

Estimating Subjective Probabilities

by

Steffen Andersen, John Fountain, Glenn W. Harrison & E. Elisabet Rutström[†]

January 2013

Department of Economics, Copenhagen Business School, Copenhagen, Denmark (Andersen);

Department of Economics, University of Canterbury, Christchurch, New Zealand (Fountain);

Department of Risk Management & Insurance and CEAR, Robinson College of Business, Georgia

State University, USA (Harrison); and Department of Economics, Andrew Young School of Policy

Studies, and Dean's Behavioral Economics Laboratory, Robinson College of Business, Georgia

State University, USA (Rutström).

Corresponding Author: Glenn W. Harrison, Center for the Economic Analysis of Risk

J. Mack Robinson College of Business, Georgia State University, P.O. BOX 4036, Atlanta, GA

30302-4036, gharrison@gsu.edu and erutstrom@gsu.edu.

[†] Acknowledgements: We thank the U.S. National Science Foundation for research support under grants NSF/HSD 0527675 and NSF/SES 0616746. We are grateful to the Editor, a referee, John Duffy, J. Todd Swarthout, Peter Wakker, Randall Walsh and Nathaniel Wilcox for discussions and helpful comments. Complete instructions, additional estimation results, and details of the econometric specification are provided in Web Appendices available in *CEAR Working Paper 2010-06* at <http://cear.gsu.edu>.

ABSTRACT. Subjective probabilities play a central role in many economic decisions, and act as an immediate confound of inferences about behavior, unless controlled for. Several procedures to recover subjective probabilities have been proposed, but in order to recover the correct latent probability one must either construct elicitation mechanisms that control for risk aversion, or construct elicitation mechanisms which undertake “calibrating adjustments” to elicited reports. We illustrate how the *joint estimation of risk attitudes and subjective probabilities* can provide the calibration adjustments that theory calls for. We illustrate this approach using data from a controlled experiment with real monetary consequences to the subjects. This allows the observer to make inferences about the latent subjective probability, under virtually any well-specified model of choice under subjective risk, while still employing relatively simple elicitation mechanisms.

JEL codes: C9, C8, D84

Subjective probabilities are operationally defined as those probabilities that lead an agent to choose some prospects over others when outcomes depend on events that are not yet actualized. These choices could be as natural as placing a bet on a horse race, or as experimentally structured as responding to the payoff prizes provided by some scoring rule. In order to infer subjective probabilities from observed choices of this kind, however, one either has to make some strong assumptions about risk attitudes or jointly estimate risk attitudes and subjective probabilities. We show how the latter can be implemented by pairing several experimental tasks together, some of which identify risk attitudes and some of which identify the interplay between risk attitudes and subjective probabilities. Joint estimation of a structural model of choice across these two types of tasks allows one to make inferences about subjective probabilities from observed behavior in relatively simple choice tasks.¹

The notion that subjective probabilities can be usefully viewed as prices at which one might trade has been a common one in statistics, and is associated with de Finetti [1937][1970] and Savage [1971]. It is also clear, of course, in the vast literature on gambling, particularly on the setting of odds by bookies and parimutuel markets (Epstein [1977; p. 298ff.]). The central insight is that subjective probabilities of events are marginal rates of substitution between contingent claims, where the contingencies are events that the probabilities refer to. There are then a myriad of ways in which one

¹ We are not the first to propose joint estimation of subjective probabilities and utility functions in general. Viscusi and Evans [1998] introduced the general idea in an examination of how subjects appear to update latent subjective probabilities when told how objective probabilities change. They infer subjective posterior probabilities and utility curvatures from stated willingness to pay for risk reductions in hypothetical surveys, testing whether subjects completely respond to the information given or are influenced by prior beliefs. Viscusi and Evans [2006] extend this approach to compare such implied subjective probabilities to stated probabilities also observed in survey responses. As we do here, they carefully consider the requirements that are needed on the observations in order to identify the set of parameters to be estimated. We extend this original idea of joint estimation to data derived from incentivized popular scoring rules for eliciting subjective probabilities and lottery choices, so that there are real monetary consequences of different reports.

can operationalize this notion of a marginal rate of substitution.²

Scoring rules are procedures that convert a “report” by an individual into a lottery defined over the outcome of some event. The formal link between scoring rules and optimizing decisions by agents is also familiar, particularly in Savage [1971], Kadane and Winkler [1987][1988], Holt [1986] and Hanson [2003]. Jose, Nau, and Winkler [2008] explore the relationship between expected scores, expected utility, and generalized information/entropy measures for several popular scoring rules and the HARA class of utility functions. Alternatives to scoring rules include procedures like those developed for utility elicitation by Becker, DeGroot and Marschak [1964]³ (BDM) and adapted to eliciting probabilities. Karni [2009], Grether [1992], Köszegi and Rabin [2008], and Holt and Smith [2009] are examples of this alternative.⁴ All of these elicitation procedures have some potential problems. The first is the poor incentive properties around the true subjective belief. This is particularly the case for BDM procedures where under-reporting leads to very small expected lost earnings. The second is that explanations to these procedures are not naturally easy to understand for subjects, such that payoff consequences to various decisions must be made transparent.

One advantage with the BDM is that the curvature of the utility function does not confound the elicited probabilities, since the same payoffs can be used in the reference lottery and the actual lottery. Under non-EUT, however, one needs to recognize that it is the decision weights and not the subjective probabilities that are elicited. In the case of scoring rules one could instead undertake

² For example, one could elicit the p that makes the subject indifferent between a lottery paying M with probability p and m with probability $(1-p)$, for $M > m$, and a lottery paying M if the event occurs and m if it does not (Marschak [1964; p. 107ff.]). This method formally requires that one elicit indifference, which raises procedural issues that can be avoided by using the type of scoring rules investigated here.

³ See for example Rutström [1998], Harstad [2000], Plott and Zeiler [2005], and Hao and Hauser [2012] for discussions of properties of BDM.

⁴ Let there be two prizes, $x > y$. The subject reports a probability ξ , and a random number ζ is selected from the unit interval. If $\zeta \leq \xi$ the subject gets the lottery that pays off x if the event occurs, and y otherwise; if $\zeta > \xi$ the subject gets the lottery that pays off x with probability ζ and y with probability $1-\zeta$.

“calibrating adjustments” to the elicited beliefs for non-linear utility functions and/or probability weighting, as carried out in Offerman, Sonnemans, van de Kuilen and Wakker [2009; §6].⁵ Their elegant approach has a reduced form simplicity, and is agnostic about which structural model of decision making under risk one uses. The maintained assumption, however, is that any deviation in reports from objective and known probabilities transfers directly to a task where probabilities are subjective, less precise and not known with certainty. While precise and agnostic about structural assumptions, the method requires the employment of identical tasks and incentives across the natural lottery for which one wants to elicit subjective beliefs and the calibration task with known objective probabilities. Our approach is more general in the sense that it does not require two identical instruments for eliciting the beliefs and the belief calibration function.

Our approach also has the advantage of allowing a structural identification of sources of imprecision about inferences over subjective probabilities. If inferred subjective probabilities are conditioned on knowing risk attitudes from earlier tasks, or assuming risk attitudes *a priori*, then any behavioral or statistical uncertainty on the calibration *should* be allowed to “propagate” into some additional uncertainty over inferences about subjective probabilities. Our joint estimation method allows these error propagation effects to occur, as theory says they should, providing more reliable estimates of subjective probabilities, even if those estimates have large standard errors. Reliability and precision are not the same thing. One can numerically generate a precise estimate based on maintained assumptions that are false, resulting in an unreliable estimator. In other words, it is possible that the choice task for eliciting subjective probabilities generates a point response that appears to be quite precise by itself, but which is actually not a very precise estimate of the latent subjective probability when one properly accounts for uncertainty over the “calibrated” risk

⁵ The need for some correction is also recognized by Offerman, Sonnemans and Schram [1996; p.824, fn.8] and Rutström and Wilcox [2009; p.11, fn.8].

attitudes.

In section 1 we briefly state the theory underlying our approach and relate this to the literature on belief elicitation. The properties of the QSR and LSR, and the fact that responses to these are affected by risk attitudes, are well known. We assume throughout that the agent is acting in what is called a “probabilistically sophisticated” manner, although our method does not restrict the characterization of risk attitudes to expected utility theory (EUT), or to specific functional forms. We also consider the inference of subjective probabilities for subjects who are assumed to make decisions according to the rank-dependent utility (RDU) model. This extension is particularly appropriate in the case of eliciting subjective probabilities, because it involves allowing for probability weighting and non-additive decision weights on the utility of final outcomes. Given that one of the probabilities to be weighted is the subjective probability being estimated, one might expect estimates of the subjective probability to be even more sensitive to the correct specification of the model of risk attitudes employed.

In section 2 we describe the experimental task we posed to 140 subjects, split roughly equally across the QSR and LSR alternatives. Our subjects made choices over a number of standard lotteries, characterized by objective uncertainty over monetary outcomes between \$0 and \$100. The lotteries vary both in prizes and in probabilities, allowing us to identify parameters for both utility functions and probability weighting functions. Subjects also gave responses to either a QSR or LSR choice task over subjective beliefs. The prizes on each of these scoring rule tasks also spanned \$0 and \$100, so that we were able to infer risk attitudes over the same prize domain as the scoring rule responses.

Section 3 formally sets out the econometric model used for estimating subjective probabilities, spelling out the manner in which we undertake joint estimation over all tasks in order to identify subjective probabilities. Section 4 then presents our estimates of the inferred subjective

probabilities from these scoring rules, after adjusting for risk. Our primary result is that subjective probabilities inferred under the assumption of risk neutrality are very different to the subjective probabilities inferred when one allows the data to say how risk averse the subjects were. This finding has immediate implications for the practical use of scoring rules,⁶ which is the focus here; it also has implications for inferences that can be drawn from prediction markets since experimental findings consistently shows that people are risk averse on average. We offer a methodology for jointly estimating risk attitudes and subjective probabilities, and show that utility curvature and probability weighting can have opposing qualitative effects on inferences about subjective probabilities. Section 5 draws conclusions.

1. Scoring Rules

For simplicity we assume throughout that the events in question only have two outcomes. A scoring rule asks the subject to make some report θ , and then defines how an elicitor pays a subject depending on their report and the outcome of the event. This framework for eliciting subjective probabilities can be formally viewed from the perspective of a trading game between two agents: you give me a report, and I agree to pay you $\$X$ if one outcome occurs and $\$Y$ if the other outcome occurs. The scoring rule defines the terms of the exchange quantitatively, explaining how the elicitor converts the report from the subject into a lottery. We use the terminology “report” because we want to view this formally as a mechanism, and do not want to presume that the report is in fact the subjective probability π of the subject. In general, it is not.

⁶ There is a large, practical literature on the “normative” elicitation of subjective probabilities, reviewed by O’Hagen et al. [2006] and illustrated well by Shephard and Kirkwood [1994]. One characteristic of those elicitation procedures is that they involve considerable real-time, one-on-one feedback between the elicitor and the elicitee, often including results from proper scoring rules. Our approach is to better model inferences from more static, impersonal applications of those scoring rules, as a prelude to the design and evaluation of normative elicitation procedures with incentives.

The QSR was apparently first used by McKelvey and Page [1990], and later by Offerman, Sonnemans and Schram [1996], McDaniel and Rutström [2001], Nyarko and Schotter [2002], Schotter and Sopher [2003], Costa-Gomes and Weizsäcker [2008] and Rutström and Wilcox [2009].⁷ In each case the subject is implicitly or explicitly assumed to be risk-neutral.⁸ Scoring rules that are linear in the absolute deviation of the estimate have been used by Dufwenberg and Gneezy [2000], and Haruvy, Lahav and Noussair [2007]. Croson [2000] and Hurley and Shogren [2005] used scoring rules that are linear in the absolute deviation as well as providing a bonus for an exactly correct prediction. Scoring rules that provide a positive reward for an “exact” prediction and zero otherwise have been used by Charness and Dufwenberg [2006].

In many cases the inferential objective has been to test hypotheses drawn from “psychological game theory,” which rest entirely on making operational the beliefs of players in strategic games. For this purpose, and also in other applications, the elicitation of beliefs is combined with some other experimental task. Other applications include testing the hypothesis that the belief elicitation task will encourage players in a game to think more strategically (Croson [2000], Costa-Gomes and Weizsäcker [2008], Rutström and Wilcox [2009]). Of course, combining tasks in this way violates the “no stakes condition” required for the QSR to elicit beliefs reliably unless one assumes that the subject is risk neutral (Kadane and Winkler [1988], Karni and Safra [1995] and Karni [1999]). Only one of these studies employ a “spectator” treatment in which players are asked to provide beliefs but do not take part in the constituent game determining the event outcome: study

⁷ Hanson [1996] contains some important corrections to some of the claims about QSR elicitation in McKelvey and Page [1990].

