and make more FOSD choices.

[28]These MPLs have a different structure from those used in the results in Table 1 and Figure 3, but the simulation methodology is the same. Full details of the simulations are reported in Appendix E.

[29]It appears that our simulations underestimate the measurement error in the survey MPLs because we do not account for participants' use of rules-of-thumb. Compared to our simulation, participants in the

4 Economic Preferences in a Representative Sample

The U.S. population is more loss tolerant than those in lab-based samples. Consistent with this finding, higher cognitive ability participants are more loss averse. This contrasts with the prior literature that suggests that those with higher cognitive ability are more "rational," or, as it applies here, more likely to make expected value maximizing choices. DOSE estimates of risk aversion, discounting, and choice consistency, are in-line with this perspective: higher cognitive ability participants are less risk averse and more patient. The literature, however, has found mixed results about the relationship between risk aversion and cognitive ability. We use DOSE estimates of choice consistency (μ) to show that these results may be driven by inconsistent choice. In particular, we show that if we examine those with above-median choice consistency we recover a correlation between MPL-based measures of risk aversion and cognitive ability that is obscured when examining all participants.

Further, prior estimates of the cross-time stability of economic preferences may be understated. Cross-time consistency of DOSE estimates of risk aversion and discounting are higher than both MPLs and prior studies. The cross-time stability of loss aversion, which has not been previously measured, is comparable to that of risk aversion and discounting, suggesting that loss attitudes are at least as stable a descriptor of preferences as these other, more common, measured preferences.

The results in this section are presented under the assumption, driven by the results in Section 3, that DOSE is capturing useful information about economic preferences. Although some of the analyses in this section provide further support for this assumption, we do not examine it in detail until the next section. There, we show that our results are robust to adjusting for a number of possible issues that may be unique to DOSE, such as misspecification of the utility function used to select questions and analyze the resulting choices.

survey were more likely to switch in rows of the MPL that are especially salient—such as the first or last rows, or those referring to the midpoint of the lottery. These choices may be capturing framing effects or heuristics in the face of the large amount of information in an MPL. DOSE avoids these issues by using simple binary choices that are likely simpler to understand—a claim supported by evidence in Section 5.2. Thus, our simulations may actually underestimate the relative advantages of DOSE.

4.1 Loss Aversion in the U.S. Population

We find far more loss tolerant participants, and thus a lower average level of loss aversion, in the general population than in lab-based samples that have used DOSE, as displayed in Figure 4. While the distribution of parameter estimates from DOSE are largely the same across both waves of our incentivized survey, those produced by lab-based samples are markedly different. This is also true for risk aversion: lab-based populations are less risk averse than the general population, in-line with prior research (see Snowberg and Yariv, 2018, and references therein).[30] We compare estimates of risk aversion in this study and prior studies when examining robustness in Appendix B.

The median estimate of the loss aversion parameter, $\lambda = 0.98$, in the U.S. population is much lower than the "standard" estimate of 2 (Fehr-Duda and Epper, 2012). However, in-line with prior studies, DOSE in the lab ($N = 439$) produces a median estimate of $\lambda = 1.99$. Our lab results come from studies using DOSE in much the same way we do here, based on our original working paper. For details on three of these studies, see Clay et al. (2017, in progress) and Krajbich et al. (2017). The fourth study is unpublished, and ran the DOSE procedure on 207 students at UCLA.[31] We use the individual choices from these studies to calculate parameter estimates as described in Section 2.

The proportion of loss-tolerant participants in the U.S. population—53%—is higher than in the eight studies, referenced in the Introduction, that have investigated heterogeneity in loss aversion (all in lab samples). In those experiments, between 13% and 30% (weighted average: 22%, $N = 1{,}023$) of participants in laboratory experiments are loss tolerant. However, as noted above, the methodologies used in several of these studies make classifying many participants impossible. Our representative sample also produces different results than other

[30] The median CRRA coefficient (1-ρ) in the general population is 0.31 vs. 0.05 in the student/lab sample. We do not have DOSE estimates of the discount rate in the lab. However, the median monthly discount factor here (0.90) is in the lowest quartile of the results of three recent laboratory studies using the Convex Time Budget method of Andreoni and Sprenger, (2012a; see Appendix Table D1 in Imai and Camerer 2018). The distribution of both the discounting and choice consistency measures are displayed in Appendix Figure D.1.

[31] The data was generously provided to us by Alec Smith.

Figure 4: Distribution of Economic Preferences within the U.S. population

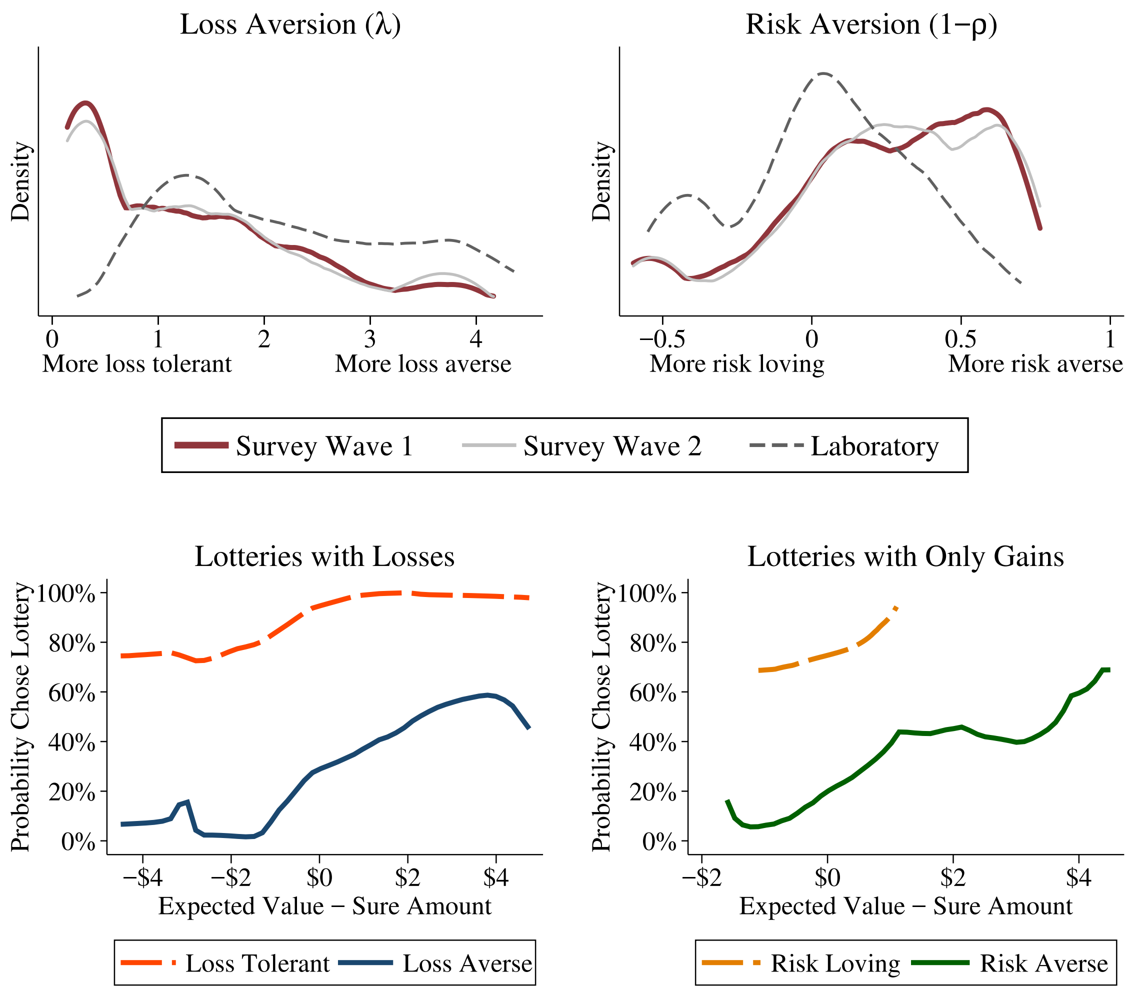

Notes: The top panel displays the kernel density of each parameter, plotted using Epanechnikov kernel with bandwidth chosen by rule-of-thumb estimator. The bottom panel displays the Nadaraya-Watson (local mean smoothing) estimator (bandwidth 0.6) with Epanechnikov kernel and without sample weights.

field studies, discussed in the literature review. Two of those studies obtain estimates for less than 30% of their participants (Booij and Van de Kuilen, 2009; Booij et al., 2010). A third produces estimates of the median level of loss aversion between 0.12 and 4.47 depending on the specification (von Gaudecker et al., 2011). As we show in Section 5.1 and Appendix F, our results are much more stable under different estimation specifications.

Choice patterns clearly illustrate the source of DOSE estimates, as shown in the bottom panel of Figure 4. The x-axis is the difference between the expected value of a lottery and the sure amount in a given choice. Loss-tolerant participants ($\lambda < 1$) are clearly more likely to choose lotteries with losses than those who are loss averse ($\lambda > 1$), with loss-tolerant participants choosing lotteries nearly 100% of the time when the expected difference is zero. Note, however, that the flat parts of both lines in the left-hand side of the bottom-left panel are due to the fact that DOSE only exposes those who have already revealed loss tolerance through prior choices of lotteries with large negative expected values. Similar patterns exist for those who are risk averse versus those who are risk loving: the latter are more likely to choose gambles with gains at every expected value difference. For all four groups of participants, the probability of choosing the lottery increases with the difference between the expected value of the lottery and the sure amount.

