Estimating Aversion to Uncertainty

by

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

May 2009

Working Paper 09-07, Department of Economics,
College of Business Administration, University of Central Florida, 2009

ABSTRACT. It is intuitive that decision-makers might have attitudes towards uncertainty just as they might have attitudes towards risk. However, it is only recently that this intuitive notion has been formalized and axiomatically characterized. We estimate the extent of uncertainty aversion in a manner that is parsimonious and consistent with theory. We demonstrate that one can *jointly* estimate attitudes towards uncertainty, attitudes towards risk, and subjective probabilities in a rigorous manner. Our structural econometric model constructively demonstrates the theoretical claims that it is possible to define uncertainty aversion in an empirically tractable manner. Our results show that attitudes towards risk and uncertainty can be different, qualitatively and quantitatively, and that allowing for these differences can have significant effects on inferences about subjective probabilities.

[†] Department of Economics and Centre for Economic and Business Research, Copenhagen Business School, Copenhagen, Denmark (Andersen); Department of Economics, University of Canterbury, Christchurch, New Zealand (Fountain); and Department of Economics, College of Business Administration, University of Central Florida, USA (Harrison and Rutström). E-mail: sa.eco@cbs.dk, john.fountain@canterbury.ac.nz, gharrison@research.bus.ucf.edu and erutstrom@bus.ucf.edu. We thank the U.S. National Science Foundation for research support under grants NSF/HSD 0527675 and NSF/SES 0616746. We are grateful to David Rivenbark, Peter Wakker, university seminar participants at Brescia, George Mason, LUISS and Pittsburgh, and participants at the 2009 *Decisions – Theory, Experiments, and Applications* conference in Paris for helpful comments.

It is intuitive that decision-makers might have attitudes towards uncertainty just as they might have attitudes towards risk. However, it is only recently that this intuitive notion has been formalized and axiomatically characterized. We estimate the extent of uncertainty aversion in a manner that is parsimonious and consistent with theory. Although some hypothetical thought experiments, due to Ellsberg [1961], were the ultimate cause of the development of these theoretical models, many of the previous experimental studies were focused on demonstrating the *qualitative existence* of uncertainty (or ambiguity) aversion. Our focus is on *quantitatively estimating* uncertainty aversion.

There are several reasons to be interested in the measurement of uncertainty aversion. One is to see if it is quantitatively significant, even if it exists qualitatively. Another is to see how large it is, and whether it is something that can be safely assumed away. We present evidence that it is quantitatively significant, and cannot be safely assumed away.

Our approach is deliberately parsimonious, in the sense of operationalizing, in experiments and econometric inference, the simplest possible models that can rigorously account for uncertainty aversion. We discuss the relationship between the approach adopted here and the most general theoretical frameworks.

In section 1 we review theoretical specifications that allow a role for uncertainty aversion, and which provide operational definitions of the concept. We discuss the issues that arise in translating these theoretical specifications into empirical measuring instruments, and the tradeoffs involved in using the alternative specifications. In section 2 we present an experimental design that allows us to estimate uncertainty aversion in a manner that is directly consistent with the theoretical specification we employ. We elicit risk attitudes with a standard lottery choice task, and elicit beliefs and uncertainty aversion defined over subjective events using scoring rules. In section 3 we formally

state the estimation problem, defining the likelihood function to be maximized. A key feature of this problem, and one that derives intimately from the theoretical specification, is the need to jointly estimate risk attitudes, the subjective probability of some uncertain event, and attitudes towards uncertainty. These three are intrinsically linked conceptually, and this is reflected in the theoretical and empirical approach employed. Section 4 presents the results of applying our approach. We find striking evidence that attitudes towards risk can be different to attitudes to uncertainty. We find evidence of quantitatively significant differences in the two, further implying significant differences in inferences about latent subjective probabilities. An appendix discusses related empirical literature in detail.

1. Uncertainty Aversion

Consider the canonical Ellsberg [1961; p. 650] example in which there are two urns, one with known probabilities and one with unknown probabilities. Each urn contains some mix of red and black balls, possibly degenerate. There are 100 balls in each, and in one urn the subject knows that there are 50 red and 50 black balls. Call the realization from drawing one ball from this urn R or B. The other urn has some mix of red and black balls, but the mix is unknown. It might be 50-50, it might be 100-0, or it might be 3-97. Call the realization from drawing one ball from this urn r or b.