⁸ McKelvey and Page [1990] augmented the scoring rule procedure with a “binary lottery” payment procedure to induce risk-neutrality. In theory the subject earns “points” in the scoring rule, which convert in a linear manner into an increased probability of winning some lottery defined over a high prize and a low prize. There is considerable controversy over the behavioral validity of this procedure, reviewed in Harrison, Martínez-Correa and Swarthout [2013].

#2 of Offerman, Sonnemans and Schram [1996].

The popular QSR is defined in terms of two positive parameters, α and β that determine a fixed reward the subject gets and a penalty for error. Assume that the possible outcomes are A or B, where B is the complement of A, that θ is the reported probability for A, and that Θ is the true binary-valued outcome for A. Hence $\Theta=1$ if A occurs, and $\Theta=0$ if it does not occur (and thus B occurs instead). The subject is paid $S(\theta | A \text{ occurs}) = \alpha - \beta(\Theta-\theta)^2 = \alpha - \beta(1-\theta)^2$ if event A occurs and $S(\theta | B \text{ occurs}) = \alpha - \beta(\Theta-\theta)^2 = \alpha - \beta(0-\theta)^2$ if B occurs. In effect, the score or payment penalizes the subject by the squared deviation of the report from the true binary-valued outcome, Θ , which is 1 and 0 respectively for A and B occurring. An omniscient seer would obviously set $\theta = \Theta$. The fixed reward is a convenience to ensure that subjects are willing to play this trading game, and the penalty function simply accentuates the penalty from not being an omniscient seer. In our experiments $\alpha = \beta = \$100$, so subjects could earn up to \$100 or as little as \$0. If they reported 1 they earned \$100 if event A occurred or \$0 if event B occurred; if they reported $\frac{3}{4}$ they earned \$93.75 or \$43.75; and if they reported $\frac{1}{2}$ they earned \$75 no matter what event occurred.

It is intuitively obvious, and also well known in the literature (e.g., Winkler and Murphy [1970] and Kadane and Winkler [1988]), that risk attitudes will affect the incentive to report one's subjective probability "truthfully" in the QSR.⁹ A sufficiently risk averse agent is clearly going to be drawn to a report of $\frac{1}{2}$, and varying degrees of risk aversion will cause varying distortions in reports from subjective probabilities. If we knew the form of the (well-behaved) utility function of the subjects, and their degree of risk aversion, we could infer back from any report what subjective

⁹ There exist mechanisms that will elicit subjective probabilities without requiring that one correct for risk attitudes, such as the procedures proposed by Köszegi and Rabin [2008; p.199], Karni [2009], Grether [1992], Holt and Smith [2009], Offerman, Sonnemans, van de Kuilen and Wakker [2009] and Hao and Hauser [2012], discussed further below. The last four employ these mechanisms in an experimental evaluation. We discussed earlier the difficulties of practical application of these methods.

probability they must have had. Indeed, this is exactly what we do below, recognizing that we only ever have estimates of their true degree of risk aversion.

The LSR is also defined in terms of two positive parameters, γ and λ , that serve as fixed rewards and penalties, respectively. The subject is paid a fixed reward less some multiple of the absolute difference between their report and what actually happened, which is also what an omniscient seer would have reported. Thus the payment is $S(\theta | A \text{ occurs}) = \gamma - \lambda(1-\theta)$ if event A occurs and $S(\theta | B \text{ occurs}) = \gamma - \lambda(\theta-0)$ if B occurs. We again set $\gamma = \lambda = \$100$, generating payoffs of \$100 or \$0 for a report of 1; \$75 and \$25 for a report of $\frac{3}{4}$; and \$50 no matter what the outcome for a report of $\frac{1}{2}$. The LSR is not a favorite of decision theorists, since a risk neutral subject would jump to corner-solution reports of 1 or 0 whenever their true beliefs were either side of $\frac{1}{2}$. But when the subject is (even modestly) risk averse, an interior solution is obtained, and we face the same issues of inference as with the QSR. The LSR is a favorite of experimental economists because of the simplicity of explaining the rule: the score for a report and an event is linear in the reported probability report, so there is no need for elaborate tables showing cryptic payoff scores for discrete reports.¹⁰

In order to avoid portfolio effects from the combined choices rewards in these scoring rule task are sometimes very, very small. For example, Nyarko and Schotter [2002] and Rutström and Wilcox [2009] gave each subject an endowment of 10 cents, from which their penalties are to be deducted. In the latter study this does not present a problem since the focus is not on analyzing the elicited beliefs but on analyzing the game play as a function of exposing subjects to a belief

¹⁰ Hanson [1996; p. 1224] provides a useful reminder that discrete implementations of proper scoring rules can also engender piecewise linear opportunity sets. He points out that certain regions of the QSR implemented by McKelvey and Page [1990] were actually LSR, and that *risk-neutral* subjects would then rationally report a probability at the extremes of that linear region, and not at the discrete alternative closest to their true belief.

elicitation process. Nevertheless, one has to worry about the incentive properties of the elicitation method once the interest is on analyzing the beliefs themselves.

The need to calibrate or control for risk aversion is often not made explicit, or is claimed to be of marginal concern. Schotter and Sopher [2003; p. 504] recognize the role of risk aversion, but appear to argue that it is not a factor behaviorally:

It can easily be demonstrated that this reward function provides an incentive for subjects to reveal their true beliefs about the actions of their opponents. Telling the truth is optimal; however, this is true only if the subjects are risk neutral. Risk aversion can lead subjects to make a “secure” prediction and place a .50 probability of each strategy. We see no evidence of this type of behavior.

Of course, evidence of subjects selecting the probability report of $\frac{1}{2}$ only shows that the subject has *extreme* risk aversion. The absence of that extreme evidence says nothing about the role that risk aversion might play in general. Only one QSR study attempts to explicitly calibrate the beliefs for non-linear utility functions and/or probability weighting: Offerman, Sonnemans, van de Kuilen and Wakker [2009; §6]. We introduce an alternative experimental method where the decision model can be identified econometrically from one set of tasks for the same subjects, or even on a different pool of subjects drawn from the same population, and then statistically integrated with the belief elicitation task, while transparently allowing error terms to propagate. The estimation of the decision model over risk and subjective probabilities is joint and simultaneous, even if one can think of the lottery choices as recursively identifying the decision model.

[INSERT FIGURE 1 HERE]

2. Experimental Design

Figure 1 illustrates the lottery choice that our subjects were given. Each subject faced 45 such choices, where prizes spanned the domain \$0 up to \$100 and probabilities for various prizes varied across each lottery. One choice was selected to be paid out at random after all choices had

been entered. Choices of indifference were resolved by rolling a die and picking one lottery, as had been explained to subject. This interface builds on the classic binary choice design of Hey and Orme [1994], and is discussed in greater detail in Harrison and Rutström [2008; Appendix B]. The lotteries were presented sequentially in 3 blocks of 15, where each block had prizes in one of three intervals between \$0 and some higher level. One level was between \$0 and \$1, the other level was between \$0 and \$10, and the third level was between \$0 and \$100. We presented the lotteries sequentially so that the subject could see that all of the lotteries in one block were for a given scale. The sequence of blocks was randomized across subjects. Complete instructions are provided in Web Appendix A and the full set of lotteries in Web Appendix D.

[INSERT FIGURE 2 HERE]

The belief tasks were presented to subjects with a novel interface that has many attractive features.¹¹ Figure 2 shows the interface for the QSR as it was presented to subjects on a computer screen and in printed instructions. The interface was explained with instructions which used a trusty old bingo cage to illustrate one underlying random process. By varying the slider the subject could choose a report, with the conditional payoffs being displayed. The formula for the scoring rule is not shown, but the subject instead just sees alternative payoffs for different reports. The subject was then taken through displays of their payoffs if they chose to report 0% or 100%.¹² Each subject participated in an unpaid training choice, in which they were told the number of orange balls in the bingo cage that was on public display, and asked to make a report and confirm it. We deliberately adopted an extremely high scale of a maximum \$1000 payoff to ensure that the subjects understood

¹¹ This interface immediately extends to other tasks with a cardinal scale that experimental economists use to elicit risk preferences, discount rates, or social preferences. We opted to use binary choices for our lottery tasks, to be consistent with the vast bulk of the literature on the elicitation of risk preferences.

¹² The display of the probability on the right side of the slider always show an integer percentage, and the earnings were always calculated on that value, thus making the continuous probability scale into an integer one. In the statistical analysis we treat the probability as taking on 101 integer values in the range 0 to 100.

that this was to be a trainer.

Each subject then participated in 7 belief elicitation tasks, knowing that one would be selected for payment. The first 3 were repetitions of the training task with orange and white ping pong balls, but with completely different distributions of orange and white ping pong balls, and serve to provide subjects with hands-on, incentivized experience in the scoring rule.¹³ We do not analyse these choices here. The fourth task was based on the outcomes of a test in psychology for empathy known as The Eyes Test (e.g., Baron-Cohen [2003]). All subjects had completed this test at the outset of the session, and the event they were asked about was whether the score that a randomly chosen man got on the Eyes Test was equal to or greater than the score that a randomly chosen woman would get. The final three tasks were based on the 2008 U.S. Presidential Election, which was to be held about one week after the session. One task was whether the outright winner of the Presidency would be a Democrat or a Republican, one task was whether the winning share of the popular vote would be at least 5 percentage points greater than the losing share, and the final task was whether the winning share of the popular vote would be at least 10 percentage points greater than the losing share.¹⁴ Our own *a priori* expectations for these subjective probabilities were just around 50%, around 80%, around 65% and less than 10%, respectively.

The events were explained in written instructions, which were also read out loud. The event based on the Eyes Test was explained as follows: “We will pick one man and one woman in the room. Do you think the man who is selected will have a higher score on the Eyes Test than the

¹³ The subjects were told that there were 60 balls in total in a publicly visible, but initially covered, bingo cage, but were not told the number of orange or white balls. The urn was uncovered and spun for 10 rotations, and then the subject had to make a report that a ball drawn at random would be orange. Applying the methods examined here to these data does not change any of our conclusions.

¹⁴ These events compare to similar events employed in popular prediction markets, inspired by Forsythe, Nelson, Neumann and Wright [1992]. See <http://www.biz.uiowa.edu/iem/> for the current version of this market, and the contracts traded in the 2008 Presidential Election.

woman who is selected?” We then explained how we would randomly select one man and one woman in the session and compare their scores. They were then asked to bet on: “That the man we select at random will have a higher score on the Eyes Test than the woman we select at random.”

After subjects had completed their bets on the bingo cage and Eyes Test tasks, the final three events, about the Presidential Elections, included the following three bets:

1. Will the next President of the United States be a Democrat?
2. Will the popular vote for the winning candidate be 5 or more percentage points greater than the popular vote for the losing candidate?
3. Will the popular vote for the winning candidate be 10 or more percentage points greater than the popular vote for the losing candidate?

We explained that the first question was about the outcome of the Electoral College vote, and not the popular vote, and that for the second and third question, we were asking if they thought that the winner of the popular vote would beat the loser by 5 or 10 percentage points or more. The subjects then completed their belief elicitation tasks for these Presidential election events, and went on to the lottery choice tasks described earlier.

We recruited 140 subjects from the student population of the University of Central Florida, split equally across the QSR and LSR treatments. The experiments were conducted in the week prior to the 2008 election, Monday October 27 through Friday October 31. Our 140 subjects earned about \$90 on average for the two paid tasks. Each session lasted around 1½ to 2 hours, and never more than 2 hours. There was considerable variation in earnings, with one subject taking home \$3 and another subject taking home \$205.

3. Econometric Model

We develop the econometric model to be estimated in three stages, with details presented in Web Appendix C. First we specify risk attitudes assuming an EUT model of latent choice, where the focus is entirely on the concavity of the estimated utility function. Second, we specify risk attitudes assuming a RDU model of latent choice, so that risk attitudes are determined by the interplay of concave utility functions and non-linear probability weighting.¹⁵ Third, we consider the joint estimation of risk attitudes and subjective probability, using either the EUT or the RDU specification.