4.2 Economic Preferences and Cognitive Ability

Our comprehensive survey allows us to document new facts about the correlates of loss aversion and choice consistency in the U.S. population. An examination of the simple correlations between economic preferences, socioeconomic characteristics, and cognitive ability shows that cognitive ability is the most important correlate of loss aversion, and the other three DOSE-estimated parameters. High cognitive ability participants are more loss averse, while those of lower cognitive ability are more loss tolerant, on average. This correlation reflects clear differences in the choices participants made during the survey: high cognitive ability participants were consistently less likely to choose lotteries involving losses. Higher cognitive ability participants are more patient (and consistent), in-line with previous studies. Higher cognitive ability participants are also less risk averse. Examining this result in further detail in the following subsection allows us to demonstrate that the mixed evidence on the relationship between risk aversion and cognitive ability in previous studies may be explained by inconsistent choice (Andersson et al., 2016b; Dohmen et al., 2018).

Cognitive ability was measured using a set of nine questions. Six questions were from the International Cognitive Ability Resource (ICAR, Condon and Revelle, 2014): three were

Table 2: DOSE preference parameters are correlated with individual characteristics.

	Loss Aversion (λ)	Risk Aversion ($1-\rho$)	Patience (δ)	Choice Consistency (μ)
Cognitive Ability	0.21***	−0.21***	0.18***	0.15***
	(.030)	(.028)	(.029)	(.026)
Income	0.15***	−0.15***	0.12***	0.06*
	(.032)	(.034)	(.034)	(.033)
Education	0.13***	−0.10***	0.17***	0.11***
	(.032)	(.033)	(.037)	(.032)
Male	0.08**	−0.10***	−0.02	0.01
	(.033)	(.032)	(.035)	(.033)
Age	−0.10***	0.02	0.18***	0.05
	(.033)	(.032)	(.036)	(.036)
Stock Investor	0.06**	−0.11***	0.10***	-0.02
	(.031)	(.029)	(.031)	(.032)

Notes: ***, **, * denote statistical significance at the 1%, 5%, and 10% level. Standard errors, in parenthesis, come from a standardized regression. Each cell corresponds to a single regression.

similar to Raven's Matrices, and the other three involved rotating a shape in space. We also administered the Cognitive Reflection Test (CRT; Frederick, 2005): three arithmetically straightforward questions with an instinctive, but incorrect, answer. Our cognitive ability score was the sum of correct answers to these questions.

The relationships in Table 2 are consistent with the analyses in the prior subsection showing that the general population is more risk averse and less loss averse than lab/student populations. In particular, more educated, higher income, and more cognitively able individuals tend to be more loss averse and less risk averse; and lab populations have higher cognitive ability than the general population (Snowberg and Yariv, 2018). Men and younger people tend to be more loss averse, and those that own stock are less loss averse.[32]

The strong correlations between cognitive ability and economic preferences are robust

[32] Appendix Table D.2 presents additional correlations with Church Attendance, Ethnicity, and Home Ownership, and shows that the correlations with the two components of cognitive ability (CRT and IQ) are similar to the correlations in Table 2.

Figure 5: Low cognitive ability participants chose more lotteries with losses.

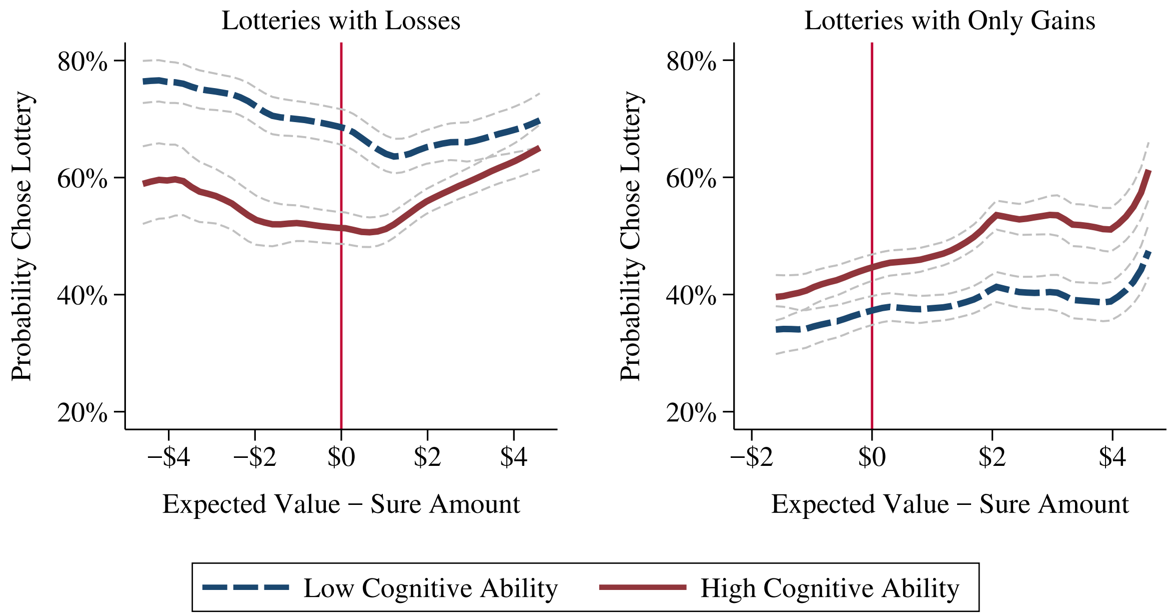

Notes: Figure displays the Nadaraya-Watson (local mean smoothing) estimator (bandwidth 1) with Epanechnikov kernel. Grey dotted lines represent 95% confidence intervals, constructed with 10,000 clustered bootstrap replications. High and low cognitive ability refer to the top and bottom terciles, respectively.

to controlling for the other individual characteristics in Table 2—see Appendix Table D.3. In fact, differences in cognitive ability appear to explain most of the relationship between education and economic preferences.

The choices participants make differ by cognitive ability, as shown in Figure 5.[33] Across the range of expected value differences, high and low cognitive ability participants exhibit different patterns of choice. In the first panel, which focuses only on lotteries with a loss, low cognitive ability participants are significantly more likely to choose the lottery than high cognitive ability participants. The u-shape of the curve for both ability terciles is driven by the fact that DOSE only presents very negative expected value difference choices to those who have already expressed significant loss tolerance. In contrast, in the second panel, which focuses on lotteries that only contain a zero payoff and a gain, low cognitive ability participants are significantly less likely to choose the lottery.

In summary, the patterns of correlation between cognitive ability and risk and loss aver-

[33] Appendix Figure D.1 presents the results in Figure 4 by cognitive ability tercile.

sion in Table 2 are clearly driven by underlying choices. Low cognitive ability participants are especially willing to accept lotteries with losses, even when those result in an expected value loss. However, low cognitive ability participants are also less willing to choose a lottery over gains, even when that results in an expected value gain.

Very few participants consistently make expected value maximizing choices, regardless of cognitive ability. Fewer than 2% of participants made all EV-maximizing choices, and fewer than 5% made more than 8 such choices (out of ten). Further, in contrast to some previous studies (for example, Burks et al., 2009; Benjamin et al., 2013), we find the proportion of choices that maximize expected value is only slightly higher for high cognitive ability participants: those in the highest tercile of cognitive ability made 57% EV-maximizing choices compared to 52% for participants in the lowest tercile.

4.3 Choice Consistency and Estimate Accuracy

Accounting for inconsistent choice can explain the mixed evidence about the relationship between risk aversion and cognitive ability in previous studies (Dohmen et al., 2018). The simulations in Section 3.2 show that, in the presence of inconsistent choice, MPLs measure risk aversion with considerable error. It is well known that error will attenuate, and potentially bias, any estimated relationship between these measures and other factors. In this subsection, we show that inconsistent choice is related to attenuation bias in our survey. The MPL measure of risk aversion on our incentivized survey is weakly associated with cognitive ability, in contrast to the DOSE measure. However when we focus only on participants that make (more) consistent choices, the MPL and DOSE measures exhibit similar correlations.

The MPL-based risk aversion measure is more weakly correlated with other characteristics than the DOSE measure, as shown in the first panel of Figure 6. For example, the correlation with cognitive ability is -0.04 (s.e. = .028), compared to -0.21 (.028) for DOSE. This pattern is consistent with our simulation results, which showed more error in MPLs than DOSE.

Inconsistent choice is related to the attenuation of correlations. Once we use the DOSE

Figure 6: DOSE measure of risk aversion is more highly correlated with individual characteristics before choice consistency is accounted for.