Define four lotteries over the possible draws from these two urns: x, y, x′ and y′. Lottery x pays ($100, $100, $0, $0) for realizations (Rr, Rb, Br, Bb), so it pays $100 depending on the realization of R or B from the urn with known probabilities, irrespective of the realization of r or b from the urn with unknown probabilities. Similarly define the lottery y as ($100, $0, $100, $0), so this lottery pays off depending on the realization from the urn with unknown probabilities. Then define lottery x′ as ($0, $0, $100, $100) and lottery y′ as ($0, $100, $0, $100).

Let us assume the standard behavioral pattern assumed here in discussions of this example, ignoring the problems of operationalizing the uncertain urn in any real experiment. Specifically, subjects are presumed to exhibit a strong preference for x over y, a strong preference for x′ over y′, and to be indifferent between x and x′ and between y and y′. In effect, they prefer to bet over lotteries defined on known probabilities, and appear to exhibit an aversion to uncertainty.

There are two models that minimally extend subjective expected utility (SEU) theory to account for this apparent aversion to uncertainty. One is called the Source-Dependent Risk Attitude (SDRA) model and the other is called the Uncertain Priors (UP) model, following the taxonomy of Nau [2007], although names vary a great deal in the literature. The SDRA is formally a special case of UP model, but provides a parsimonious specification for estimation purposes and is worth writing out separately. We present each model, and show how it can rationalize the stylized pattern of choices for this Ellsberg example using the notion of uncertainty aversion. Connections to the theoretical literature are discussed subsequently, because there have been several indirect paths to these two formalizations.

Assume that one can *a priori* identify two stochastic processes that are more or less certain than the other, or at least "different" for now. In the Ellsberg example this is obvious: one urn has outcomes generated by known probabilities, and the other does not. Assume for the moment that we know *a priori* which is which. Adapting slightly the notation in Nau [2007] let there be two logically independent sets of events $(U_{i1},..., U_{iJ})$ and $(C_{i1},..., C_{iK})$ for each type of process, where the U events are for the uncertain process with unknown probabilities over J possible events, the C events are for the certain process with known probabilities over K possible events, and i denotes one of I prior subjective probability distributions that the decision-maker might hold. The probabilities over these I priors can be denoted $\boldsymbol{\rho} = (\rho_1,..., \rho_I)$. To connect to the central metaphor from the literature,

think of the U events as the horse races of Anscombe and Aumann [1963] and the C events as their roulette wheels.[1]

The decision-maker has probabilistic beliefs over each of these events. For the U events and prior i, denote these $\pi^U = (\pi^U_{i1},..., \pi^U_{ij})$. For the C events and prior i, think of these as conditional on a realization from the U process: conditional on some event U_{ij}, we assume probabilities $\pi^C = (\pi^C_1,..., \pi^C_K)$ for the K events in process C. To pursue the horse race metaphor, the decision-maker might have one prior over the performance of horses if it rains and the track is heavy, and another prior over the performance of the horses if the track is dry. This is quite natural: you might be betting over the subjectively uncertain performance of a horse in U, but your payoffs are specific amounts of money realized in process C with known probabilities (possibly degenerate), and your priors depend on the expected state of the world (or weather). For our purposes we will assume that the final outcomes are all amounts of money.

The theoretical literature provides a rich array of alternative axiomatizations of these representations. The general insight is to modify the standard axioms of Savage [1972] so that the horse lotteries and roulette wheel lotteries of Anscombe and Aumann [1963] are treated differently by the decision-maker, and not reduced to one compound lottery.[2] The UP model was independently axiomatized by Neilsen [1993] and Klibanoff, Marinacci and Mukerji [2005], and discussed explicitly by Nau [2001] as a generalization of the SDRA model.

---

[1] According to that metaphor, individuals place bets on horses that are based on their subjective probabilities of each horse winning, and are paid off in lotteries defined over realizations from a roulette wheel with objective probabilities of generating a finite set of terminal outcomes. The roulette wheel might be degenerate, in the sense of having one terminal outcome for each horse.