3.1 Estimating the EUT and RDU Models

We assume an Expo-Power (EP) utility function originally proposed by Saha [1993]. Following Holt and Laury [2002], the EP function is defined as

$$u(y) = [1 - \exp(-\alpha y^{1-r})] / \alpha, \quad (1)$$

where α and r are parameters to be estimated, and y is income from the experimental choice. The EP function can exhibit increasing or decreasing relative risk aversion (RRA), depending on the parameter α : RRA is defined by $r + \alpha(1-r)y^{1-r}$, so RRA varies with income if $\alpha \neq 0$ and the estimate of r defines RRA at a zero income. This function nests CRRA (as $\alpha \rightarrow 0$) and CARA (as $r \rightarrow 0$). The functional form of utility employed here is of no importance, and any monotonic increasing function of $u(\cdot)$ could have been implemented. We can identify both of these parameters because of the variation in prizes that are included in the full lottery design, as shown in Web Appendix D.

The RDU model extends the EUT model by allowing for decision weights on lottery

¹⁵ We could just develop an RDU model and test if the estimated probability weighting is the identity function, in which case the RDU model collapses to an EUT model. However, the exposition is, in our view, simpler if one develops the models separately because of the familiarity of EUT to most economists.

outcomes. We use the same utility function specification (1), but allow for decision weights generated by a probability weighting function. We adopt the simple “power” probability weighting function proposed by Quiggin [1982], with curvature parameter γ :

$$\omega(p) = p^\gamma \tag{2}$$

So $\gamma \neq 1$ is consistent with a deviation from the conventional EUT representation.¹⁶ Any other functional form could have been used, as long as the probability function maps the probability $p \in [0,1]$ into the unit interval $[0,1]$ and the function is strictly increasing in p .

[INSERT FIGURE 3 HERE]

Figure 3 shows the manner in which the parameter γ characterizes the probability weighting function and the decisions weights used to evaluate lottery choices, when the outcomes are ranked from best to worst. Since we assume $\gamma=0.77 < 1$ in this illustration, to anticipate our estimates, the probability weighting function $\omega(p)$ is. For simplicity here we assume lotteries with 2, 3 or 4 prizes that are equally likely when we generate the decision weights. So for the case of 2 prizes, each prize has $p=1/2$; with 3 prizes, each prize has $p=1/3$; and with 4 prizes, each prize has $p=1/4$. We see the usual result, that the decision weights on the largest prizes are relatively greater than the true probability, and the decision weights on the smallest prizes are relatively smaller than the true probability, reflecting optimism over the outcomes.¹⁷

Each panel in Figure 3 is important for our analysis. For the purposes of estimating γ from the observed lottery choices with known probabilities we only need the decision weights in the right

¹⁶ We compared our results to estimations based on the inverse-S shaped function, popularized by Tversky and Kahneman [1992] and the flexible Prelec [1998] function that allows a variety of shapes including concave, convex, S-shaped and inverse S-shaped. We conclude that the power function fits better for these data. It has a better log-likelihood, and it rejects a linear probability weighting function. The shape is confirmed in our estimates of the flexible Prelec function, while the inverse-S function is close to linear, which is the best it can do to proxy a power function.

¹⁷ Hence a concave probability weighting function, as in the left panel of Figure 3, implies risk-seeking behavior, *ceteris paribus* the curvature of the utility function.

panel of Figure 3. But for the purposes of recovering a subjective probability π subject to probability weighting, as distinct from p , the induced objective probability, we instead only need the probability weighting function. In fact, we need its inverse function, since it is the π in the $\omega(\pi)$ function that we are seeking to recover in that case. We do not directly observe $\omega(p)$ or $\omega(\pi)$, but we can estimate $\omega(\cdot)$ as part of the latent structure generating the observed choices in the two types of task, implicitly assuming that $\omega(p) = \omega(\pi)$. To anticipate slightly the exposition below, once we have $\omega(\cdot)$ we can then recover π by directly applying the estimated probability weighting function, such as the one shown, for a typical γ , in the left panel of Figure 3.¹⁸

3.2. Estimating the Subjective Probability

To estimate the subjective probability π that each subject holds from LSR or QSR responses we have to assume something about how they make decisions under risk. This is obvious in theory, and the only issue then is how to operationalize that property of these scoring rules.

If they are assumed to be risk neutral, then we can directly infer the subjective probability from the report of the subject.¹⁹ This result is immediate under the QSR, but raises a problem of interpretation under the LSR if the reports are not at the corner solutions of 0% and 100%. On the other hand, any minimal level of risk aversion will suffice, under the LSR, to generate interior responses, so we assume that the subjects indeed have some minimal level of risk aversion when we

¹⁸ There is a caveat to this intuition when subjective probabilities are close to 0.5, due to a jump discontinuity in decision weights as reports under the QSR vary around 0.5. At the point where the report is 0.5 the rank ordering of the choice options reverses. For reports less than 0.5 the A option implies higher monetary rewards, but for reports greater than 0.5 it is the B option that offers the higher reward. For sufficiently large probability weighting this jump discontinuity can wreak havoc with the ability to infer latent subjective probabilities. In our applications the probability weighting is not severe, and we can initialize the maximum likelihood estimation at the solution values obtained by assuming EUT (so that directional derivatives do not need to be numerically evaluated in the problematic regions of γ and π).

¹⁹ The expression “risk neutral” here should be understood to include the curvature of the utility function *and* the curvature of the probability weighting function. So it is not just a statement about the former, unless one assumes EUT.

report “risk neutral subjective beliefs” for the LSR.

Moving to the models that allow for varying risk attitudes, we jointly estimate the subjective probability and the parameters of the core model. Assume for the moment that we have an EUT specification. The subject that selects report θ from a given scoring rule receives the following EU

$$EU_{\theta} = \pi_A \times u(\text{payout if A occurs} \mid \text{report } \theta) + (1-\pi_A) \times u(\text{payout if B occurs} \mid \text{report } \theta) \quad (3)$$

where π_A is the subjective probability that A will occur. The payouts that enter the utility function are defined by the scoring rule and of course the specific report θ , and span the interval $[\$0, \$100]$. For the QSR and a report of 75%, for example, we have

$$EU_{75\%} = \pi_A \times u(\$93.75) + (1-\pi_A) \times u(\$43.75) \quad (3')$$

For the LSR, and the same report, we have:

$$EU_{75\%} = \pi_A \times u(\$75) + (1-\pi_A) \times u(\$25) \quad (3'')$$

and so on for other possible reports. We observe the report made by the subject for QSR or LSR.

This report can take 101 different integer values defined over percentage points. Then we can calculate the likelihood of that choice given values of τ , π_A and μ , where the likelihood is the multinomial analogue of the binary logit specification used for lottery choices. We define

$$eu_{\Theta} = \exp[(EU_{\Theta}/\nu)/\mu] \quad (4)$$

for any report Θ , where μ is a Fechner error and ν is a contextual utility transformation, both discussed in Web Appendix C, and then

$$\nabla EU = eu_{\theta} / (eu_{0\%} + eu_{1\%} + \dots + eu_{100\%}) \quad (5)$$

for the specific report θ observed, analogously to the comparable expression for EUT or RDU.

We need τ and α to evaluate the utility function in (3), we need π_A to calculate the EU_{θ} in (3) for each *possible* report Θ in $\{0\%, 1\%, 2\%, \dots, 100\%\}$ once we know the utility values, and we need μ to calculate the latent indices (4) and (5) that generate the subjective probability of observing the

choice of specific report θ when we allow for some noise in that process. The *joint* maximum likelihood problem is to find the values of these four parameters that best explain observed choices in the belief elicitation tasks as well as in the lottery tasks. We estimate these parameters simultaneously, even if the lottery tasks *could* be used solely to estimate the r and α parameters, and are needed to ensure identification of all parameters.

Exactly the same logic extends to the model in which we assume an RDU latent structure instead of an EUT latent structure. In effect, the lottery task allows us to identify r and α under EUT, and r , α and γ under RDU, thanks to the variations in both prizes and probabilities in this task. Individual heterogeneity is allowed for by estimating both risk attitudes and subjective probabilities as linear functions of the demographic characteristics defined earlier.

4. Results

[INSERT FIGURE 4 HERE]

[INSERT TABLE 1 HERE]

4.1. Raw Elicited Beliefs

Figure 4 displays the raw responses from each of the scoring rules for each event, in the form of kernel densities, and Table 1 shows summary statistics of the elicited responses. The four events are the “Eyes Test,” the “President,” “win by 5%,” and “win by 10%.” The summary statistics suggest that the QSR and LSR provided roughly the same responses, but the densities in Figure 4 do have some differences in shape. In part this simply alerts us to be aware of the non-Gaussian shape of these distributions.

The general location of the densities corresponds with our qualitative priors on the subjective beliefs that were to be expected for these events. Recall that the Eyes Test asked for the probability that a typical male would score better than a typical female. For the Eyes Test it

appears, from the observation that the modal response is around 0.4 for both LSR and QSR treatments, that the sample did not expect the male score to exceed the female score, but that there was a wide variability around this modal belief. The sample appeared confident that Barack Obama would indeed win the election outright, but displayed a healthy sense of perspective on what the winning margin would be. For our purposes, the choices of a 5% and 10% threshold for the popular vote could not have worked out better, with a majority believing that a 5% margin would be attained but that a 10% margin would not. These results show, at a minimum, that responses were at least correlated with what we believe to be reasonably coherent subjective beliefs for these events.

The fact that the responses to the LSR are not at “corner” values of 0 or 1 shows that the subjects were not exactly risk neutral. But it does not show much more, because one would observe some interior response even for small amounts of risk aversion, as noted earlier.

With the exception of the bets on the outright winner of the Presidential election, the distribution of responses for the two scoring rules are roughly the same. This conclusion is supported by a Kolmogorov-Smirnov test of the null hypothesis that the two distributions are equal. The exception is the case of the outright winner event, where the p -value is only 0.011, so we can reject that hypothesis in this instance. This finding provides some support for those that would prefer to use the LSR on the grounds that it is simpler to explain to subjects than the QSR. Of course, the real issue is whether they generate the same estimates of subjective probability when one allows for risk attitudes.

4.2. Characterizing Risk Attitudes

Looking just at the lottery choices under a maintained hypothesis of EUT for now, we find evidence of modest risk aversion at low stakes (since $r=0.3>0$, and r defines RRA at $y=0$), and evidence of slightly increasing relative risk aversion as the prizes climb to \$100 (since $\alpha=0.03>0$).

Detailed results are provided in Web Appendix B, since they are only of indirect interest here. Given these parameter estimates we can calculate RRA at various prize levels: at \$25, \$50, \$75 and \$100 the RRA is estimated to be 0.49, 0.61, 0.71 and 0.81, respectively.²⁰ Thus, despite the relatively low estimate of α the implied risk aversion at higher stakes is much stronger than for the lower stakes.

When we allow for an array of covariates to better characterize the heterogeneity of risk attitudes, we observe females to be significantly more risk averse, with RRA 0.14 higher at the \$0 level. As expected, allowing for covariates has no significant average effect on the RRA by prize level.

These results suggest that one might see somewhat different effects of “risk-conditioning” in the reports for scoring rules depending on the stakes involved. Our stakes are large in relation to the literature: a maximum prize of \$100, compared to common implementations in experiments of a maximum prize of less than \$1.²¹ In applications where the interest is directly on eliciting beliefs, such small stakes will not generally result in precise estimates, so larger stakes are necessary. Our results on risk attitudes for low stakes and high stakes imply that the extent of adjustment for risk attitudes is much greater for higher stake elicitation. It is a factor for both, since we estimate RRA to be positive for the lowest stakes, but it is not as serious a factor as when the stakes are lower.

The results from estimating the RDU model are slightly different. Detailed estimates are again reported in Web Appendix B. The estimates of the utility function parameters \hat{r} and $\hat{\alpha}$ are both larger than their counterparts (r and α) under EUT, implying greater concavity of the utility function. We estimate the probability weighting parameter under RDU to be $\gamma=0.72$ without covariates, and

²⁰ We simply use the estimated model to predict the point estimate of the RRA for each subject and each of these prizes, and report the average of those RRA point estimates here.

²¹ To be fair, those low-stake implementations are often in the context of the probability elicitation task being paired with another task, such as the choice of a strategy in a game, and the stake for the probability elicitation task is kept small to avoid the subject attempting to construct a portfolio of paired responses across the two tasks.

can reject the hypothesis that this is equal to 1 (p -value <0.001). A likelihood ratio test of the hypothesis that the EUT model and the RDU model are the same when there are no covariates has a $\chi^2_1 = 35.47$ (p -value <0.01), so we reject that null. The same conclusion is true when we account for covariates and heterogeneity of responses. There are no strikingly different coefficient estimates due to demographic characteristics (even female is not significant here), and the small changes across the board do not add up to a statistically significant difference. A likelihood ratio test of the hypothesis that the EUT model and the RDU model are the same when including covariates has a $\chi^2_{14} = 47.71$ (p -value <0.01). Thus, we find support for the hypothesis of probability weighting in this case, although here we are not investigating to what extent this is generally true or just holds for portions of our subjects and/or tasks. Since the estimated γ is less than 1, we conclude that subjects are optimistic (or, which is the same, risk seeing), which explains why the utility function is more concave than for EUT when estimated on the same choice data.