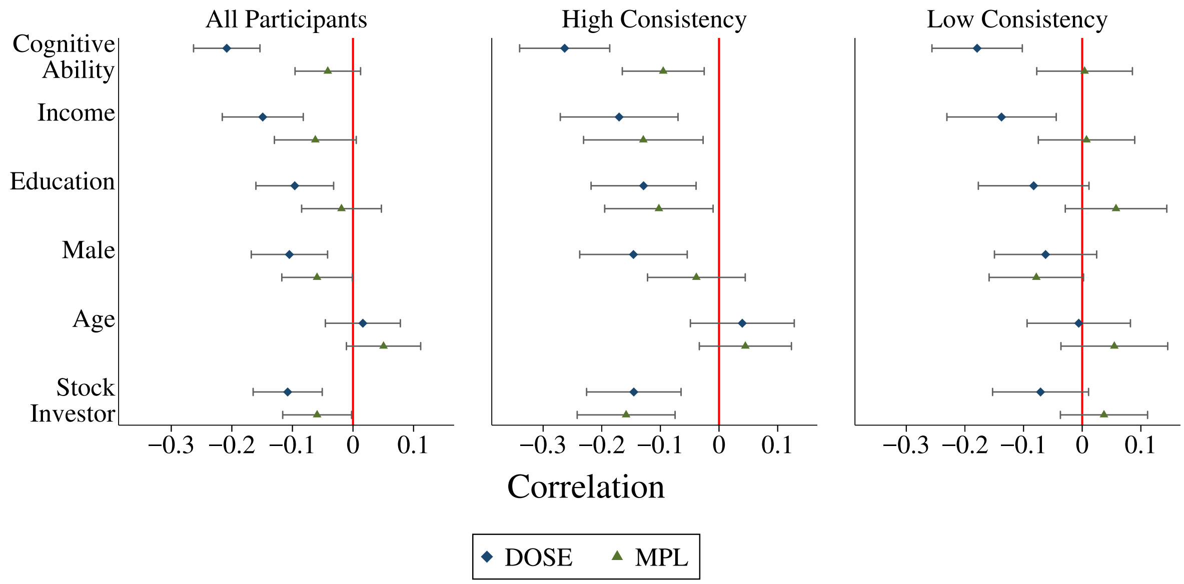

Notes: Figure displays correlations between the DOSE and MPL measures of risk aversion and individual characteristics. The left-hand panel includes all participants, the middle contains those with above median choice consistency, and the right-hand panel contains those with below median choice consistency. The survey contained two MPL measures of risk preference. Correlations are estimated by stacking the two and clustering standard errors by participant.

consistency measure to exclude inconsistent participants, there is a strong negative relationship between cognitive ability and the MPL risk aversion measure—see the middle panel of Figure 6, which contains only those with above median choice consistency parameters μ.[34] The magnitude of the correlations is consistently higher for both risk aversion measures; however the contrast is particularly striking for the MPL measure, where a number of relationships—including with cognitive ability—are now statistically significant. This is despite the fact that standard errors are increased by only using half the sample. As shown in the right-hand panel of the figure, DOSE estimates exhibit similar correlations even for very inconsistent participants, while correlations with MPL estimates are almost zero. This is also in-line with our simulation results.

The patterns in Figure 6 also demonstrate that the choice consistency parameter can

[34] As we discuss in Section 5.2, DOSE estimates may also contain less measurement error because the binary choice questions are easier to understand than MPLs.

identify individuals that make more mistakes even outside of the DOSE module, and thus help researchers address survey noise. This information is difficult to obtain through other easily available measures. For example, as we demonstrate in Appendix D.2, the correlations in Figure 6 cannot be recovered by truncating the sample based on response time rather than consistency. In fact, the consistency measure helps distinguish whether fast responses reflect a lack of attention: restricting the sample to high-consistency participants recovers correlations even among the subgroup of participants with particularly fast response times.

The value of the choice consistency measure is particularly striking when we consider that our DOSE design did not focus primarily on eliciting this measure. Simple design tweaks—such as allowing the procedure to ask questions multiple times—could allow the variable to be measured more accurately, and hence provide even more information to researchers.

4.4 Within-person Stability of Loss Aversion

The within-person stability of the DOSE estimates of risk and time preference is higher than other behavioral elicitations, both in our survey and in most previous studies. That is, they are more highly correlated within-person across time, consistent with the fact that DOSE reduces measurement error in parameter estimates. The correlation of DOSE estimates across survey waves was 0.40 (s.e. = .04) for loss aversion (λ), 0.45 (.04) for risk aversion (ρ), and 0.47 (.05) for discounting (δ). In comparison, the inter-temporal correlation between choices in the two risk MPLs were 0.29 and 0.26 (.04 for both), and for choices in a risky project measure (Gneezy and Potters, 1997) the correlation was 0.33 (.04). The stability in the two time preference MPLs was 0.28 and 0.20 (.06 for both).[35] These findings are consistent with higher measurement error in the MPL measures, as suggested by both the simulation results (Section 3.2) and the survey results discussed in the previous subsection.

Interestingly, there is no clear relationship between the stability of the DOSE estimates of

[35] The stability of the consistency parameter (μ) was 0.22 (S.E.=.05); lower than the other DOSE measures but similar to the MPLs discussed above. Part of the explanation for this relatively low correlation is that our DOSE implementation was designed to update more on other parameters: the relatively small number of questions made it harder to identify inconsistent choices.

preference parameters and choice consistency (μ), suggesting that by accounting for mistakes DOSE blunts the impacts of them. For participants with above-median choice consistency the over time correlations are 0.41 for loss aversion, 0.43 for risk aversion, and 0.47 for discounting. For participants with below-median choice consistency correlations are 0.38, 0.47, and 0.46, respectively. This is consistent with evidence in the prior subsection and section that the accuracy of DOSE estimates is relatively constant across much of the observed range of choice consistency—except for the most inconsistent.

The over-time correlations of DOSE estimates also compare favorably with methods in prior studies. The only study we are aware of that measures stability of risk attitudes in the loss domain is Levin et al. (2007), who report over-time correlations for 62 participants of 0.29 for a risk measure over gains, and 0.20 for a measure of differential risk-taking between the gain and loss domain.[36] Two further studies use incentivized methods to investigate the stability of risk aversion (over gains) over lengthy periods, both finding lower over-time correlations than the DOSE estimates. Gillen et al. (Forthcoming) find an inter-temporal correlation of 0.32 for both of two risk MPLs and 0.36 and 0.47 for two risky project questions. Lönnqvist et al. (2015) report a within-subject correlation of 0.21 for an MPL measure across a year.[37] There is a similar pattern when comparing the stability of the DOSE-measured time preferences to the prior literature, although differences in methodology and samples make it harder to compare (see discussion in Appendix B). Thus, DOSE is an answer to Meier and Sprenger's (2015, p. 286) challenge to develop, "A more precise experimental technique for eliciting time preferences...to make further study of stability."

Notably, the DOSE estimates are more stable, despite these previous studies occurring in the laboratory or with high-IQ college populations (or both)—suggesting that DOSE can obtain levels of measurement error similar to a laboratory environment in an online survey. Economists have often shied away from using incentivized measures in large samples because

[36]Levin et al. (2007) report correlations for 62 pairs of parents and children. The figures above are from the adults, for comparability. For the children, over-time correlations are 0.38 and 0.30 respectively.

[37]Andersen et al. (2008b) elicit risk aversion over time, however, they do not report over-time correlations.

of high measurement error and the prohibitive cost of implementing multiple elicitations (Schildberg-Hörisch, 2018). DOSE overcomes this constraint and, as we show in Section 5.2, the procedure is faster to complete than an MPL module.

5 Robustness

In this section we use our data to demonstrate the robustness of DOSE parameter estimates, providing further evidence for our finding that a significant proportion of the U.S. population is loss tolerant. DOSE obtains accurate parameter estimates even if the utility function used in the question selection procedure is misspecified. Consequently, it is straightforward to test robustness to different utility specifications: our survey conclusions are unchanged assuming either exponential (CARA) utility or allowing curvature to vary across gains and losses. Further tests suggest that our results are not driven by inattention or misunderstanding.

5.1 Robustness to Misspecification

DOSE is largely robust to misspecification of the parametric form used to select questions because the choice data itself can be used to estimate the parameters of different functional forms. This re-estimation can be done in response to new information, or simply as a robustness check. In this subsection, we focus on the latter application, and assess the robustness of our loss aversion results to different forms of the utility function over gains and losses. Our estimates are largely robust. This should be unsurprising based on the patterns of choice described in Sections 4.1 and 4.2.

Simulations indicate that using a misspecified utility function produces little change in our results. In Appendix F.1, we run the DOSE question selection procedure with a exponential (CARA) utility function for the simulated participants used in Section 3.2. As the simulated participants make choices using a power (CRRA) utility function, this means that the utility function used to select questions is misspecified. The Spearman rank correlation between the (misspecified) CARA risk aversion parameter estimates and the (true) CRRA parameters is almost exactly the same as when using the correct utility function–0.75 (versus 0.79) for risk aversion and 0.90 (versus 0.90) for loss aversion. Moreover, we are able to recover equally accurate CRRA parameter estimates using the choice data. Even when misspecified in this way, DOSE outperforms other measures by similar margins to those shown in Table 1.

Applying this procedure to our survey data, in Figure 7, shows that our conclusions regarding loss aversion are similar using different utility functions. In addition to the CRRA utility function used in (1), we add a CARA utility function, and a CRRA utility function with different curvature parameters over losses and gains (Tversky and Kahneman, 1992). The latter specification produces the biggest difference. However, the finding that a large portion of the population are loss tolerant is unchanged: the proportion of participants with estimated $\lambda < 1$ ranges from 46% to 53% across the three models. The variation in the median value of the loss aversion parameter λ is very narrow across the three models, ranging from 0.98 to 1.12.[38] The second panel of Figure 7 breaks down both the one- and two-parameter CRRA utility parameters by cognitive ability. Again, there is the same pattern as in the prior section: those with higher cognitive ability are more loss averse, and those with lower cognitive ability are more loss tolerant.