[2] Savage [1972] defined preferences over "acts," which are lotteries with subjective probabilities defined over final payoffs. Thus the object of preferences in his axiom set collapses the horse lottery and roulette wheel compound lottery. Smith [1969] and Gärdenfors and Sahlin [1982][1983] were the first to propose the use of compound lotteries to understand behavior in the Ellsberg task. Segal [1987] was the first to prove how a compound lottery representation and the relaxation of the reduction axiom could account for the Ellsberg behavior, resulting in the first "recursive" representation.

Let there be I priors, where each prior is a set of probabilities over the U-events. Let $(\sigma_{ij1},...,$ $\sigma_{ijK})$ denote the subjective conditional probability of events U_{ij} *and* C_k occurring under prior i, so σ_{ijk} $\equiv \pi^U_{ij}\pi^C_k$. The overall evaluation of an act z is then accomplished by evaluating over all I priors in the natural manner:

$$W(z) = \sum_{i=1,I} \rho_i \, u[\sum_{j=1,J} \sum_{k=1,K} \pi^U_{ij}\pi^C_k \, v(z_{jk})] = \sum_{i=1,I} \rho_i \, u[\sum_{j=1,J} \sum_{k=1,K} \sigma_{ijk} \, v(z_{jk})] = \sum_{i=1,I} \rho_i \, u[\Lambda_{ijk}] \quad (1)$$

where Λ_{ijk} is the SEU of the final outcomes if events j and k occur and the decision-maker uses prior i. Thus we have a prior-probability weighted average of the "SEU evaluations of SEU."

To take a specific parametric example, assume simple power specifications

$$v(z) = z^\alpha \quad (2)$$

and

$$u(z) = z^\beta \quad (3)$$

for some argument z. The argument of $u(\cdot)$ in (1) is actually the SEU defined over final monetary outcomes, Λ_{ijk}, so there is a natural sense in which this is a recursive SEU model.[3] Thus α<1 (>1) implies risk aversion (loving), β<1 (>1) implies uncertainty aversion (loving), α=1 implies risk neutrality, and β=1 implies uncertainty neutrality. When there is uncertainty neutrality (1) behaves just like a conventional SEU characterization. Assume α=0.5, consistent with modest risk aversion, and β=0.9, consistent with very slight uncertainty aversion. Call β the Constant Relative Uncertainty Aversion (CRUA) coefficient, in parallel with α as the Constant Relative Risk Aversion (CRRA) coefficient.

Assume I=3, and let the three priors over r and b be (0.5, 0.5), (0.4, 0.6) and (0.6, 0.4). Thus the first prior treats the uncertain urn as being a fair urn, and the other two priors assume some

---

[3] The final SEU is defined with respect to objective probabilities, so we could drop the "subjective" in SEU and refer to (1) as an SEU evaluation of an EU. Some of the theoretical literature substitutes the certainty-equivalent of the lottery for the EU, perhaps to stress the similarity between the arguments of (2) and (3). If the individual is not satiated, then the CE is an order-preserving transform of the EU, and nothing is formally lost with this switch, but it can be confusing.

small bias in favor of one of the colors. The probabilities on these three priors are assumed to be $\rho_1=0.6$, $\rho_2=0.2$ and $\rho_3=0.2$. We evaluate the conditional probabilities σ_{ijk} as (0.25, 0.25, 0.25, 0.25) for $(\sigma_{1Rr}, \sigma_{1Rb}, \sigma_{1Br}, \sigma_{1Bb})$, as (0.2, 0.3, 0.2, 0.3) for $(\sigma_{2Rr}, \sigma_{2Rb}, \sigma_{2Br}, \sigma_{2Bb})$, and as (0.3, 0.2, 0.3, 0.2) for $(\sigma_{3Rr}, \sigma_{3Rb}, \sigma_{3Br}, \sigma_{3Bb})$. For each lottery and prior we calculate the vector of interim SEU evaluations, $\Lambda_{ijk} = \sum_{j=1,J} \sum_{k=1,K} \sigma_{ijk} v(z_{jk})$, using the probabilities σ_{ijk} under each prior i as weights and applying the risk aversion function $v(\cdot)$ from (2). Spelling this out, for the first prior i=1, and for lottery x, we have $\Lambda_{1jk}=\sum_{j=1,J} \sum_{k=1,K} \sigma_{1jk} v(z_{jk}) = 0.25\times100^\alpha+0.25\times100^\alpha+0.25\times0^\alpha+0.25\times0^\alpha = 5$. For the two other priors for lottery x we calculate $\Lambda_{2jk}= 5$ and $\Lambda_{3jk}=5$. For each lottery we can thus calculate the vector of interim SEU values as $\Lambda_{ijk} = (5, 5, 5)$ for lottery x, $\Lambda_{ijk} = (5, 4, 6)$ for lottery y, $\Lambda_{ijk} = (5, 5, 5)$ for lottery x′, and $\Lambda_{ijk} = (5, 6, 4)$ for lottery y′. Finally, we take the weighted average of these SEU evaluations, using the probabilities ρ_i over each prior i as a weight and applying the uncertainty aversion function $u(\cdot)$ from (3) to the interim SEU. So we have