One noteworthy feature of these estimates is that one can reject the CRRA specification for both EUT and RDU models in this case. Of course, one might accept CRRA if estimating risk attitudes over a much smaller income domain, such as between \$0 and \$10, or when the variation in stakes as a percentage is relatively small.

4.3. Estimating Subjective Probabilities

[INSERT TABLE 2 HERE]

Table 2 lists the main results from estimating subjective probabilities for each of the four events considered here, and assuming either an EUT or RDU specification. We pool data over LSR and QSR, and control for the effect of the scoring rule with a binary dummy variable.

Assume the EUT specification for now. Given that we find evidence of risk aversion in our subjects over the domain of prizes used in the belief elicitation tasks, our estimated subjective

probabilities are all translations of the raw responses away from the 50% response. In both the LSR and QSR we expect risk averse subjects to make choices biased towards 50%, so that when we correct for their risk attitudes the inferred probabilities should move away from 50%. The reason, again, is that risk averse subjects are drawn to respond toward 50% simply to reduce the uncertainty over payoffs, so evidence of risk aversion implies that their true, latent, subjective probabilities must be further away from 50% than their raw responses. Our maximum likelihood estimates simply impose some parametric structure on that theoretical structure, to be able to quantify the extent of the required translation and the precision of the resulting inference about the latent subjective probability.

For the RDU specification, where we find a stronger concavity of the utility function and therefore a stronger effect moving the report closer to the 50% point, there is also an *opposite effect* moving the report away from the 50% point due to the concavity of the probability function. Our results demonstrate that the net effect of these two forces varies with the closeness to the 50% point. The closer we are to a true, latent subjective probability of 50% the smaller is the utility correction effect but the stronger is the decision weight correction effect.²²

For the Eyes Test, we observe a very small movement in subjective probabilities away from the raw responses under the assumption of EUT. For any given utility curvature, the closer the true subjective probabilities are to 50% the smaller is the absolute bias generated by the tendency for aversion to outcome variability to move reports towards 50%. The effect of probability weighting in the RDU model is statistically significant, and the net effect of probability weighting and utility

²² The magnitude of the probability weighting correction close to 50% is a function of the use of a power function. We also estimated the inverse-S probability weighting function proposed by Tversky and Kahneman [1992] as well as a flexible Prelec [1998] function. The inverse-S function has both concave and convex portions and our estimates move this function as close to the everywhere-concave range as is possible, but with a much lower log likelihood than the power function. The flexible Prelec function is estimated to be concave. We therefore report results only for the power function.

curvature is to further adjust the inferred probabilities away from the 50% responses. However, the net effect masks the fact that there are two underlying forces moving in opposite direction.

Correcting for probability weighting will *increase* the estimated probability, while correcting for utility curvature will *decrease* the estimated probability. The reason that the estimated probability increases when correcting for probability weighting in this case is because the report we ask subjects to make is for what they perceive as the unlikely event, with a latent subjective probability less than 0.5. The scoring rule pays a lower prize for the unlikely event, and with optimistic probability weighting the worse outcome has a decision weight that understates the subjective belief. The estimated value of the probability weighting parameter for the Eyes test is $\gamma=0.77$, which is the value used to generate the illustration in Figure 3. Notice also that the closer we are to reports of 50% the larger is the probability weighting effect on the report. The strong decrease in the latent probability due to utility curvature is because the estimated utility curvature is so much stronger under RDU than under EUT. The 95% confidence interval does not include the raw report for this test under RDU, but does under EUT.

For the three election events, we see more interesting effects of adjusting for risk aversion. In the “win by 5%” event, which is again one in which the raw responses are relatively close to 50%, we again infer a small translation from the raw response average of 59% to 62% under EUT, and from 59% to 57% under RDU. Because we are now on the opposite side of the 50% divide, the utility correction will increase the latent probability, while the decision weight correction will cause it to decrease. The other two election events, on the other hand, illustrate when one might expect to see the effect of utility curvature exert a more significant quantitative effect because we are further away from the 50% point, and the spread in the payoffs is larger. On the other hand, we expect a smaller effect from decision weight corrections since we are closer to the tails of the cumulative distributions shown in Figure 3, and hence closer to the fixed points $\omega(\pi) = \pi$ when π is equal to 0 or

1. In the case of the outright winner event, labeled “President” in Table 2, we estimate latent subjective probabilities to be 80% (EUT) or 87% (RDU), rather than the raw response average of 71%. Moreover, the 95% confidence interval on these estimates does not include the raw response. Similarly, in the case of the chance of the popular vote for the winner being more than 10% of the popular vote of the loser, we estimate subjective probabilities of 18% (EUT) or 0.4% (RDU), rather than the raw response average of 28%.

It is possible to estimate the structural model of risk attitudes and subjective beliefs allowing for observable covariates on each parameter. To illustrate, Table B3 in Appendix B lists detailed EUT estimates for the model of the “win by 5%” event when we include covariates.²³ One reason for controlling for covariates is to allow for sample composition differences between treatments.²⁴ We observe no statistically significant effect from using the QSR or LSR for this event, but there is no reason to expect one *after we properly condition for risk attitudes*. That is, the two scoring rules simply provide subjects with different lotteries with which to place bets about their subjective beliefs. So the same subjective belief should be estimated from each treatment, once one has conditioned on risk attitudes. We find a significant difference between QSR and LSR only for the event “win by 10%”, which generates higher probability estimates in the LSR.

To put these results into perspective, particularly those for the outright winner, it is important to note that in the week of the experiments the tide of public opinion clearly favored

²³ The estimates for the subjective probability π refer to a non-linear transform in which we actually estimate the parameter κ and then convert κ to π using $\pi = 1/(1+\exp(\kappa))$. Thus κ can vary between $\pm\infty$ and π is constrained to the open unit interval. To interpret these coefficients, $\kappa=0$ implies $\pi=1/2$, $\kappa>0$ implies $\pi<1/2$, and $\kappa<0$ implies $\pi>1/2$. The estimated subjective probabilities we report have been converted back from κ to π using this non-linear function and the “delta method” to correctly calculate standard errors (Oehlert [1992]). In addition, it is the linear function of κ that is constrained by this transform to be in the unit interval, not each element of that function. Thus, in Table B3 the constant term for π has a statistically significant coefficient of -0.7, which would violate that constraint *if there were no covariates*.

²⁴ For example, men and women might have different risk attitudes or subjective beliefs, and the mix of men and women could vary from treatment to treatment. This can occur even with randomization to treatment, particularly when considering a wide range of covariates.

Barack Obama to win in a landslide. Eliciting a raw belief of only a 70% chance of Obama winning is therefore puzzling, and an *a priori* challenge for those that would use the raw results from a QSR or LSR procedure. For example, consider the Iowa Presidential Market, the current version of the prediction market developed by Forsythe, Nelson, Neumann and Wright [1992]. Average daily prices on this market for the month of October 2008, and specifically the week in which our experiments were conducted, implied that the market probability of Obama winning was around 85%. Moreover, this was the prevailing sense of the market for at least 2 weeks prior to our experiments.²⁵ Indeed, the online Irish betting house, PaddyPower.com, was already paying out on over €1 million in pro-Obama bets as early as October 17! However, calibrating these inferred market probabilities to reflect the fact that most populations are risk averse on average would imply an even higher subjective probability.²⁶

The prices from the Iowa Presidential Market on the popular vote share suggested that a 10% difference was possible. On the other hand, one striking feature of the two contracts, the “winner take all” contract and the “popular vote share” contract, is that the latter was poorly traded in comparison to the former as measured by volume. Of course, that could reflect a market that has found an equilibrium price, but it also could reflect a market in which there is too much uncertainty about the outcome for traders to feel safe making a bet. Additionally, correcting for the risk preferences of the average market participant would imply a much lower probability, consistent again with what we report here.

²⁵ And the same was true generally, and consistently, for 9 months prior to the 2012 Presidential election, even if the popular vote was relatively close.

²⁶ Fountain and Harrison [2011] examine many ways on which averages or medians from prediction markets might not reflect the average or median of the aggregate distribution of beliefs, apart from heterogeneity of risk attitudes.

5. Conclusions

We demonstrate how one can apply the theory of subjective probability elicitation by means of scoring rules. Our experimental design shows how one can pair different types of choice tasks to allow estimation of risk attitudes, which can then be used to condition the inferences from responses to the scoring rule tasks. Our method involves two simple types of choice tasks that allow us to identify not just the curvature of utility but also of probability weighting, and to control for these when inferring subjective beliefs. Our structural econometric model shows how maximum likelihood methods can then be used to estimate subjective probabilities while properly controlling for both the heterogeneous risk attitudes and for the statistical errors of the latter. We applied this approach to elicit subjective probabilities over four naturally occurring events from a sample of 140 subjects. We find that it is important to correct for risk attitudes both through the utility curvature and probability curvature. For this data we find that both utility functions and probability functions are concave, so that the corrections of probability reports work in opposite directions, at least for likely events. In addition, the magnitude of the utility correction is smallest close to 50%, while this is where the magnitude of the probability correction is the strongest. Once we correct for risk attitudes using our joint estimation approach we find no consistent significant difference between the Quadratic and Linear Scoring rules. This makes the Linear Scoring rule appealing due to the simplicity of explaining it to subjects.

Our results show that one has to be sensitive to the risk attitudes of subjects before drawing inferences about subjective probabilities from responses to scoring rules. One cannot just directly treat the responses to the scoring rule as if it is a subjective probability, unless one is willing *a priori* to make striking assumptions about risk attitudes. Those assumptions are rejected in our data.

Quite apart from inferring the correct point estimate of subjective probability, uncertainty about risk attitudes affects the confidence that one should have in any point estimate. Even if

subjects are “approximately risk neutral” on average, and the QSR is used, uncertainty about the precise level of risk attitudes should be properly reflected in uncertainty about inferences over subjective probabilities. Our analysis has demonstrated how to combine theory and econometrics to do just that. The choice task for eliciting subjective probabilities generates a point response that might appear to be quite precise by itself, but which is actually not a very precise estimate of the latent subjective probability when one properly accounts for uncertainty over risk attitudes.

Of course, although the estimation of subjective probabilities is an important objective in itself, the issue of how to best characterize subjective uncertainty, and attitudes towards it, involve deeper issues. Our estimation approach is clearly within the conventional Bayesian subjective probability framework of Savage [1971][1972]. That framework provides one obvious point of departure for criticisms of EUT based on ideas that subjective *beliefs* should not be represented by subjective *probabilities*. Exactly how one then models subjective beliefs is an open and important area of research (e.g., Smith [1969], Gilboa and Schmeidler [1989], Ghirardoto, Maccheroni and Marinacci [2004], Klibanoff, Marinacci and Mukerji [2005], Nau [2007] and Gilboa, Postlewaite and Schmeidler [2008]). Further, it seems plausible that subjects exhibit some degree of “uncertainty aversion” in addition to traditional risk aversion when faced with making decisions about events that involve subjective probabilities rather than objective probabilities. We leave these issues to future research.

References

- Baron-Cohen, S. (2003). *The Essential Difference: The Truth About the Male and Female Brain*. New York: Basic Books.
- Becker, G. M., DeGroot, M. H. & Marschak, J. (1964). Measuring Utility By A Single-Response Sequential Method. *Behavioral Science*, 9, 226-232.
- Charness, G. & Dufwenberg, M. (2006). Promises and Partnership. *Econometrica*, 74, 1579-1601.
- Costa-Gomes, M. A., & Weizsäcker, G. (2008). Stated Beliefs and Play in Normal-Form Games. *Review of Economic Studies*, 75, 729-762.
- Croson, R. (2000). Thinking like a game theorist: Factors affecting the frequency of equilibrium play. *Journal of Economic Behavior and Organization*, 41, 299-314.
- de Finetti, B. (1937) La Prévision: Ses Lois Logiques, Ses Sources Subjectives, *Annales de l'Institut Henri Poincaré*, 7, 1-68; English translation as: Foresight: Its Logical Laws, Its Subjective Sources. In H.E. Kyburg and H.E. Smokler (eds.), *Studies in Subjective Probability*. Huntington, NY: Robert E. Krieger.
- de Finetti, B. (1970). Logical Foundations and Measurement of Subjective Probability. *Acta Psychologica*, 34, 129-145.
- Dufwenberg, M. & Gneezy, U. (2000). Measuring beliefs in an experimental lost wallet game. *Games & Economic Behavior*, 30, 163-182.
- Epstein, R. A. (1977). *The Theory of Gambling and Statistical Logic*. San Diego: Academic Press.
- Forsythe, R., Nelson, F., Neumann, G. R. & Wright, J. (1992). Anatomy of an Experimental Political Stock Market. *American Economic Review*, 82, 1142-1161.
- Fountain, J., & Harrison, G. W. (2011). What Do Prediction Markets Predict?. *Applied Economics Letters*, 18, 267-272.
- Ghirardoto, P., Maccheroni, F. & Marinacci, M. (2004). Differentiating Ambiguity and Ambiguity Attitude. *Journal of Economic Theory*, 118, 133-173.
- Gilboa, I. & Schmeidler, D. (1989). Maxmin Expected Utility with a Non-Unique Prior. *Journal of Mathematical Economics*, 18, 141-153.
- Gilboa, I., Postlewaite, A. P. & and Schmeidler, D. (2008). Probability and Uncertainty in Economic Modeling. *Journal of Economic Perspectives*, 22, 173-188.
- Grether, D. M., (1992). Testing Bayes Rule and the Representativeness Heuristic: Some Experimental Evidence. *Journal of Economic Behavior & Organization*, 17, 31-57.