The CRRA model we used in prior sections fits the data best: it predicts 89% of choices correctly, compared to less than 85% for the other two functions. Moreover, we estimate that most (68%) participants are risk averse over gains and risk loving over losses, in-line with prior experiments and Prospect Theory (Kahneman and Tversky, 1979). The average difference in the curvature between the domains is close to zero, offering support for specification (1), although there is considerable individual heterogeneity—see Appendix F.

The individual choice data demonstrate widespread loss tolerance without making any parametric assumptions. A loss-averse participant should never accept a lottery with neg-

[38]The correlation between the loss aversion parameters under different utility functions is greater than 0.8. Risk aversion over gains is also highly correlated across the different specifications (Spearman correlations of 0.98 or above).

Figure 7: Results on loss aversion are robust to different parametric specifications.

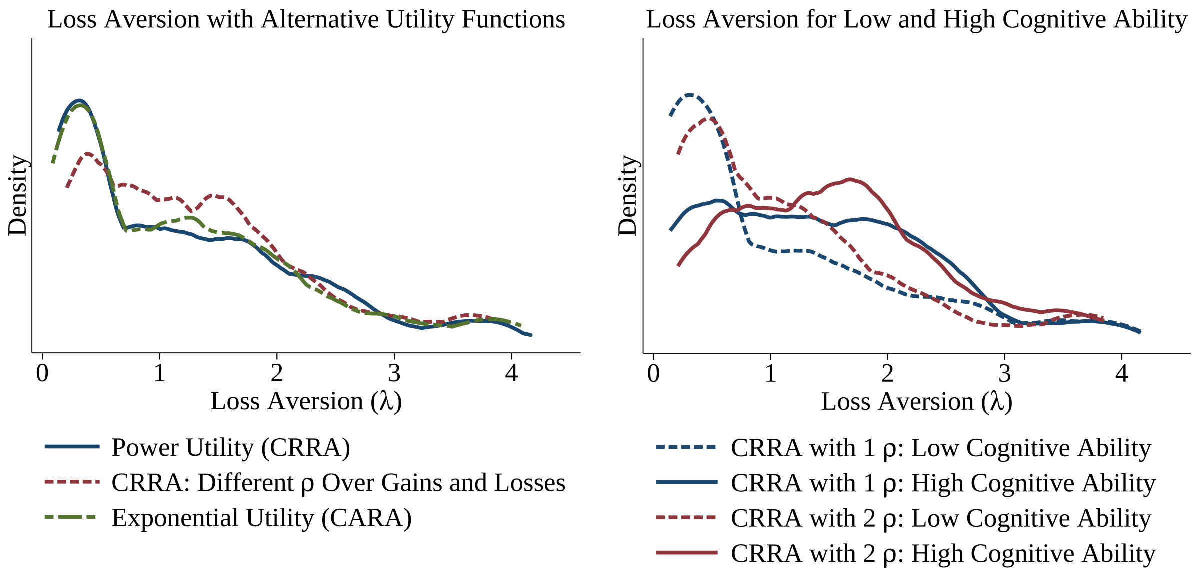

Notes: Plotted using Epanechnikov kernel, with bandwidth chosen by rule-of-thumb estimator. Figure displays the estimated loss aversion parameter after re-estimating the model parameters with alternative functional forms for the utility function—see Appendix F for full details. "Power utility (CRRA)" does not allow for differential curvature over gains and losses—see (1), whereas "CRRA: Different ρ over gains and losses" allows for this possibility.

ative expected value in our implementation, as they could always choose a certain amount of $0 instead. Yet, many participants do, as shown in Figure 4. Accepting a lottery with a negative expected value implies that the magnitude of the negative value for a loss is less than that for an equally-sized gain—implying loss tolerance. The fact that a large proportion of the population is loss tolerant is thus apparent without any parametric assumptions.

The results in this subsection clearly demonstrate that the DOSE estimates reflect participants' choices. However, they cannot speak to the extent to which those choices accurately reflect individual preferences or, specifically, whether participants comprehended or paid attention during the DOSE module. Thus, in the next subsection, we address concerns that our results are driven by a lack of understanding or care in completing the survey.

5.2 Response Times

The questions we used in DOSE appear to be easy to understand. First, respondents across the cognitive ability spectrum complete DOSE questions in similar amounts of time, as opposed to more complicated methods, which take lower cognitive ability participants longer. Second, the speed of response is not evidence of participants "giving up," or misunderstanding the questions. Removing participants that complete DOSE more quickly has no effect on estimates of the distributions of parameters.

Earlier results provide some evidence that participants did, in fact, find DOSE relatively simple. The higher correlation of DOSE elicitations across time (compared to MPL elicitations) suggests that it was easier to maintain similar response patterns even more than six months later. Moreover, the fact that correlations between DOSE estimates of risk aversion, discounting, and sociodemographic characteristics largely mirror the existing literature, and are stronger than the correlations with MPL measures, also suggests that they more reliably extract information on preferences. However, the fact that lab and general population results differ could be interpreted as evidence of possible confusion in the representative sample.

As a reminder, we designed our questions in order to make expected value comparisons as simple as possible, as discussed in Section 2.2.2. In particular, all questions contained only one 50/50 lottery, and at least one of the three payoffs (the sure payoff, and the two lottery payoffs) was zero. This implies that computing the expected value of a gain lottery only required dividing the lottery by two. Comparing a gain/loss lottery to a sure payoff of zero required only comparing the magnitude of the gain and the loss.

A way to assess the complexity of a question is the amount of time participants take to answer: if participants struggle to understand a question they will usually take longer to answer it (or much shorter if they give up). The DOSE module (including instructions), was, however, fast to complete. The median time on the risk-loss DOSE module was 115 seconds, and on the time-discounting module was 107 seconds. In comparison, the median time taken to complete the MPL instructions and first elicitation was 259 seconds.[39]

[39]The first MPL measured time preferences. The DOSE modules always appeared prior to this MPL.

Figure 8: Low cognitive ability participants take longer on MPL questions, but not on DOSE.

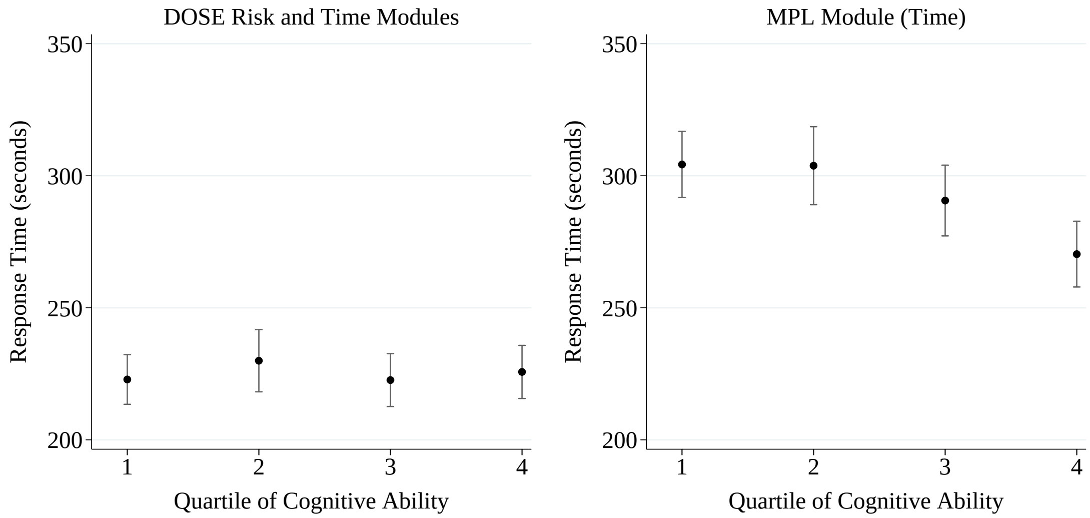

Notes: DOSE module includes 20 questions addressing both risk and time preferences. MPL module includes two MPLs assessing time preferences, which was the first MPL module on the survey. Respondents with response times greater than 15 minutes for either module are excluded. Both panels include time taken for questions and explanatory text.

Additional evidence that DOSE was simple to understand is that low cognitive ability participants took the same amount of time as high cognitive ability participants to complete the DOSE module, in contrast to the MPL—see Figure 8. Participants in the top cognitive ability quartile took, on average, 34 seconds less in the MPL module than those in the bottom quartile. For DOSE, in contrast, high cognitive ability participants took 3 seconds longer on average (p-value=0.68). Moreover, the variance of time taken was relatively constant across quartiles, as indicated by the confidence intervals in the figure. Together, these facts suggests that DOSE was equally easy for participants all along the cognitive ability spectrum.

However, just because participants, in general, seem to have found DOSE easier to understand than the MPL doesn't mean there were *no* participants confused by DOSE, or that all were paying attention. We examine this possibility next by re-analyzing our data while leaving out those who were most likely to have given up and rushed through the survey. The same results hold, implying that neither confusion nor inattentiveness, nor giving up, is likely to explain many of the choices we see.

Our results are largely unchanged when removing the fastest responses. As shown in

Figure 9: Distributions are similar when removing participants with short response times.