$$W(x) = 0.6 \times 5^\beta + 0.2 \times 5^\beta + 0.2 \times 5^\beta = 4.26 \tag{4}$$
$$W(x′) = 0.6 \times 5^\beta + 0.2 \times 5^\beta + 0.2 \times 5^\beta = 4.26 \tag{4′}$$
$$W(y) = 0.6 \times 5^\beta + 0.2 \times 4^\beta + 0.2 \times 6^\beta = 4.25 \tag{5}$$
$$W(y′) = 0.6 \times 5^\beta + 0.2 \times 6^\beta + 0.2 \times 4^\beta = 4.25 \tag{5′}$$

These results qualitatively explain the stylized Ellsberg pattern of choices, since W(x)>W(y), W(x′)>W(y′), W(x)=W(x′) and W(y)=W(y′). When the decision-maker is uncertainty-neutral, and $\beta=1$, these evaluations collapse to the SEU representation, and W(x) = W(y) = W(x′) = W(y′).

The historically important "maxmin multiple priors" model of Gilboa and Schmeidler [1989] can be viewed as a special case of this model in which the prior probabilities $(\rho_1,...,\rho_I)$ are uniform and the uncertainty aversion function (3) exhibits extreme concavity. In one sense the maxmin multiple priors model conflates *perceptions of* uncertainty and *attitudes towards* uncertainty, whereas we want to be able to identify each.

Two special cases of the UP model are worth noting. One arises when I=1, and there is

simply one prior. Thus we can remove the index i and (1) becomes

$$W(z) = u[\ \textstyle\sum_{j=1,J} \sum_{k=1,K}\ \pi^{U}_{j}\pi^{C}_{k}\ v(z_{jk})\ ] \tag{6}$$

In some respects we can think of this as the "natural special case" of the UP model, since it smoothly collapses the subjective belief *distribution* down to a subjective *probability*. It does not thereby collapse further to a Savage-consistent SEU representation, which is

$$W(z) = \textstyle\sum_{j=1,J} \sum_{k=1,K}\ \pi^{U}_{j}\pi^{C}_{k}\ v(z_{jk}) \tag{7}$$

Of course, (7) is the same as (6) when the u(·) function exhibits neutrality towards uncertainty.

The other special case of the UP model is the SDRA model. To see how this formally emerges as a special case, first set I=J, so that we can remove the index j and only have one value of π^{U}_{i}, so that (1) becomes

$$W(z) = \textstyle\sum_{i=1,I}\ \rho_{i}\ u[\ \sum_{k=1,K}\ \pi^{U}_{i}\ \pi^{C}_{k}\ v(z_{jk})\ ] \tag{8}$$

Then set $\pi^{U}_{i} = \{1, 0\}$ for each i, and we get

$$W(z) = \textstyle\sum_{i=1,I}\ \rho_{i}\ u[\ \sum_{k=1,K}\ \pi^{C}_{k}\ v(z_{jk})\ ] \tag{9}$$

In effect, we have assumed that the only class of priors that are held in the Ellsberg example are those in which there are only "all red" or "all black" balls in the uncertain urn. This follows from assuming π^{U}_{ij} is either 0 or 1, and implicitly that J=1. Hence there are only two priors and I=2, so ρ_{1} is the probability that the uncertain urn is full of red balls such that the draw will always be r, in which case ρ_{2} would be the probability that the uncertain urn is full of black balls such that the draw will always be b. Although (9) is formally a special case of (1), it is not the same special case as (6), nor, of course, is it the same as the SEU representation (7). The SDRA model is formally developed by Nau [2006; Model II, Theorem 2, p. 143] and Ergin and Gul [2009; Theorem 3] from slightly different perspectives and axioms.