- Hanson, R. (1996). Correction to McKelvey and Page, 'Public and Private Information: An Experimental Study of Information Pooling.' *Econometrica*, 64, 1223-1224.
- Hanson, R. (2003). Combinatorial Information Market Design. *Information Systems Frontiers*, 107-119.
- Hao, L. & Houser, D. (2012). Belief Elicitation in the Presence of Naïve Respondents: An Experimental Study. *Journal of Risk and Uncertainty*, 44(2), April 2012, 161-180.
- Harrison, G. W. (1992). Theory and Misbehavior of First-Price Auctions: Reply. *American Economic Review*, 82, 1426-1443.
- Harrison, G. W., Martínez-Correa, J. & Swarthout, J. T. (2013). Inducing Risk Neutral Preferences with Binary Lotteries: A Reconsideration. *Journal of Economic Behavior & Organization*, forthcoming.
- Harrison, G. W., & Rutström, E. E. (2008). Risk Aversion in the Laboratory. In J.C. Cox and G.W. Harrison (eds.), *Risk Aversion in Experiments*, 12. Bingley : Emerald Group Publishing Limited.
- Harstad, R. M. (2000). Dominant Strategy Adoption and Bidders' Experience with Pricing Rules. *Experimental Economics*, 3, 261-280.
- Haruvy, E., Lahav, Y. & Noussair, C. (2007). Traders' Expectations in Asset Markets: Experimental Evidence. *American Economic Review*, 97, 1901-1920.
- Hey, J. D. & Orme, C. (1994). Investigating Generalizations of Expected Utility Theory Using Experimental Data. *Econometrica*, 62, 1994, 1291-1326.
- Holt, C. A. (1986). Scoring Rule Procedures for Eliciting Subjective Probability and Utility Functions. In P.K. Goel and A. Zellner (eds.), *Bayesian Inference and Decision Theory: Essays in Honor of Bruno de Finetti*. Amsterdam: North-Holland.
- Holt, C. A. & Laury, S. K. (2002). Risk Aversion and Incentive Effects. *American Economic Review*, 92, 1644-1655.
- Holt, C. A., & Smith, A. M. (2009). An Update on Bayesian Updating. *Journal of Economic Behavior & Organization*, 69, 125-134.
- Hurley, T. M., & Shogren, J. F. (2005). "An Experimental Comparison of Induced and Elicited Beliefs. *Journal of Risk and Uncertainty*, 30,169-188.
- Jose, V. R. R., Nau, R. F. & Winkler, R. L. (2008). Scoring Rules, Generalized Entropy, and Utility Maximization. *Operations Research*, 56, 1146-1157
- Kadane, J. B. & Winkler, R. L. (1987). de Finetti's Method of Elicitation. In R. Viertl (ed.), *Probability and Bayesian Statistics*. New York: Plenum.

- Kadane, J. B. & Winkler, R. L. (1988). Separating Probability Elicitation from Utilities. *Journal of the American Statistical Association*, 83, 357-363.
- Karni, E. (1999). Elicitation of Subjective Probabilities when Preferences are State-Dependent. *International Economic Review*, 40, 479-486.
- Karni, E. (2009). A Mechanism for Eliciting Probabilities. *Econometrica*, 77, 603-606.
- Karni, E. & Safra, Z. (1995). The Impossibility of Experimental Elicitation of Subjective Probabilities. *Theory and Decision*, 38, 313-320.
- Klibanoff, P., Marinacci, M. & Mukerji, S. (2005). A Smooth Model of Decision Making Under Ambiguity. *Econometrica*, 73, 1849-1892.
- Köszegi, B. & Rabin, M. (2008). Revealed Mistakes and Revealed Preferences. In A. Caplin and A. Schotter (eds.), *The Foundations of Positive and Normative Economics: A Handbook*. New York: Oxford University Press.
- Marschak, J. (1964). Actual Versus Consistent Decision Behavior. *Behavioral Science*, 9, 103-110.
- McDaniel, T. M. & Rutström, E. E. (2001). Decision Making Costs and Problem Solving Performance. *Experimental Economics*, 4, 145-161.
- McKelvey, R. D. & Page, T. (1990). Public and Private Information: An Experimental Study of Information Pooling. *Econometrica*, 58, 1321-1339.
- Nau, R. F. (2007). Extensions of the Subjective Expected Utility Model. In W. Edwards, R. Miles, Jr., and D. von Winterfeldt (eds.), *Advances in Decision Analysis: From Foundations to Applications*. New York: Cambridge University Press.
- Nyarko, Y. & Schotter, A. (2002). An Experimental Study of Belief Learning Using Elicited Beliefs. *Econometrica*, 70, 2002, 971-1005.
- Oehlert, G. W. (1992). A Note on the Delta Method. *The American Statistician*, 46, 27-29.
- O'Hagen, A., Buck, C. E., Daneshkhah, A., Eiser, J. R., Garthwaite, P. H., Jenkinson, D. J., Oakley, J. E. & Rakow, T. (2006). *Uncertain Judgements: Eliciting Experts' Probabilities*. Hoboken, NJ: Wiley.
- Offerman, T., Sonnemans, J. & Schram, A. (1996). Value Orientations, Expectations and Voluntary Contributions in Public Goods. *Economic Journal*, 106, 817-845.
- Offerman, T., Sonnemans, J., van de Kuilen, G. & Wakker, P. P. (2009). A Truth-Serum for Non-Bayesians: Correcting Proper Scoring Rules for Risk Attitudes. *Review of Economic Studies*, 76, 1461-1489.

- Plott, C. R. & Zeiler, K. (2005). The Willingness to Pay-Willingness to Accept Gap, the 'Endowment Effect,' Subject Misconceptions, and Experimental Procedures for Eliciting Valuations. *American Economic Review*, 95, 530-545.
- Prelec, D. (1998). The Probability Weighting Function, *Econometrica*, 66, 497-527.
- Quiggin, J. (1982). A Theory of Anticipated Utility. *Journal of Economic Behavior & Organization*, 3, 323-343.
- Rutström, E. E. (1998). Home-Grown Values and the Incentive Compatible Auction Design. *International Journal of Game Theory*, 27, 427-441.
- Rutström, E. E. & Wilcox, N. T. (2009). Stated Beliefs Versus Empirical Beliefs: A Methodological Inquiry and Experimental Test. *Games and Economic Behavior*, 67, 616-632.
- Saha, A. (1993). Expo-Power Utility: A Flexible Form for Absolute and Relative Risk Aversion. *American Journal of Agricultural Economics*, 75, 905-913.
- Savage, L. J. (1971). Elicitation of Personal Probabilities and Expectations. *Journal of American Statistical Association*, 66, 783-801.
- Savage, L. J. (1972). *The Foundations of Statistics*. New York: Dover Publications.
- Schotter, A. & Sopher, B. (2003). Social Learning and Coordination Conventions in Intergenerational Games: An Experimental Study. *Journal of Political Economy*, 111, 498-529.
- Shephard, G. G. & Kirkwood, C. W. (1994). Managing the Judgemental Probability Elicitation Process: A Case Study of Analyst/Manager Interaction. *IEEE Transactions on Engineering Management*, 41, 414-425.
- Smith, V. L. (1969). Measuring Nonmonetary Utilities in Uncertain Choices: the Ellsberg Urn. *Quarterly Journal of Economics*, 83, 324-329.
- Tversky, A. & Kahneman, D. (1992). Advances in Prospect Theory: Cumulative Representations of Uncertainty. *Journal of Risk & Uncertainty*, 5, 297-323.
- Viscusi, W. K. & Evans, W. N. (1998). Estimation of Revealed Probabilities and Utility Functions for Product Safety Decisions. *Review of Economics and Statistics*, 80, 28-33.
- Viscusi, W. K. & Evans, W. N. (2006). Behavioral Probabilities. *Journal of Risk and Uncertainty*, 32, 5-15.
- Winkler, R. L. & Murphy, A. H. (1970). Nonlinear Utility and the Probability Score. *Journal of Applied Meteorology*, 9, 143-148.

Table 1: Descriptive Statistics for Scoring Rule Responses

Event	Scoring Rule	Mean	Median	Standard Deviation
Eyes Test	Quadratic	0.43	0.4	0.19
	Linear	0.43	0.4	0.19
	Both	0.43	0.4	0.19
President	Quadratic	0.69	0.7	0.2
	Linear	0.74	0.75	0.22
	Both	0.71	0.7	0.21
Win by 5%	Quadratic	0.59	0.6	0.23
	Linear	0.58	0.6	0.26
	Both	0.59	0.6	0.24
Win by 10%	Quadratic	0.28	0.3	0.2
	Linear	0.29	0.26	0.21
	Both	0.28	0.3	0.21

Table 2: Estimated Subjective Probabilities

Event	Specification (Log-Likelihood)	Point Estimate	Standard Error or Standard Deviation	95% Confidence Interval
Eyes Test	EUT (-4563.2)	0.42	0.022	0.38/0.46
	RDU (-4550.6)	0.31	0.037	0.24/0.39
	Raw Responses	0.43	0.19	0/0.95
President	EUT (-4558.1)	0.8	0.033	0.74/0.87
	RDU (-4542.1)	0.87	0.079	0.72/1
	Raw Responses	0.71	0.21	0.20/1
Win by 5%	EUT (-4593.8)	0.62	0.030	0.56/0.68
	RDU (-4580.3)	0.57	0.048	0.47/0.67
	Raw Responses	0.59	0.24	0.10/1
Win by 10%	EUT (-4560.9)	0.18	0.038	0.11/0.25
	RDU (-4539.5)	0.004	0.024	0/0.05
	Raw Responses	0.28	0.2	0/0.77

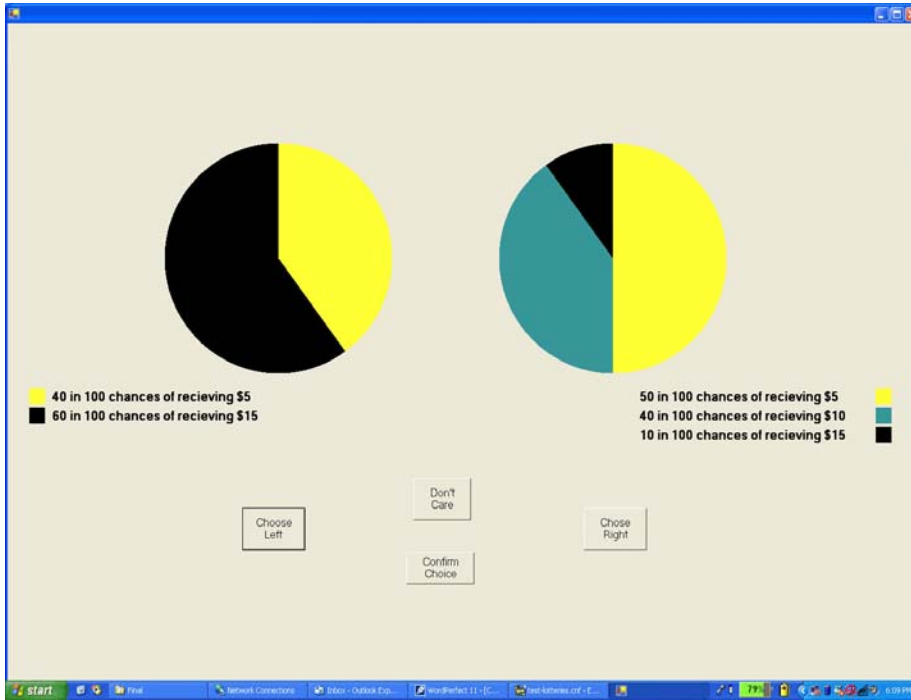


Figure 1: Illustrative Lottery Choice




Figure 2: Illustrative Quadratic Scoring Rule Interface




Figure 3: Probability Weighting and Decision Weights




Figure 4: Elicited Responses from QSR and LSR

Appendix A: Instructions (NOT FOR PUBLICATION)

We provide complete instructions for the introduction to the session (marked I in the top, right corner of the first page of instructions, the quadratic scoring rule task trainer (marked q), the linear scoring rule trainer (marked L), the actual belief elicitation tasks for which subjects are paid (marked sr100), and then the lottery choice tasks (marked LOT). Copies of the exact instructions, which were printed in color, are available on request. Each subject received either the QSR or the LSR instructions.