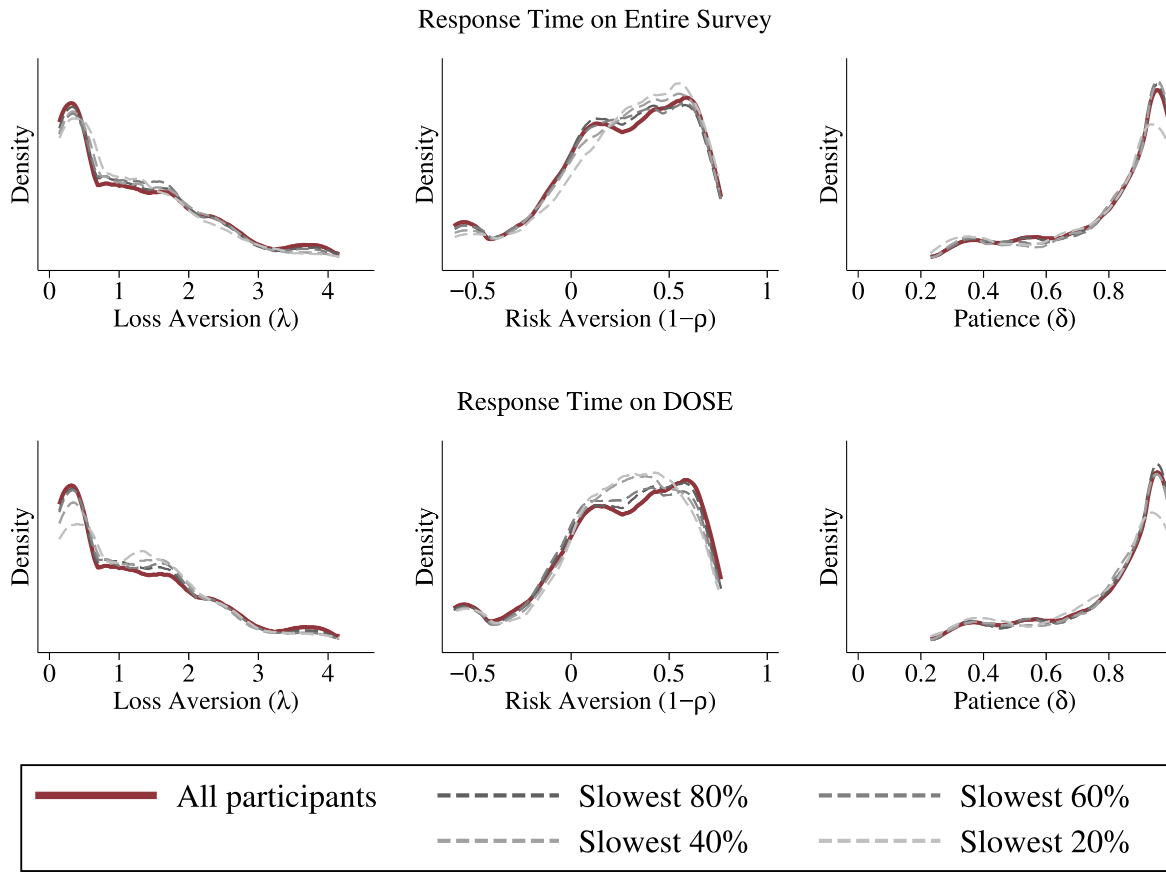

Notes: Plotted using Epanechnikov kernel, with bandwidth chosen by rule-of-thumb estimator.

Figure 9, the distribution of economic preferences is similar when restricting the sample by removing participants according to the quintile of response time. That is, we first look at the slowest 80% of respondents, then the slowest 60%, and so on. The distributions overlap almost entirely—and the median loss aversion parameter consistently remains very similar in the whole sample (1.02 or below): fast response or inattention cannot be said to be the explanation for loss tolerance. Moreover, correlations with other characteristics are also similar when removing the fastest respondents—see Appendix Table D.9.

A final manifestation of inattention might be choosing the same option in each question: either the lottery (always listed first), or the sure amount (always listed second). However, there is little evidence of this pattern of inattention in our results: fewer than 6% participants chose the same option in all ten question rounds. While we cannot rule out that some participants rapidly clicked through the DOSE module, such behavior does not appear to affect our results.

6 Discussion

In this paper, we introduce DOSE—Dynamically Optimized Sequential Experimentation—and use it to study loss aversion and other economic preferences in a representative sample of the U.S. population. Our results are summarized in Table 3. A few are worth highlighting. First, we find that around 50% of the U.S. population is loss tolerant over small stakes, differing from prior studies that have found a strong majority of loss averse participants, usually in lab/student samples. Second, those with greater cognitive ability, education, and income are more likely to be loss averse, and those with lower cognitive ability are more likely to be loss tolerant. This, along with the fact that DOSE in lab/student samples produces similar results to prior studies, suggests that differences in samples are likely the source of the difference between our results and prior studies. Third, using DOSE's choice consistency parameter we show that those with high consistency exhibit a correlation on MPL-based measures between higher cognitive ability and less risk aversion. This suggests that the mixed results about this relationship in the literature (Dohmen et al., 2018) may be due to measurement error (Gillen et al., Forthcoming). Fourth, across a range of evaluations, DOSE produces better measures: more accurate, more stable, faster, and so on.

Our findings about loss aversion diverge significantly from conventional wisdom, raising the possibility that the literature may have been influenced by factors beyond the inadvertent sample selection mentioned above. Hints can be found in Fehr-Duda and Epper (2012, p. 576), who observe, "Since the publication of Tversky and Kahneman (1992), any estimates of loss aversion that deviate significantly from the value of two have been eyed with great suspicion, notwithstanding the fact that the original estimate was based on 25

Table 3: Comparison of DOSE with Other Elicitation Methods

		DOSE	MPL	Risky Project / Lottery Menu (†)
Substantive	**Loss Attitudes**			
	Percent Loss Tolerant	53%	n.a.	n.a.
	Median Loss Aversion Parameter	0.98	n.a.	n.a.
	Correlations w/Cognitive Ability			
	Risk Aversion	-0.21	≈ 0	-0.07
	Loss Aversion	0.40	n.a.	n.a.
	Patience	0.18	0.17	n.a.
	Correlations w/ Demographics			
	Risk Aversion	✓	≈ 0	✓
	Loss Aversion	✓	n.a.	n.a.
	Patience	✓	✓	n.a.
Methodological	**Representative Survey**			
	Speed	115 secs	259 secs	33 secs
	Stability: Risk Aversion	0.45	0.28	0.33
	Stability: Loss Aversion	0.40	n.a.	n.a.
	Stability: Patience	0.47	0.24	n.a.
	Parameter Recovery Analysis			
	Inaccuracy: Risk Aversion	15%	37%	35%†
	Inaccuracy: Loss Aversion	15%	36%	n.a.†
	Correlation: Risk Aversion	0.79	0.45	0.28†
	Correlation: Loss Aversion	0.91	0.65	n.a.†

Notes: "Inaccuracy" is the average absolute percentage difference between the estimated and true parameter values. "Correlation" is the correlation between the estimated and true parameter values. "Stability" is the correlation across survey waves. ✓ indicates that a pattern of statistically significant correlations were identified. Observations denoted with a † are from the Lottery Menu elicitation—which was included in our simulations, but not our survey—rather than the Risky Project—which was included in our survey, but not our simulations.

subjects, hypothetical decisions over relatively large stakes, and that no standard errors were reported." Relatedly, the one study that examines loss aversion in a representative sample, von Gaudecker et al. (2011), report a median estimate of $\lambda = 2.38$, although specifications in the appendix have medians ranging from 0.12 to 4.47. More directly, Walasek et al. (in progress) analyze 19 studies of loss aversion in lab/student populations, and find evidence of publication bias, and Yechiam (2018, p. 1) notes in a review of the loss aversion literature that, "[T]he findings of some of these studies have been systematically misrepresented to reflect loss aversion, though they did not find it."

Our DOSE implementation provides estimates of risk aversion and time preferences that also appear to dominate MPL-based elicitations. DOSE risk aversion measures show stronger correlations with cognitive ability and other characteristics than MPL-based measures, as shown in Table 3. Further, DOSE-based measures of risk and time preferences show greater stability than MPL-based measures. These facts indicate, in-line with our simulations, that DOSE produces estimates with lower measurement error in these domains as well.

A common concern about DOSE is that, like most dynamic methods, they are not incentive compatible. For example, participants could misleadingly say they prefer a lottery to a sure amount in the first question in order to increase the magnitude of the sure amounts offered in the future. However, this is possible with *any* sequence of questions that participants believe are dynamic—and few experiments explicitly rule out this possibility (our survey did not explicitly mention the dynamic nature of our 20 questions). However, in practice this is of little concern: Ray et al. (2012) find minimal possible benefits of manipulation, and little evidence that even very sophisticated participants engage in this behavior even when explicitly informing participants that the question sequence is manipulable. For more on this concern, see Appendix A.

The accuracy, speed, and simplicity of DOSE potentially expands the range of research settings in which incentivized preference elicitation is viable. The procedure may be particularly valuable in field experiments, where it is difficult to provide participants with detailed instructions. Similarly, it may also be easier to implement than more complex or time consuming designs when conducting experiments with low literacy participants, children, patients with medical disorders, or even animals. DOSE performs better with low cognitive ability, low-education, and low-income participants, suggesting that it could be particularly useful in development environments, where current elicitation methods can be plagued by inconsistent choice (Jacobson and Petrie, 2009; Charness and Viceisza, 2012) and new techniques are particularly needed (Berry et al., 2015). DOSE can also be used to discriminate between models on an individual level in real time—an application developed for time preferences in Imai and Camerer (2018), based on an earlier working paper version of this manuscript (Wang et al., 2010).