C. Estimation

In general terms, we can see how the SDRA and general UP models differ in terms of their

implications for estimation from (ideal) experimental data. For both models we need to estimate the

utility function $v(\cdot)$ defined over final outcomes, and this is a relatively simple matter since we can

confront the subject with an array of lotteries defined over objective probabilities, which are the C-

events of these models. So we assume throughout that $v(\cdot)$ is estimable using familiar methods

surveyed in Harrison and Rutström [2008].[4]

The SDRA model is the most parsimonious to estimate. It requires that we jointly estimate

the subjective probabilities ρ and the utility function $u(\cdot)$, in addition to $v(\cdot)$.[5] In most of the

applications of interest, and certainly in the experiments we design, we can reduce the number of

events to two, so there is only one parameter to estimate in order to recover the subjective belief

distribution. Note that we are estimating the point estimate of the prior probability ρ_1 from (9),

consistent with the parsimonious assumption that there are only two subjective probability

distributions over the uncertain events, reflecting the extreme, degenerate underlying priors of each

of the two possible outcomes. So if we know ρ_1 we know $\rho_2=1-\rho_1$, and we can recover the

underlying subjective belief distribution as a 2-point discrete distribution with mass ρ_1 at 1 and mass

---

[4] If we were willing to restrict the specification of uncertainty aversion to models that reflect violations of the Reduction of Compound Lotteries axiom, then one could contemplate an experimental design in which choices over simple lotteries are used to identify $v(\cdot)$ and choices over compound lotteries are used to identify $u(\cdot)$ conditional on the estimates of $v(\cdot)$. This would "free up" the belief elicitation choices to identify the subjective probability distribution, conditional on the estimates of $v(\cdot)$ and $u(\cdot)$.

[5] The need to have an experimental procedure that "pins down" the utility function defined over final outcomes, in order to identify the SDRA (or UP) model, is also recognized clearly by Hey, Lolito and Maffioletti [2008]. Explaining why they did not consider recursive EU models, they note (p.24): "... if one follows the route of these two-stage probability distributions, we would then have to estimate the first-stage probability distribution (in the absence of an independent procedure to elicit them). Without any restriction on the form of these first-stage distributions, we have a serious problem with degrees of freedom." We agree completely, which is why we deliberately implemented an experiment in which we did have "an independent procedure" to elicit the parameters of the first stage of the recursive EU models. Of course, identification issues remain with the UP models, as we discuss below.

ρ_2 at 0. Figure 1 illustrates this representation for the case in which $\rho_1=\rho_2=\frac{1}{2}$. To provide the obvious contrast, Figure 2 illustrates in the same space what the traditional SEU model assumes: that all of the mass is concentrated at one point, in this case a subjective probability of 0.71.[6]

In the *general* UP model we would have to estimate much more. If we allow arbitrary distributions of priors and subjective beliefs, this would be a daunting estimation task. For the Ellsberg example we considered, we would have to estimate two of ρ_1, ρ_2 and ρ_3. In addition to the probability of each prior, we would have to estimate the subjective probability of the uncertain event π^U_{ir} for each possible prior i. So there would be five parameters to estimate, just for this simple example.

However, if we make some parametric assumption about the form of the distribution of the priors ρ and the uncertain process π^U, the difficulty of estimating the UP model can be reduced considerably. There are two parsimonious approaches one could follow here.

One approach is to assume some tractable continuous density function for the subjective beliefs that is defined by a (very) finite number of parameters. For example, if we assume that the subjective belief is characterized by some Normal distribution, truncated at 0 or 1, then we would get the distribution of the priors ρ and the uncertain process π^U when we estimate the (hyper-parameters of the) Normal distribution.[7] Hence the number of parameters to estimate would drop down to two, the mean and standard deviation of the Normal distribution. Figure 3 illustrates this type of representation, with a distribution that has mean of 0.6. This approach remains a challenge for estimation, demanding the use of maximum simulated likelihood methods and raising some new

---

[6] There is no sense in which the quantitative magnitudes in Figures 1 and 2 (or Figures 3 and 4) are meant to be alternative representations of the same specific subjective beliefs. These are intended purely to illustrate qualitatively the type of representation that each model assumes.

[7] One could also replace the assumption of a Normal distribution with a two-parameter Beta distribution, which is naturally constrained to have a domain between 0 and 1.

issues, but is feasible; we explore this estimation approach in Andersen, Fountain, Harrison, Hole and Rutström [2009].