We employed the following rough timing of each session, to ensure that subjects were finished within the 2-hour time period we had promised when they were recruited:

- 10 minutes demographics
- 15 minutes Eyes Test
- 15 minutes demo of ping pong
- 20 minutes 7 belief tasks including resolving the 3 ping pong cases
- 60 minutes total

- 10 minutes first walk through recording and rolling dice start Hey-Orme
- 10 minutes read Hey-Orme instructions
- 15 minutes Hey-Orme task
- 35 minutes total

- 95 minutes grand total.

YOUR INSTRUCTIONS

I

Welcome! Today we will be asking you several types of questions. Some of these will earn you cash, which we will pay you today. And some may earn you cash which we will pay you in a few weeks. You have already earned \$5 for showing up and agreeing to participate.

There are basically four stages today:

1. We will ask you a series of questions about yourself, such as some basic information about your age. The computer will prompt you for these questions, and you should just work through them at your own pace when we log you in.
2. We will then pause, and provide some instructions on the next task, which involves you making some judgements about what someone in a picture is thinking. We will explain that task when we come to it.
3. We will then pause, and provide more instructions on the next task, which will involve you placing some bets on things that have yet to happen. In this stage we will take small breaks between the bets you place, so that we may explain the next specific thing that you are to bet on. These choices will directly affect your earnings. Nothing comes out of your own pocket.
4. We will then pause, and provide more instructions on some choices you are to make over different amounts of money that have different chances of occurring. These choices will also directly affect your earnings.

The instructions for the second, third and fourth stage will provide more information on the type of choices you are being asked to make.

The experimenters will then collate all of your earnings and pay you for the money you have earned, as well as provide a receipt for any earnings that will be paid in the future.

Your choices are private, and will only be associated with an ID that we will enter when we log you in to the computer. So your name, address and SSN will not be linked to any choices you make. We will pay you privately, one at a time, at the end to keep your earnings private.

Are there any questions? If not, go ahead and answer the questions until the computer pauses and asks for a password. When everyone is finished this stage we will announce the password and we can go on to the second stage. There is no hurry, so take your time.

YOU MAY NOW PROCEED WITH THE FIRST STAGE.

INSTRUCTIONS (CONTINUED)

9

In this stage we will give you tasks where you will place bets on the outcome of events that will happen today or in the future. For example, who will be the next U.S. President? You can make more money the more accurately you can predict these outcomes.

You place these bets on a screen like the one below. In a moment we will let you practice with this screen on your computer. Remember, any betting you do today is with our money, not your money.

APPENDIX FIGURE 1 HERE

You place your bets by sliding the bar at the bottom of the screen until you are happy with your choice of a probability report. The computer will start at some point on this bar at random: in the above screen it started at 71%, but you are free to change this as much as you like. In fact, you should slide this bar and see the effects on earnings, until you are happy to confirm your choice.

ENTER THE PASSWORD THAT IS BEING ANNOUNCED NOW.

In this hypothetical example the maximum payoff you can earn is \$1,000. In the actual tasks the maximum payoff will be lower than that, and we will tell you what it is when we come to those tasks. But the layout of the screen will be the same.

In this demonstration, the event you are betting on is whether a Ping Pong ball drawn from a bingo cage will be Orange or White. We have a bingo cage here, and we have 20 ping-pong balls. 15 of the balls are white, and 5 of the balls are orange. We will give the cage a few turns to scramble them up, and then select one ball at random.

What we want you to do is place a bet on which color will be picked. At the top of your screen we tell you what the event is: in this case, it is **Picking a Ping-Pong Ball**, and you need to bet on whether you think it will be **Orange or White**.

Your possible earnings are displayed in the two bars in the main part of the screen, and also in numbers at the bottom of each bar. For example, if you choose to report 71% you can see that you would earn \$915.90 if the Ping Pong Ball was ORANGE, and \$495.90 if the Ping Pong Ball was WHITE.

Lets see what happens if you make different reports. If you chose to report 0% or 100% here is what you would see, and earn:

APPENDIX FIGURE 2 HERE + APPENDIX FIGURE 3 BESIDE HERE

These screens are a little small, but you can see that these two reports lead to extreme payoffs. The “good news” is the possible \$1,000 payoff, but the “bad news” is the possible \$0 payoff. In between the reports of 0% and 100% you will have some positive payoff no matter what happens, but it will vary, as you can see from the report of 71%.

INSTRUCTIONS (CONTINUED)

L

In this stage we will give you tasks where you will place bets on the outcome of events that will happen today or in the future. For example, who will be the next U.S. President? You can make more money the more accurately you can predict these outcomes.

You place these bets on a screen like the one below. In a moment we will let you practice with this screen on your computer. Remember, any betting you do today is with our money, not your money.

APPENDIX FIGURE 4 HERE

You place your bets by sliding the bar at the bottom of the screen until you are happy with your choice of a probability report. The computer will start at some point on this bar at random: in the above screen it started at 71%, but you are free to change this as much as you like. In fact, you should slide this bar and see the effects on earnings, until you are happy to confirm your choice.

ENTER THE PASSWORD THAT IS BEING ANNOUNCED NOW.

In this hypothetical example the maximum payoff you can earn is \$1,000. In the actual tasks the maximum payoff will be lower than that, and we will tell you what it is when we come to those tasks. But the layout of the screen will be the same.

In this demonstration, the event you are betting on is whether a Ping Pong ball drawn from a bingo cage will be Orange or White. We have a bingo cage here, and we have 20 ping-pong balls. 15 of the balls are white, and 5 of the balls are orange. We will give the cage a few turns to scramble them up, and then select one ball at random.

What we want you to do is place a bet on which color will be picked. At the top of your screen we tell you what the event is: in this case, it is **Picking a Ping-Pong Ball**, and you need to bet on whether you think it will be **Orange or White**.

Your possible earnings are displayed in the two bars in the main part of the screen, and also in numbers at the bottom of each bar. For example, if you choose to report 71% you can see that you would earn \$710 if the Ping Pong Ball was ORANGE, and \$290 if the Ping Pong Ball was WHITE.

Lets see what happens if you make different reports. If you chose to report 0% or 100% here is what you would see, and earn:

APPENDIX FIGURE 5 HERE + APPENDIX FIGURE 6 BESIDE HERE

These screens are a little small, but you can see that these two reports lead to extreme payoffs. The “good news” is the possible \$1,000 payoff, but the “bad news” is the possible \$0 payoff. In between the reports of 0% and 100% you will have some positive payoff no matter what happens, but it will vary, as you can see from the report of 71%.

Summarizing, then, there are two important points for you to keep in mind when placing your bets:

1. **Your belief about the chances of each outcome is a personal judgment that depends on information you have about the different events.** In this practice example, the information you have consists of the total number of Orange balls and White balls.
2. **Your choices might also depend on your willingness to take risks or to gamble.** There is no right choice for everyone. For example, in a horse race you might want to bet on the longshot since it will bring you more money if it wins. On the other hand, you might want to bet on the favorite since it is more likely to win something.

For each task, your choice will depend on two things: your judgment about how likely it is that each outcome will occur, and how much you like to gamble or take risks.

You will now make your report in this practice round. When you have chosen the report, confirm your bet by clicking on the OK tab.

After you click OK, a special box will come up which causes the program to pause. We will do this after every series of bets, and then explain what the next few bets are about. We will tell you what the password is when we are all ready to proceed. There is plenty of time, so there is no need to rush.

When everyone has placed their bets we will pick the ball and you will see what your earnings would have been if this had been for money. After that we will go on with the bets for which you can earn real money.

Does anyone have any questions?

ENTER THE PASSWORD THAT IS BEING ANNOUNCED NOW.

We are now ready to begin the choices for which you will be paid. There will be several sets of choices. In each case we will describe the bet to you, and then you can place your bets. Then we will explain the next couple of bets to you, and you place those bets.

Some of these bets will be about outcomes we know today, here and now, and some will be about outcomes we will only know in a few weeks. There will be 7 bets in all. We will pay you for one of these 7 bets. We will pick this bet at random after all bets are made, and tell you which one will be paid. You should view each bet as if it could be the one to determine your payoffs, since one of them actually will.

The maximum payoff for your bets today will be \$100.

Ping Pong Balls Again

We will now repeat the task with Ping Pong balls a few times.

We have a number of ping pong balls in each of three bingo cages, which we have labeled Cage A, Cage B and Cage C. Some of the ping pong balls are Orange and some are White. We will roll each bingo cage and you can decide for yourself what fraction of Orange balls you think are in the cage. Of course, the balls will be rolling around, and you may not be able to tell exactly how many Orange balls are in the cage. You will be asked to bet on the color of one ping pong ball, selected at random after you all place your bets. For example, if there are 20 Orange balls and 80 White balls, the chance of an Orange ball being picked at random is $20 \div 100$, or 20%.

We will do this task 3 times, with 3 different bingo cages. Just be sure that you check which cage you are placing a bet on. You can see this listed in the top left corner of your screen, where it refers to Cage A, Cage B or Cage C. We will show you each cage one at a time, and allow you to place your bets after we show it to you.

Do you have any questions?

The Eyes Test

At the beginning of today's experimental session we asked you to answer 36 questions, called **The Eyes Test**. These questions were designed by psychologists to measure a person's ability to "read someone else's mind" by just looking at their eyes.

Each and everyone of you were given the same 36 Eyes Test questions in today's experiment and a total score was recorded for each and every one of you in this experiment.

Now we come to the outcome we want you to place bets on in this portion of the experiment. We will pick one man and one woman in the room. Do you think the man who is selected will have a higher score on the Eyes Test than the woman who is selected?

After everyone in the experiment has made their bets for this event we will randomly select one man, and randomly select one woman from this experimental session. We will use the cards we collected, and sort them into one pile for men and one pile for women. Each pile will be shuffled, and one card drawn from each pile. We will then compare the score for the man that is drawn with the score for the woman that is drawn, and write these scores up on the board for you to see.

We therefore pose the following outcome for you to bet on now:

That the man we select at random will have a higher score on the Eyes Test than the woman we select at random.

Do you have any questions?

2008 Presidential Elections

We want you to place bets on some questions about the U.S. Presidential Elections being held in a few weeks:

5. Will the next President of the United States be a Democrat or a Republican?
6. Will the popular vote for the winning candidate be 5 or more percentage points greater than the popular vote for the losing candidate?
7. Will the popular vote for the winning candidate be 10 or more percentage points greater than the popular vote for the losing candidate?

It is important that you understand that the first question is about the outcome of the Electoral College vote, and not the popular vote. The popular vote is just the sum of all votes across the United States. We are only referring to the Presidential Election, and not to any other elections that might occur on the same day.

For the second and third question, we are asking if you think that the winner of the popular vote will beat the loser by 5 or 10 percentage points or more. For example, if the winner of the popular vote gets 51% of the vote and the loser gets 49%, then this is a 2 percentage point difference. If the winner gets 53% and the loser gets 47%, then this is a 6 percentage point difference.

The election will be on Tuesday, November 4, 2008. To use a widely respected public source for the outcome, we will use the New York Times of Friday, November 7, 2008 as the official source used to determine your payoffs. In the event that there is a drawn out determination of the next President, such as in the 2000 election, we will delay payments until Inauguration Day, which is on January 20, 2009.

You will be paid for your bets in checks that will be mailed out on Monday, November 10, assuming we know who the next President will be at that time.

Please go ahead now and place your bets for this event, unless you have any questions.

This is the final stage of today's experiment. You will be asked to choose between lotteries with varying prizes and chances of winning. You will be presented with a series of lotteries where you will make choices between pairs of them. There are 45 pairs in the series. For each pair of lotteries, you should indicate which of the two lotteries you prefer to play. You will actually get the chance to play one of the lotteries you choose, and will be paid according to the outcome of that lottery, so you should think carefully about which lotteries you prefer.