There is more to be learned about loss aversion with DOSE simply by broadening the types of questions participants are asked. Offering participants lotteries with prizes only in the loss domain would allow better identification of differences in the curvature of the utility function in the gain and loss domain. Further, the degree of loss aversion may be affected by stake size and "zero avoidance" (Ert and Erev, 2013). The questions used in this paper involved only small stakes and always contained a prize of zero payoff—the sure amount in questions involving a lottery with a loss, and one of the lottery prizes in the questions only involving gains. These design choices could straightforwardly be varied to obtain a richer view of the factors affecting loss aversion, and something like "zero aversion" could be modeled directly and relevant parameters estimated.

Even with these questions outstanding, our findings relate to current psychological and neuroscientific work on gain-loss differences. The fact that loss aversion depends on cognitive skill suggests it may have some basis in effortful cognitive processes. Recent evidence indicates aversion or tolerance to loss is associated with mental information accumulation (Clay et al., 2017) and attention paid to losses and gains (Bhatia and Golman, 2015; Yechiam and Hochman, 2013). The clearest evidence for the role of attention is that visual attention to losses correlates modestly (0.31) with the loss aversion parameter inferred from choices

(Pachur et al., 2018). In addition, exogenously manipulating attention by visually presenting it for an extra 0.6 seconds increased estimated loss-aversion by about 10%. It would be useful to know whether people with lower cognitive ability pay less attention to losses.

A small amount of neuroscientific evidence indicates stronger encoding in the amygdala in response to loss (Yacubian et al. 2006; De Martino et al. (2010); though not in Tom et al. 2007). The anti-anxiety drug Propranalol lowers loss aversion by about 15% (in low-BMI subjects; Sokol-Hessner et al., 2015). Given the role of the amygdala in rapid "vigilance" threat processing, these results suggest lower cognitive ability participants may simply feel less fear or anxiety about losses. Experiments have shown that differences in the distributions that are faced can influence choices—losses are more tolerated when many choices include possible losses, due to a general cognitive process called "adaptive coding" (Walasek and Stewart, 2015). Thus, it is possible that people with lower cognitive ability are more routinely exposed to everyday gain and loss distributions, which create a tolerance for loss.

Obviously, our incentivized survey data cannot test any of these mechanistic hypotheses. Given the striking behavioral association between loss tolerance and cognitive skill, however, it would be useful to further explore the roles of attention, anxiety, and distributional experience with low and high cognitive ability groups.

References

Abdellaoui, Mohammed, Han Bleichrodt, and Corina Paraschiv, "Loss Aversion under Prospect Theory: A Parameter-Free Measurement," *Management Science*, 2007, *53* (10), 1659–1674.

_ , _ , **and Olivier l'Haridon**, "A Tractable Method to Measure Utility and Loss Aversion under Prospect Theory," *Journal of Risk and Uncertainty*, 2008, *36* (3), 245.

_ , **Olivier L'Haridon, and Corina Paraschiv**, "Experienced vs. Described Uncertainty: Do We Need Two Prospect Theory Specifications?," *Management Science*, 2011, *57* (10), 1879–1895.

Abernethy, Jacob, Theodoros Evgeniou, Olivier Toubia, and Jean-Philippe Vert, "Eliciting Consumer Preferences Using Robust Adaptive Choice Questionnaires," *IEEE Transactions on Knowledge and Data Engineering*, 2008, *20* (2), 145–155.

Aigner, Dennis J., "A Brief Introduction to the Methodology of Optimal Experimental Design," *Journal of Econometrics*, 1979, *11* (1), 7–26.

Allen, Eric J. Patricia M. Dechow, Devin G. Pope, and George Wu, "Reference-Dependent Preferences: Evidence from Marathon Runners," *Management Science*, 2016, *63* (6), 1657–1672.

Andersen, Steffen, Glenn W. Harrison, Morten I. Lau, and E.Ẽlisabet Rutström, "Eliciting Risk and Time Preferences," *Econometrica*, 2008, *76* (3), 583–618.

_ , _ , **Morten I. Lau, and E. Elisabet Rutström**, "Lost in State Space: Are Preferences Stable?," *International Economic Review*, 2008, *49* (3), 1091–1112.

_ , _ , **Morten Igel Lau, and E. Elisabet Rutström**, "Elicitation Using Multiple Price List Formats," *Experimental Economics*, 2006, *9* (4), 383–405.

Andersson, Ola, Håkan J. Holm, Jean-Robert Tyran, and Erik Wengström, "Deciding For Others Reduces Loss Aversion," *Management Science*, 2016, *62* (1), 29–36.

_ , _ , _ , **and** _ , "Risk Aversion Relates to Cognitive Ability: Preferences or Noise?," *Journal of the European Economic Association*, 2016, *14* (5), 1129–1154.

Andreoni, James and Charles Sprenger, "Estimating Time Preferences from Convex Budgets," *The American Economic Review*, 2012, *102* (7), 3333–3356.

_ **and** _ , "Risk Preferences Are Not Time Preferences," *The American Economic Review*, 2012, *102* (7), 3357–3376.

_ , **Michael Callen, Yasir Khan, Karrar Jaffar, and Charles Sprenger**, "Using Preference Estimates to Customize Incentives: An Application to Polio Vaccination Drives in Pakistan," Technical Report, National Bureau of Economic Research 2016.

Åstebro, Thomas, José Mata, and Luís Santos-Pinto, "Skewness Seeking: Risk Loving, Optimism or Overweighting of Small Probabilities?," *Theory and Decision*, 2015, *78* (2), 189–208.

Atlas, Stephen A., Eric J. Johnson, and John W. Payne, "Time Preferences and Mortgage Choice," *Journal of Marketing Research*, 2017, *54* (3), 415–429.

Azrieli, Yaron, Christopher P. Chambers, and Paul J. Healy, "Incentives in Experiments: A Theoretical Analysis," *Journal of Political Economy*, 2018, *126* (4), 1472–1503.

Barsky, Robert B. F. Thomas Juster, Miles S. Kimball, and Matthew D. Shapiro, "Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study," *The Quarterly Journal of Economics*, 1997, *112* (2), 537–579.

Beattie, Jane and Graham Loomes, "The Impact of Incentives upon Risky Choice Experiments," *Journal of Risk and Uncertainty*, March 1997, *14* (2), 155–168.

Benartzi, Shlomo and Richard H. Thaler, "Myopic Loss Aversion and the Equity Premium Puzzle," *The Quarterly Journal of Economics*, 1995, *110* (1), 73–92.

Benjamin, Daniel J., Sebastian A. Brown, and Jesse M. Shapiro, "Who Is 'Behavioral'? Cognitive Ability and Anomalous Preferences," *Journal of the European Economic Association*, 2013, *11* (6), 1231–1255.

Berry, James, Greg Fischer, and Raymond P. Guiteras, "Eliciting and Utilizing Willingness to Pay: Evidence from Field Trials in Northern Ghana," 2015. University of Delaware, *mimeo.*

Bhatia, Sudeep and Russell Golman, "Attention and Reference Dependence," 2015. University of Pennsylvania, *mimeo.*

Booij, Adam S. and Gijs Van de Kuilen, "A Parameter-Free Analysis of the Utility of Money for the General Population under Prospect Theory," *Journal of Economic Psychology*, 2009, *30* (4), 651–666.

_ , Bernard M.S. Van Praag, and Gijs Van De Kuilen, "A Parametric Analysis of Prospect Theory's Functionals for the General Population," *Theory and Decision*, 2010, *68* (1-2), 115–148.

Brooks, Peter and Horst Zank, "Loss Averse Behavior," *Journal of Risk and Uncertainty*, 2005, *31* (3), 301–325.

Burks, Stephen V., Jeffrey P. Carpenter, Lorenz Goette, and Aldo Rustichini, "Cognitive Skills Affect Economic Preferences, Strategic Behavior, and Job Attachment," *Proceedings of the National Academy of Sciences*, 2009, *106* (19), 7745–7750.

Camerer, Colin, Linda Babcock, George Loewenstein, and Richard Thaler, "Labor Supply of New York City Cabdrivers: One Day at a Time," *The Quarterly Journal of Economics*, 1997, *112* (2), 407–441.

Cavagnaro, Daniel R., Gabriel J. Aranovich, Samuel M. McClure, Mark A. Pitt, and Jay I. Myung, "On the Functional Form of Temporal Discounting: An Optimized Adaptive Test," *Journal of Risk and Uncertainty*, 2016, *52* (3), 233–254.

_ , Jay I. Myung, Mark A. Pitt, and Janne V. Kujala, "Adaptive Design Optimization: A Mutual Information-Based Approach to Model Discrimination in Cognitive Science," *Neural Computation*, 2010, *22* (4), 887–905.

_ , Mark A. Pitt, Richard Gonzalez, and Jay I. Myung, "Discriminating Among Probability Weighting Functions using Adaptive Design Optimization," *Journal of Risk and Uncertainty*, 2013, *47* (3), 255–289.

_ , Richard Gonzalez, Jay I. Myung, and Mark A. Pitt, "Optimal Decision Stimuli for Risky Choice Experiments: An Adaptive Approach," *Management Science*, 2013, *59* (2), 358–375.

Chaloner, Kathryn and Isabella Verdinelli, "Bayesian Experimental Design: A Review," *Statistical Science*, 1995, *10* (3), 273–304.

Chang, Alvin, "Four Ways the Lottery Preys on the Poor," January 2016. https://www.vox.com/identities/2016/1/13/10763268/lottery-poor-prey.

Chapman, Jonathan, Mark Dean, Pietro Ortoleva, Erik Snowberg, and Colin F. Camerer, "Willingness to Pay and Willingness to Accept are Probably Less Correlated Than You Think," 2017. NBER Working Paper #23954.

_ , _ , _ , _ , and _ , "Econographics," 2018. NBER Working Paper #24931.

Charness, Gary and Angelino Viceisza, "Comprehension and Risk Elicitation in the Field: Evidence from Rural Senegal," 2012. University of California Santa Barbara, *mimeo.*

_ , Uri Gneezy, and Brianna Halladay, "Experimental Methods: Pay One or Pay All," *Journal of Economic Behavior & Organization*, 2016, *131*, 141–150.

Chuang, Yating and Laura Schechter, "Stability of Experimental and Survey Measures of Risk, Time, and Social Preferences: A Review and Some New Results," *Journal of Development Economics*, 2015, *117*, 151–170.

Clay, S.N. A. Harris, and C.L. Reed, "Loss Aversion Affects Inhibitory Processes for Reward and Loss as Indicated by Inhibition of Return," in progress.

Clay, Summer N. John A. Clithero, Alison M. Harris, and Catherine L. Reed, "Loss Aversion Reflects Information Accumulation, Not Bias: A Drift-Diffusion Model Study," *Frontiers in Psychology*, 2017, *8:1708.*

Condon, David M. and William Revelle, "The International Cognitive Ability Resource: Development and Initial Validation of a Public-Domain Measure," *Intelligence*, 2014, *43*, 52–64.

Cubitt, Robin, Chris Starmer, and Robert Sugden, "On the Validity of the Random Lottery Incentive System," *Experimental Economics*, 1998, *1* (2), 115–131.

Dave, Chetan, Catherine C. Eckel, Cathleen A. Johnson, and Christian Rojas, "Eliciting Risk Preferences: When is Simple Better?," *Journal of Risk and Uncertainty*, 2010, *41* (3), 219–243.

De Martino, Benedetto, Colin F Camerer, and Ralph Adolphs, "Amygdala Damage Eliminates Monetary Loss Aversion," *Proceedings of the National Academy of Sciences*, 2010, *107* (8), 3788–3792.

Dean, Mark and Anja Sautmann, "Credit Constraints and the Measurement of Time Preferences," 2014. Columbia University, *mimeo.*

DellaVigna, Stefano and M. Daniele Paserman, "Job Search and Impatience," *Journal of Labor Economics*, 2005, *23* (3), 527–588.

Dohmen, Thomas, Armin Falk, Armin Huffman, and Uwe Sunde, "Are Risk Aversion and Impatience Related to Cognitive Ability?," *The American Economic Review*, 2010, *100* (3), 1238–1260.

_ , _ , David Huffman, and Uwe Sunde, "On the Relationship between Cognitive Ability and Risk Preference," *Journal of Economic Perspectives*, 2018, *32* (2), 115–34.

Dunn, Lucia F., "Loss Aversion and Adaptation in the Labor Market: Empirical Indifference Functions and Labor Supply," *The Review of Economics and Statistics*, 1996, *78* (3), 441–450.

Eckel, Catherine C. and Philip J. Grossman, "Sex Differences and Statistical Stereotyping in Attitudes toward Financial Risk," *Evolution and Human Behavior*, 2002, *23* (4), 281–295.

_ and _ , "Men, Women and Risk Aversion: Experimental Evidence," in Charles R. Plott and Vernon L. Smith, eds., *Handbook of Experimental Economics Results*, Vol. 1, North-Holland, 2008, pp. 1061–1073.

El-Gamal, Mahmoud A. and Thomas R. Palfrey, "Economical Experiments: Bayesian Efficient Experimental Design," *International Journal of Game Theory*, 1996, *25* (4), 495–517.

El-Gamal, Mahmoud A., Richard D. McKelvey, and Thomas R. Palfrey, "A Bayesian Sequential Experimental Study of Learning in Games," *Journal of the American Statistical Association*, 1993, *88* (422), 428–435.

Ert, Eyal and Ido Erev, "On the Descriptive Value of Loss Aversion in Decisions under Risk: Six Clarifications," *Judgment and Decision Making,*, 2013, *8* (3), 214–235.

Falk, Armin, Anke Becker, Thomas Dohmen, Benjamin Enke, David Huffman, and Uwe Sunde, "Global Evidence on Economic Preferences," *The Quarterly Journal of Economics*, 2018, p. https://doi.org/10.1093/qje/qjy013.

Fehr-Duda, Helga and Thomas Epper, "Probability and Risk: Foundations and Economic Implications of Probability-Dependent Risk Preferences," *Annual Review of Economics*, 2012, *4* (1), 567–593.

Frederick, Shane, "Cognitive Reflection and Decision Making," *Journal of Economic Perspectives*, 2005, *19* (4), 25–42.

Freeman, David, Yoram Halevy, and Terri Kneeland, "Eliciting Risk Preferences using Choice Lists," 2015. University of British Columbia, *mimeo.*

Frydman, Cary, Colin Camerer, Peter Bossaerts, and Antonio Rangel, "MAOA-L Carriers are Better at Making Optimal Financial Decisions under Risk," *Proceedings of the Royal Society of London B: Biological Sciences*, 2011, *278* (1714), 2053–2059.

Gächter, Simon, Eric J. Johnson, and Andreas Herrmann, "Individual-level Loss Aversion in Riskless and Risky Choices," 2007. *mimeo.*

Gillen, Ben, Erik Snowberg, and Leeat Yariv, "Experimenting with Measurement Error: Techniques and Applications from the Caltech Cohort Study," *Journal of Political Economy*, Forthcoming.

Gneezy, Uri and Jan Potters, "An Experiment on Risk Taking and Evaluation Periods," *The Quarterly Journal of Economics*, 1997, *112* (2), 631–645.

Goette, Lorenz, David Huffman, and Ernst. Fehr, "Loss Aversion and Labor Supply," *Journal of the European Economic Association*, 2004, *2* (2-3), 216–228.

_ , Thomas Graeber, Alexandre Kellogg, and Charles Sprenger, "Heterogeneity of Loss Aversion and Expectations-Based Reference Points," *Working Paper*, 2018.

Hardie, Bruce G.S. Eric J. Johnson, and Peter S. Fader, "Modeling Loss Aversion and Reference Dependence Effects on Brand Choice," *Marketing Science*, 1993, *12* (4), 378–394.

Harrison, Glenn W. Morten I. Lau, and E. Elisabet Rutström, "Estimating Risk Attitudes in Denmark: A Field Experiment," *The Scandinavian Journal of Economics*, 2007, *109* (2), 341–368.

_ , _ , and Melonie B. Williams, "Estimating Individual Discount Rates in Denmark: A Field Experiment," *The American Economic Review*, 2002, *92* (5), 1606–1617.

Hey, John and Jinkwon Lee, "Do Subjects Separate (or Are They Sophisticated)?," *Experimental Economics*, September 2005, *8* (3), 233–265.

Holt, Charles A. and Susan K. Laury, "Risk Aversion and Incentive Effects," *The American Economic Review*, 2002, *92* (5), 1644–1655.

Imai, Taisuke and Colin F. Camerer, "Estimating Time Preferences from Budget Set Choices Using Optimal Adaptive Design," 2018. Ludwig Maximilian University of Munich, *mimeo.*

Jacobson, Sarah and Ragan Petrie, "Learning from Mistakes: What Do Inconsistent Choices Over Risk Tell Us?," *Journal of Risk and Uncertainty*, 2009, *38* (2), 143–158.

Kahneman, Daniel and Amos Tversky, "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, 1979, *47* (2), 263–291.

_ **and Dan Lovallo**, "Timid Choices and Bold Forecasts: A Cognitive Perspective on Risk Taking," *Management Science*, 1993, *39* (1), 17–31.

Karni, Edi and Zvi Safra, "'Preference Reversal' and the Observability of Preferences by Experimental Methods," *Econometrica*, 1987, *55* (3), 675–685.

Kirby, Kris N., "One-year Temporal Stability of Delay-Discount Rates," *Psychonomic Bulletin & Review*, 2009, *16* (3), 457–462.

Kirby, Kris N. Ricardo Godoy, Victoria Reyes-Garcıa, Elizabeth Byron, Lilian Apaza, William Leonard, Eddy Perez, Vincent Vadez, and David Wilkie, "Correlates of Delay-Discount Rates: Evidence from Tsimane'Amerindians of the Bolivian Rain Forest," *Journal of Economic Psychology*, 2002, *23* (3), 291–316.

Köbberling, Veronika and Peter Wakker, "An Index of Loss Aversion," *Journal of Economic Theory*, 2005, *122*, 119–131.

Kornotis, George M. and Alok Kumar, "Cognitive Abilities and Financial Decisions," in H. Kent Baker and John Nofsinger, eds., *Behavioral Finance*, Hoboken, NJ: John Wiley & Sons, Inc., 2010.

Kőszegi, Botond and Matthew Rabin, "A Model of Reference-Dependent Preferences," *The Quarterly Journal of Economics*, 2006, *121* (4), 1133–1165.

Krajbich, Ian, Colin Camerer, and Antonio Rangel, "Exploring the Scope of Neurometrically Informed Mechanism Design," *Games and Economic Behavior*, 2017, *101*, 49–62.