Here is an example of what the computer display of such a pair of lotteries will look like. The display on your screen will be bigger and easier to read.

APPENDIX FIGURE 7 HERE

The outcome of the lotteries will be determined by the draw of a random number between 1 and 100. Each number between (and including) 1 and 100 is equally likely to occur. In fact, you will be able to draw the number yourself using two 10-sided dice.

In the above example the left lottery pays five dollars (\$5) if the number on the dice that is rolled is between 1 and 40, and pays fifteen dollars (\$15) if the number is between 41 and 100. The yellow color in the pie chart corresponds to 40% of the area and illustrates the chances that the number on the dice rolled will be between 1 and 40 and your prize will be \$5. The black area in the pie chart corresponds to 60% of the area and illustrates the chances that the number on the dice rolled will be between 41 and 100 and your prize will be \$15.

We have selected colors for the pie charts such that a darker color indicates a higher prize. White will be used when the prize is zero dollars (\$0).

Now look at the pie in the chart on the right. It pays five dollars (\$5) if the number on the dice rolled is between 1 and 50, ten dollars (\$10) if the number is between 51 and 90, and fifteen dollars (\$15) if the number is between 91 and 100. As with the lottery on the left, the pie slices represent the fraction of the possible numbers which yield each payoff. For example, the size of the \$15 pie slice is 10% of the total pie.

Each pair of lotteries is shown on a separate screen on the computer. On each screen, you should indicate which of the lotteries you prefer to play by clicking on one of the three boxes beneath the lotteries. You should click the LEFT box if you prefer the lottery on the left, the RIGHT box if you prefer the lottery on the right, and the DON'T CARE box if you do not prefer one or the other.

You should approach each pair of lotteries as if it is the one out of the 45 that you will play out. If you chose DON'T CARE in the lottery pair that we play out, you will pick one using a 10-sided die, where the numbers 1-5 correspond to the left lottery and the numbers 6-10 to the right lottery.

After you have worked through all of the pairs of lotteries, raise your hand and an experimenter will come over. You will then roll two 10-sided die to determine which pair of lotteries that will be played out. You roll the die until a number between 1 and 45 comes up, and that is the lottery pair to be played. If you picked DON'T CARE for that pair, you will use the 10-sided die to

decide which one you will play. Finally, you will roll the two 10-sided dice to determine the outcome of the lottery you chose.

For instance, suppose you picked the lottery on the left in the above example. If the random number you rolled was 37, you would win \$5; if it was 93, you would get \$15. If you picked the lottery on the right and drew the number 37, you would get \$5; if it was 93, you would get \$15.

Therefore, your payoff is determined by three things:

- by which lottery pair is chosen to be played out in the series of 45 such pairs using the two 10-sided die;
- by which lottery you selected, the left or the right, for that pair; and
- by the outcome of that lottery when you roll the two 10-sided die.

This is not a test of whether you can pick the best lottery in each pair, because none of the lotteries are necessarily better than the others. Which lotteries you prefer is a matter of personal taste. The people next to you will have different lotteries in front of them when you make your choices, and may have different tastes, so their responses should not matter to you. Please work silently, and make your choices by thinking carefully about each lottery.

All payoffs are in cash, and are in addition to the \$5 show-up fee that you receive just for being here and any earnings from the previous stage.

We will now come around to your computer and get you started. When you are finished, please signal someone to come around to play out your lottery and record your earnings. As soon as you have finished the actual series, and after you have rolled the necessary dice, you will be asked to check with someone in the next room to make sure that your earnings sheet is complete, and then you will be paid. You are then free to go. Thanks again for your participation today!

Appendix B. Additional Statistical Results (NOT FOR PUBLICATION)

Table B1: Estimates of the Utility Function Under EUT

Variable	Description	Estimate	Standard Error	<i>p</i> -value	Lower 95% Confidence Interval	Upper 95% Confidence Interval
A. No Covariates						
r		0.30	0.03	<0.01	0.23	0.36
α		0.03	0.01	<0.01	0.02	0.04
μ		0.07	0.00	<0.01	0.07	0.08
B. Adding Covariates						
r						
Constant		0.23	0.06	<0.01	0.12	0.35
female	Female	0.14	0.07	0.04	0.01	0.28
Over22	Over 22 Years of age	0.09	0.10	0.37	-0.11	0.29
GPAhi	GPA great than 3.75	0.05	0.09	0.57	-0.13	0.23
Graduate	Graduate Student	-0.16	0.12	0.17	-0.40	0.07
α						
Constant		0.02	0.01	0.02	0.00	0.03
female	Female	0.02	0.01	0.18	-0.01	0.05
Over22	Over 22 Years of age	0.01	0.01	0.39	-0.02	0.04
GPAhi	GPA great than 3.75	0.00	0.02	0.96	-0.03	0.03
Graduate	Graduate Student	0.00	0.02	0.95	-0.03	0.03
μ		0.07	0.00	<0.01	0.07	0.08

Table B2: Estimates of the Utility Function Under RDU

Variable	Description	Estimate	Standard Error	p -value	Lower 95% Confidence Interval	Upper 95% Confidence Interval
A. No Covariates						
\hat{r}		0.49	0.05	<0.01	0.39	0.58
$\hat{\alpha}$		0.07	0.02	<0.01	0.03	0.12
$\hat{\gamma}$		0.72	0.06	<0.01	0.60	0.83
$\hat{\mu}$		0.08	0.00	<0.01	0.07	0.09
B. Adding Covariates						
\hat{r}						
Constant		0.45	0.07	<0.01	0.31	0.58
female	Female	0.08	0.10	0.93	-0.18	0.20
Over22	Over 22 Years of age	0.02	0.15	0.87	-0.28	0.33
GPAhi	GPA great than 3.75	0.07	0.10	0.48	-0.12	0.26
Graduate	Graduate Student	0.18	0.25	0.46	-0.30	0.67
$\hat{\alpha}$						
Constant		0.05	0.02	0.04	0.00	0.09
female	Female	-0.01	0.04	0.82	-0.08	0.06
Over22	Over 22 Years of age	0.03	0.08	0.72	-0.12	0.18
GPAhi	GPA great than 3.75	0.07	0.07	0.30	-0.06	0.21
Graduate	Graduate Student	0.29	0.37	0.42	-0.42	1.01
$\hat{\gamma}$						
Constant		0.72	0.07	<0.01	0.59	0.86
female	Female	0.15	0.10	0.15	-0.05	0.35
Over22	Over 22 Years of age	0.05	0.25	0.85	-0.45	0.54
GPAhi	GPA great than 3.75	-0.09	0.13	0.46	-0.34	0.16
Graduate	Graduate Student	-0.47	0.41	0.25	-1.28	0.33
$\hat{\mu}$		0.07	0.01	<0.01	0.06	0.08

Table B3: Detailed EUT Estimates for Bets on Win Margin Being 5% or More

Variable	Description	Estimate	Standard Error	<i>p</i> -value	Lower 95% Confidence Interval	Upper 95% Confidence Interval
\mathbf{r}						
Constant		0.22	0.07	<0.01	0.07	0.33
female	Female	0.16	0.09	0.05	0.00	0.33
Over22	Over 22 Years of age	0.12	0.12	0.32	-0.12	0.36
GPAhi	GPA great than 3.75	0.06	0.10	0.57	-0.14	0.26
Graduate	Graduate Student	-0.19	0.14	0.16	-0.47	0.07
α						
Constant		0.01	0.01	0.14	0.00	0.03
female	Female	0.02	0.02	0.27	-0.01	0.05
Over22	Over 22 Years of age	0.01	0.02	0.61	-0.03	0.03
GPAhi	GPA great than 3.75	0.00	0.02	0.98	-0.03	0.03
Graduate	Graduate Student	0.00	0.02	0.87	-0.04	0.04
μ_{RA}	Fechner Error	0.09	0.01	<0.01	0.07	0.10
π						
Constant		-0.74	0.28	<0.01	-1.30	-0.19
female	Female	0.32	0.24	0.20	-0.17	0.81
Over22	Over 22 Years of age	-0.17	0.39	0.66	-0.93	0.59
GPAhi	GPA great than 3.75	0.16	0.23	0.48	-0.28	0.60
Graduate	Graduate Student	-0.16	0.48	0.73	-1.12	0.78
LSR	Linear Scoring Rule tasks	0.34	0.27	0.21	-0.19	0.86

Figure B1: Alternative Probability Weighting Functions
with Rank Dependent Utility Model of Choices
over Objective Risk

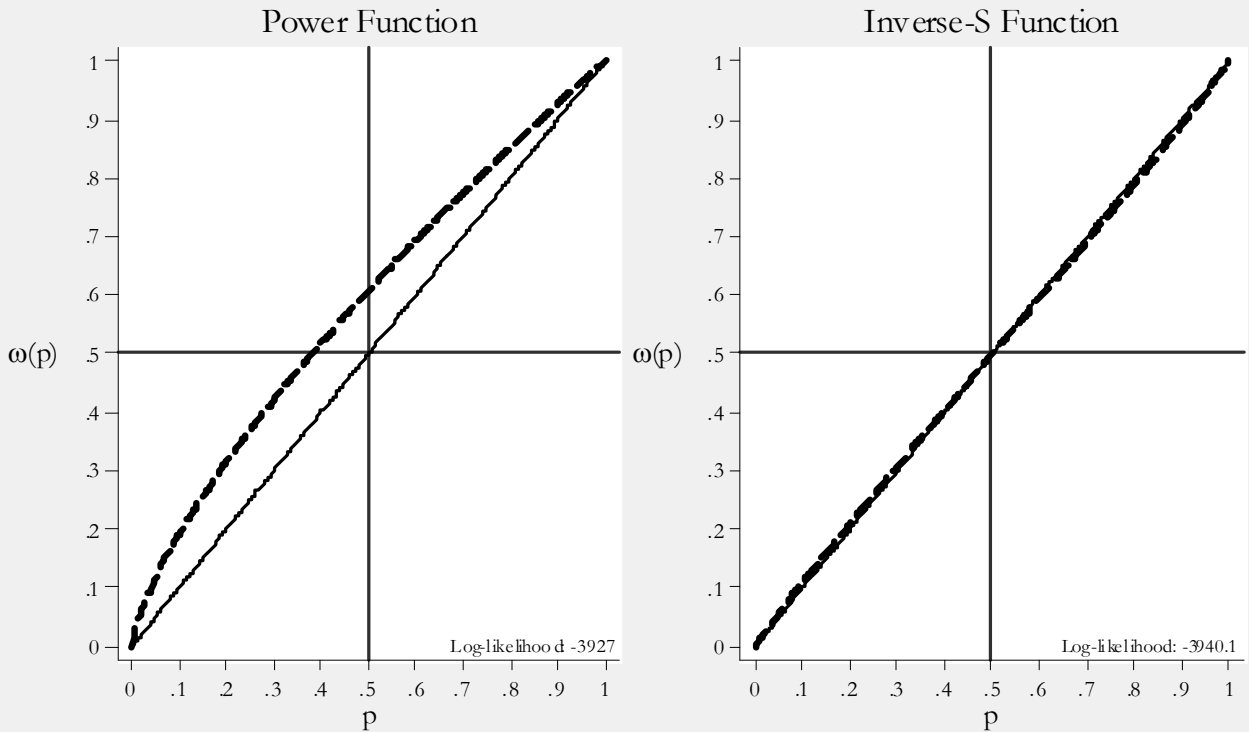



Figure B2: Power-Prelec Probability Weighting Functions
Estimated over Objective and Subjective Risk Tasks



Appendix C: Econometric Specification (NOT FOR PUBLICATION)

A. Risk Attitudes under Expected Utility Theory

We assume an Expo-Power (EP) utility function originally proposed by Saha [1993]. Following Holt and Laury [2002], the EP function is defined as

$$u(y) = [1 - \exp(-\alpha y^{1-r})] / \alpha, \quad (1)$$

where α and r are parameters to be estimated, and y is income from the experimental choice. The EP function can exhibit increasing or decreasing relative risk aversion (RRA), depending on the parameter α : RRA is defined by $r + \alpha(1-r)y^{1-r}$, so RRA varies with income if $\alpha \neq 0$ and the estimate of r defines RRA at a zero income. This function nests CRRA (as $\alpha \rightarrow 0$) and CARA (as $r \rightarrow 0$). The functional form of utility employed here is of no importance, and any monotonic increasing function of $u(\cdot)$ could have been implemented.