Kujala, Janne V. and Tuomas J. Lukka, "Bayesian Adaptive Estimation: The Next Dimension," *Journal of Mathematical Psychology*, 2006, *50* (4), 369–389.

Kurata, Hiroshi, Hiroshi Izawa, and Makoto Okamura, "Non-Expected Utility Maximizers Behave as if Expected Utility Maximizers: An Experimental Test," *Journal of Socio-Economics*, 2009, *38* (4), 622–629.

Lesmes, Luis Andres, Seong-Taek Jeon, Zhong-Lin Lu, and Barbara Anne Dosher, "Bayesian Adaptive Estimation of Threshold Versus Contrast External Noise Functions: The Quick TvC Method," *Vision Research*, 2006, *46* (19), 3160–3176.

Levin, Irwin P., Stephanie S. Hart, Joshua A. Weller, and Lyndsay A. Harshman, "Stability of Choices in a Risky Decision-Making task: A 3-year Longitudinal Study with Children and Adults," *Journal of Behavioral Decision Making*, 2007, *20* (3), 241–252.

Lewi, Jeremy, Robert Butera, and Liam Paninski, "Sequential Optimal Design of Neurophysiology Experiments," *Neural Computation*, 2009, *21* (3), 619–687.

Lönnqvist, Jan-Erik, Markku Verkasalo, Gari Walkowitz, and Philipp C. Wichardt, "Measuring Individual Risk Attitudes in the Lab: Task or Ask? An Empirical Comparison," *Journal of Economic Behavior & Organization*, 2015, *119*, 254–266.

McFadden, Daniel, "Economic Choices," *The American Economic Review*, 2001, *91* (3), 351–378.

Mehra, Rajnish and Edward C. Prescott, "The Equity Premium: A Puzzle," *Journal of Monetary Economics*, 1985, *15* (2), 145–161.

Meier, Stephan and Charles D. Sprenger, "Temporal Stability of Time Preferences," *Review of Economics and Statistics*, 2015, *97* (2), 273–286.

Moffatt, P., "Optimal Experimental Designs in Models of Decision and Choice," in M. Boumans, ed., *Measurement in Economics: A Handbook*, Elsevier, 2007.

Müller, Peter, Don A. Berry, Andy P. Grieve, Michael Smith, and Michael Krams, "Simulation-Based Sequential Bayesian Design," *Journal of Statistical Planning and Inference*, 2007, *137* (10), 3140–3150.

O'Donoghue, Ted and Charles Sprenger, "Reference-Dependent Preferences," in Bernheim B. Douglas, DellaVigna Stefano, and Laibson David, eds., *Handbook of Behavioral Economics: Foundations and Applications*, Vol. 1, Elsevier, 2018.

Oechssler, Jörg, Andreas Roider, and Patrick W. Schmitz, "Cognitive Abilities and Behavioral Biases," *Journal of Economic Behavior & Organization*, 2009, *72*, 147–152.

Pachur, Thorsten, Michael Schulte-Mecklenbeck, Ryan O Murphy, and Ralph Hertwig, "Prospect Theory Reflects Selective Allocation of Attention," *Journal of Experimental Psychology: General*, 2018, *147* (2), 147–169.

Paninski, Liam, "Asymptotic Theory of Information-Theoretic Experimental Design," *Neural Computation*, 2005, *17* (7), 1480–1507.

Peirce, Charles S., "Note on the Theory of the Economy of Research," 1879. United States Coast Survey for the fiscal year ending June 1876, US Government Printing Office 1879, reprinted in *Operations Research*, 1967, vol. XV, 643–648.

Pew Research Center, *Evaluating Online Nonprobability Surveys*, www.pewresearch.org, May 2016.

Pope, Devin and Uri Simonsohn, "Round Numbers as Goals: Evidence from Baseball, SAT Takers, and the Lab," *Psychological Science*, 2011, *22* (1), 71–79.

Ray, Debajyoti, Daniel Golovin, Andreas Krause, and Colin Camerer, "Bayesian Rapid Optimal Adaptive Design (BROAD): Method and Application Distinguishing Models of Risky Choice," 2012. California Institute of Technology, *mimeo.*

Rees-Jones, Alex, "Quantifying Loss-Averse Tax Manipulation," *The Review of Economic Studies*, 2017, *85* (2), 1251–1278.

Ryan, Elizabeth G. Christopher C. Drovandi, James M. McGree, and Anthony N. Pettitt, "A Review of Modern Computational Algorithms for Bayesian Optimal Design," *International Statistical Review*, 2016, *84* (1), 128–154.

Schildberg-Hörisch, Hannah, "Are Risk Preferences Stable?," *Journal of Economic Perspectives*, 2018, *32* (2), 135–54.

Schmidt, Ulrich and Stefan Traub, "An Experimental Test of Loss Aversion," *Journal of Risk and Uncertainty*, 2002, *25* (3), 233–249.

Shamosh, Noah A and Jeremy R Gray, "Delay Discounting and Intelligence: A Meta-Analysis," *Intelligence*, 2008, *36* (4), 289–305.

Snowberg, Erik and Leeat Yariv, "Testing the Waters: Behavior across Subject Pools," 2018. NBER Working Paper #24781.

Sokol-Hessner, Peter, Ming Hsu, Nina G. Curley, Mauricio R. Delgado, Colin F Camerer, and Elizabeth A. Phelps, "Thinking Like a Trader Selectively Reduces Individuals' Loss Aversion," *Proceedings of the National Academy of Sciences*, 2009, *106* (13), 5035–5040.

_ , Sandra F. Lackovic, Russell H. Tobe, Colin F. Camerer, Bennett L. Leventhal, and Elizabeth A. Phelps, "Determinants of Propranolols Selective Effect on Loss Aversion," *Psychological Science*, 2015, *26* (7), 1123–1130.

Sprenger, Charles, "An Endowment Effect for Risk: Experimental Tests of Stochastic Reference Points," *Journal of Political Economy*, 2015, *123* (6), 1456–1499.

Tanaka, Tomomi, Colin F Camerer, and Quang Nguyen, "Risk and Time Preferences: Linking Experimental and Household Survey Data from Vietnam," *The American Economic Review*, 2010, *100* (1), 557–71.

Tom, Sabrina M. Craig R. Fox, Christopher Trepel, and Russell A. Poldrack, "The Neural Basis of Loss Aversion in Decision-Making under Risk," *Science*, 2007, *315* (5811), 515–518.

Toubia, Olivier, Eric Johnson, Theodoros Evgeniou, and Philippe Delquié, "Dynamic Experiments for Estimating Preferences: An Adaptive Method of Eliciting Time and Risk Parameters," *Management Science*, 2013, *59* (3), 613–640.

_ , John R. Hauser, and Duncan I. Simester, "Polyhedral Methods for Adaptive Choice-Based Conjoint Analysis," *Journal of Marketing Research*, 2004, *41* (1), 116–131.

Tovar, Patricia, "The Effects of Loss Aversion on Trade Policy: Theory and Evidence," *Journal of International Economics*, 2009, *78* (1), 154–167.

Tversky, Amos and Daniel Kahneman, "Advances in Prospect Theory: Cumulative Representation of Uncertainty," *Journal of Risk and Uncertainty*, 1992, *5* (4), 297–323.

Urminsky, Oleg and Gal Zauberman, "The Psychology of Intertemporal Preferences," in Gideon Keren and George Wu (Editor), eds., *Handbook of Judgment and Decision Making*, Vol. 2, Wiley-Blackwell, 2016, pp. 141–181.

von Gaudecker, Hans-Martin, Arthur Van Soest, and Erik Wengström, "Heterogeneity in Risky Choice Behavior in a Broad Population," *The American Economic Review*, 2011, *101* (2), 664–694.

Wakker, Peter P., *Prospect Theory: For Risk and Ambiguity*, Cambridge University Press, 2010.

Walasek, Lukasz and Neil Stewart, "How to make Loss Aversion Disappear and Reverse: Tests of the Decision by Sampling Origin of Loss Aversion.," *Journal of Experimental Psychology: General*, 2015, *144* (1), 7–11.

_ , Timothy L. Mullett, and Stewart Neil., "A Meta-Analysis of Loss Aversion in Risky Contexts," in progress. *mimeo.*

Wald, Abraham, *Statistical Decision Functions*, John Wiley and Sons, New York; Chapman and Hall, London., 1950.

Wang, Stephanie W. Michelle Filiba, and Colin F. Camerer, "Dynamically Optimized Sequential Experimentation (DOSE) for Estimating Economic Preference Parameters," 2010. *mimeo.*

Wölbert, Eva and Arno Riedl, "Measuring Time and Risk Preferences: Reliability, Stability, Domain Specificity," 2013. *CESifo Working Paper Series No. 4339.*

Yacubian, Juliana, Jan Gläscher, Katrin Schroeder, Tobias Sommer, Dieter F Braus, and Christian Büchel, "Dissociable Systems for Gain- and Loss-Related Value Predictions and Errors of Prediction in the Human Brain," *Journal of Neuroscience*, 2006, *26* (37), 9530–9537.

Yechiam, Eldad, "Acceptable Losses: The Debatable Origins of Loss Aversion," *Psychological Research*, 2018, *https://doi.org/10.1007/s00426-018-1013-8.*

_ and Guy Hochman, "Losses as Modulators of Attention: Review and Analysis of the Unique Effects of Losses over Gains," *Psychological Bulletin*, 2013, *139* (2), 497–518.