The utility function (1) can be estimated using maximum likelihood and a latent EUT structural model of choice. Let there be K possible outcomes in a lottery; in our lottery choice task $K \leq 4$. Under EUT the probabilities for each outcome k in the lottery choice task, p_k , are those that are induced by the experimenter, so expected utility is simply the probability weighted utility of each outcome in each lottery i :

$$EU_i = \sum_{k=1,K} [p_k \times u_k]. \quad (2)$$

The EU for each lottery pair is calculated for a candidate estimate of r and α , and the index

$$\nabla EU = eu_R / (eu_L + eu_R) \quad (3)$$

calculated, where

$$eu_i = \exp(EU_i), \quad (4)$$

for $i = \{R, L\}$, and where EU_L is the “left” lottery and EU_R is the “right” lottery, as displayed to the subject and illustrated in Figure 1. This latent index, based on latent preferences, is already in the form of a probability.²⁷

²⁷ It is well known, but useful to note, that (3) is equivalent to $\Lambda(EU_R - EU_L)$ where $\Lambda(\cdot)$ is the logistic cumulative density function. Thus (3) embodies a statistical “link function” between the difference in the EU of the two lotteries and the probability of the observed choice. Again, the specific “link function” is of no importance here, and we could have implemented any link function mapping the utility difference into a choice probability.

Thus the likelihood of the observed responses, conditional on the EUT and EP specifications being true, depends on the estimates of r and α given the above statistical specification and the observed choices. If we ignore responses that reflect indifference²⁸ the log-likelihood is then

$$\ln L(r, \alpha; y, \mathbf{X}) = \sum_i [(\ln (\nabla EU) \times \mathbf{I}(y_i = 1)) + (\ln (1 - \nabla EU) \times \mathbf{I}(y_i = -1))] \quad (5)$$

where $\mathbf{I}(\cdot)$ is the indicator function, $y_i = 1(-1)$ denotes the choice of the Option R (L) lottery in risk aversion task i , and \mathbf{X} is a vector of individual characteristics reflecting age, sex, race, and so on. When we pool responses over subjects the \mathbf{X} vector will play an important role to allow for some heterogeneity of preferences.

To allow for observed subject heterogeneity with respect to risk attitudes, the parameters r and α are each modeled as linear functions of observed individual characteristics of the subject.²⁹ For example, assume that we only had information on the age and sex of the subject, denoted Age (years of age) and Female (0 for males, and 1 for females). Then we would estimate the coefficients r_0 , r_1 and r_2 in $r = r_0 + r_1 \times \text{Age} + r_2 \times \text{Female}$. Therefore, each subject would have a different estimated r , \hat{r} , for a given set of estimates of r_0 , r_1 and r_2 to the extent that the subject had distinct individual characteristics. So if there were two subjects with the same sex and age, to use the above example, they would literally have the same \hat{r} , but if they differed in sex and/or age they would generally have distinct \hat{r} . In fact, we use four individual characteristics to model heterogeneity in our estimated risk preferences. Apart from a dummy for people aged over 22 years and gender, these include binary indicators for subjects that self-declare as having a high GPA (over 3.75) and self-declare as graduate students.

An important extension of the core model is to allow for subjects to make some errors. The notion of error is one that has already been encountered in the form of the statistical assumption (4) that the probability of choosing a lottery is not 1 when the EU of that lottery exceeds the EU of the other lottery.³⁰ By varying the shape of the link function implicit in (4), one can informally imagine subjects that are more sensitive to a given difference in the index ∇EU and subjects that are not so

²⁸ In our lottery experiments the subjects are told at the outset that any expression of indifference would mean that the experimenter would toss a fair coin to make the decision for them if that choice was selected to be played out. Hence one can modify the likelihood to take these responses into account either by recognizing this is a third option, the compound lottery of the two lotteries, or alternatively that such choices imply a 50:50 mixture of the likelihood of choosing either lottery, as illustrated by Harrison and Rutström [2008; p.71]. We do not consider indifference here because it was an extremely rare event.

²⁹ Alternatively we could allow for unobserved individual heterogeneity in our maximum likelihood estimation by specifying fixed effects, random effects or random coefficients, and constrain the variance matrix accordingly. Andersen, Fountain, Harrison, Hole and Rutström [2012] illustrate how to apply the random coefficients approach to estimating “imprecise” subjective probabilities.

³⁰ This assumption is clear in the use of a link function from the difference between the EU of each option to the probability of picking one or other lottery; in the case of the logistic CDF that is implied by our approach in (3), this link function is $\Lambda(EU_R - EU_L)$. If the subject exhibited no errors from the perspective of EUT, this link function would instead be a step function: zero for all values of $(EU_R - EU_L) < 0$, anywhere between 0 and 1 for $(EU_R - EU_L) = 0$, and 1 for all values of $(EU_R - EU_L) > 0$. Harrison [2008; p.326] illustrates the implied CDF, referring to it as the CDF of a “Hardnose Theorist.”

sensitive. We use the contextual error specification proposed by Wilcox [2011]. It posits the latent index

$$eu_i = \exp[(EU_i/v)/\mu], \quad (4')$$

instead of (4), where v is a normalizing term for each lottery pair L and R , and $\mu > 0$ is a structural “noise parameter” used to allow some errors from the perspective of the deterministic EUT model. The normalizing term v is defined as the maximum utility over all prizes in this lottery pair minus the minimum utility over all prizes in this lottery pair, and ensures that the *normalized* EU difference $[(EU_R - EU_L)/v]$ remains in the unit interval. As $\mu \rightarrow \infty$ this specification collapses ∇EU to 0 for any values of EU_R and EU_L , so the probability of either choice converges to $1/2$. So a larger μ means that the difference in the EU of the two lotteries, conditional on the estimate of r , has less predictive effect on choices. Thus μ can be viewed as a parameter that flattens out, or “sharpens,” the link functions implicit in (4). This is just one of several different types of error story that could be used, and Wilcox [2008] provides a masterful review of the implications of the strengths and weaknesses of the major alternatives.

Thus we extend the likelihood specification to include the noise parameter μ and maximize $\ln L(r, \alpha, \mu; y, \mathbf{X})$ by estimating r , α and μ , given observations on y and \mathbf{X} .³¹ Additional details of the estimation methods used, including corrections for “clustered” errors when we pool choices over subjects and tasks, is provided by Harrison and Rutström [2008; p.69ff].

B. Risk Attitudes under Rank Dependent Utility Theory

The RDU model extends the EUT model by allowing for decision weights on lottery outcomes. Instead of (1) we have

$$u(y) = [1 - \exp(-\acute{\alpha}y^{1-\acute{r}})]/\acute{\alpha}, \quad (6)$$

where \acute{r} and $\acute{\alpha}$ are the coefficient for the curvature of the utility function, comparable to r and α in (1). To calculate decision weights under RDU one replaces expected utility defined by (2) with RDU

$$RDU_i = \sum_{k=1, K} [w_k \times u_k] \quad (7)$$

where

$$w_i = \omega(p_i + \dots + p_n) - \omega(p_{i+1} + \dots + p_n) \quad (8)$$

for $i=1, \dots, n-1$, and

$$w_n = \omega(p_n) \quad (9)$$

³¹ The normalizing term v is given by the value of r and the lottery parameters, which are part of \mathbf{X} .

for $i=n$, the subscript indicating outcomes ranked from worst to best, and where $\omega(p)$ is some probability weighting function. In general, we adopt the simple “power” probability weighting function proposed by Quiggin [1982], with curvature parameter γ :

$$\omega(p) = p^\gamma \tag{10}$$

So $\gamma \neq 1$ is consistent with a deviation from the conventional EUT representation.

An alternative probability weighting function is the “inverse-S” function popularized by Tversky and Kahneman [1992]:

$$\omega(p) = p^\gamma / (p^\gamma + (1-p)^\gamma)^{1/\gamma} \tag{11}$$

This function exhibits inverse-S probability weighting (optimism for small p , and pessimism for large p) for $\gamma < 1$, and S-shaped probability weighting (pessimism for small p , and optimism for large p) for $\gamma > 1$.

The final probability weighting function is a general functional form proposed by Prelec [1998] that exhibits considerable flexibility, nesting the first two specifications (10) and (11). This function is

$$\omega(p) = \exp\{-\eta(-\ln p)^\varphi\}, \tag{12}$$

and is defined for $0 < p \leq 1$, $\eta > 0$ and $0 < \varphi < 1$. When $\varphi = 1$ this function collapses to the Power function $\omega(p) = p^\eta$. Andersen, Harrison, Lau and Rutström [2011] point out that this function is not well-behaved numerically, in the sense of having reliable numerical derivatives, when $\varphi \rightarrow 1$.³² Indeed, when $\varphi = 1$ this function collapses to the Power function

$$w(p) = p^\eta \tag{10'}$$

There are several solutions to this problem of numerical instability, and the simplest is to estimate a generalization that nests (12) inside a power function with parameter γ :

$$w(p) = [\exp\{-\eta(-\ln p)^\varphi\}]^\gamma, \tag{13}$$

so that we get the numerical stability of (10) when (12) collapses because $\varphi \rightarrow 1$, and yet we do not rule out the flexibility of (12) when appropriate.³³ We call this the Power-Prelec probability

³² The numerical problems are easy to see. To ensure that $0 < \varphi < 1$, we estimate the parameter φ and then define φ by $1/(1+\exp(\varphi))$ within the maximum likelihood evaluator. Hence the algorithm searching for maximum likelihood parameter values can vary φ between $\pm\infty$ and ensure that the constraint on φ is always satisfied. However, if the maximum likelihood value of φ is 1, then arbitrarily large negative values of φ will generate numerical values of φ close to 1. Thus large swings in φ can occur with the same likelihood value, disrupting the gradient methods used to locate a maximum.

³³ The other numerical solution is to generate “profile likelihoods” in which φ is set parametrically to 1, 0.99, 0.98... 0.9, and evaluate the likelihood conditional on that assumed value and the maximum likelihood

weighting function. We reference the Prelec function in the main text, because the Power-Prelec is just an extension to achieve the same specification in a more robust numerical manner.

We explain in the text that the Power function has a better likelihood than the Inverse-S function, and that the Prelec function essentially validates the use of the Power. Figure B1 in Appendix B shows that the estimate of γ for the Power function is 0.71 when applied solely to the risk choices over objective probabilities, with a log-likelihood value of -3927. The comparable estimate of γ for the Inverse-S function is 0.95, with a worse log-likelihood value of -3940. These qualitative results persist when one jointly estimates objective and subjective probabilities. Figure B3 considers the flexible Prelec specification³⁴ when we undertake that joint estimation, showing that the preferred specification is clearly “power-like.”

The same behavioral noise specification is used as in the EUT model, replacing EU_i with RDU_i , leading to the definition of a comparable latent index ∇RDU . The likelihood specification for the RDU model is therefore

$$\ln L(\acute{r}, \acute{\alpha}, \gamma, \mu; y, \mathbf{X}) = \sum_i [(\ln (\nabla RDU) \times \mathbf{I}(y_i=1)) + (\ln (1 - \nabla RDU) \times \mathbf{I}(y_i = -1))] \quad (11)$$

and entails the estimate of \acute{r} , $\acute{\alpha}$, γ and μ using maximum likelihood. Individual heterogeneity is allowed for by estimating the parameters \acute{r} , $\acute{\alpha}$ and γ as linear functions of the demographic characteristics defined earlier.

Additional References

- Andersen, S., Harrison, G. W., Lau, M. I., and Rutström, E. E. (2011). Discounting Behavior: A Reconsideration. *Working Paper 2011-03*, Center for the Economic Analysis of Risk, Robinson College of Business, Georgia State University.
- Andersen, S., Fountain, J., Harrison, G. W., Hole, A. R., & Rutström, E. E. (2012). Inferring Beliefs as Subjectively Imprecise Probabilities. *Theory and Decision*, 73, 161-184.
- Harrison, G. W. (2008). Neuroeconomics: A Critical Reconsideration. *Economics & Philosophy*, 24, 303-344.
- Wilcox, N. T. (2008). Stochastic Models for Binary Discrete Choice Under Risk: A Critical Primer and Econometric Comparison. In J. Cox and G.W. Harrison (eds.), *Risk Aversion in Experiments*. Bingley: Emerald Publishing Limited.
- Wilcox, N. T. (2011). ‘Stochastically More Risk Averse?’ A Contextual Theory of Stochastic Discrete Choice Under Risk. *Journal of Econometrics*, 162, 89–104.

values of other parameters. Inspection of these likelihoods can then verify that the maximum likelihood value of φ is indeed 1.

³⁴ For reasons just explained, we used the numerically robust Power-Prelec function in all estimations, but refer to it as the Prelec function.