The other parsimonious approach is to assume that the prior distribution ρ is a κ-point, discrete, uniform distribution, and that each π^{U}_{ij} is a degenerate distribution with mass at just one value. We could then estimate the κ subjective probability values. When $\kappa=1$ we have the natural special case (6) referred to above. When $\kappa=2$ *and* we further constrain $\pi^{U}_{1j}=1$, $\pi^{U}_{2j}=0$ and J=1, we have the SDRA special case (9). Indeed, for some larger κ one could generate arbitrarily good discrete approximations to continuous density functions such as the one shown in Figure 3 (e.g., $\kappa=10$ would allow us to estimate deciles of the underlying distribution, which might be adequate for some inferential purposes). This discrete approximation approach is particularly attractive when there is reason to expect a latent density function for subjective beliefs that is not well-characterized by a finite number of parameters, such as multi-modal distributions.

For $\kappa>1$ we have to make some further parametric assumptions about the form of the priors and the distribution of subjective beliefs about the uncertain process. For example, for $\kappa=2$ we could assume that ½ of the subjective distribution is attached to each distinct estimated probability, so $\rho_1 = \rho_2 = $ ½ by assumption, and estimate two subjective probability values π^{U}_{1j} and π^{U}_{2j}.[8] This approach only involves the estimation of κ parameters.[9] Figure 4 illustrates, for the case in which

---

[8] This representation has a formal similarity to the definition of a "subjective mixture" of acts provided by Ghirardoto, Maccheroni, Marinacci and Siniscalchi [2003], which they use to axiomatize many popular models of ambiguity aversion.

[9] Although this approach would result in estimates that are consistent with the subject behaving in accordance with the UP model, alternative interpretations are not ruled out. Any differences in the subjective probabilities π^{U}_{1j} and π^{U}_{2j} might be due to the heterogeneity of subjective probabilities that are each held with certainty by different individuals. For example, if we have equal samples of men and women, men might believe one thing with certainty and women might believe another thing with certainty – not a rare phenomenon, we have been led to believe. Of course, one solution in this instance is to model the heterogeneity by including a dummy for sex. But then the problem remains in the form of unobserved individual heterogeneity. It is important to recall that we are making statements about the sample, not about any individuals, and that these alternative interpretations are not ruled out.

π^{U}_{1j}=0.42 and π^{U}_{2j}=1. Alternatively, one could assume κ specific subjective probability values π^{U}_{ij}, and estimate the prior probability ρ_i for each. To use the above example, we might fix the subjective probability values for r at 0.5, 0.4 and 0.6, and then estimate the prior probabilities as 0.6, 0.2 and 0.2.

The upshot is that one can identify uncertainty aversion in a manner that is consistent with the UP model, as long as some parametric assumptions are made and alternative interpretations are acknowledged. For present purposes, we make the assumption of a uniform κ-point prior probability $\rho_i = 1/\kappa$ and of homogeneity of subjective beliefs within our sample. We appreciate that these assumptions are strong, but some such assumptions are needed in order to operationalize the UP model.[10] These assumptions do allow one to evaluate the extent to which the estimates of uncertainty aversion appear to be sensitive to allowing this additional degree of uncertainty.

2. Experimental Design

We recruited 140 subjects from the student population of the University of Central Florida in October 2008 to participate in these experiments. Complete instructions are provided in an

---

[10] Experimental procedures that are designed to generate estimates of theories of decision-making under uncertainty at the level of the individual are developed by Abdellaoui, Baillon, Placido and Wakker [2008], Ahn, Choi, Gale and Kariv [2009] and Hey, Lotito and Maffioletti [2007][2008]. We discuss their approaches in more detail in Appendix B, although we note the cautionary comment by Hey et al. [2007; p.9] that "... one could argue that many of these second-order models simply are not identifiable." We certainly agree with this viewpoint if one is not willing to make some strong identifying assumptions, such as the ones we suggest. It is also arguable that many of the alternative models of ambiguity aversion are conceptually indistinguishable without further identifying structure. Consider, for example, the Maxmin Expected Utility (MEU) model of Gilboa and Schmeider [1989] and the α-MEU representation due to Ghirardoto, Maccheroni and Marinacci [2004]. Siniscalchi [2006] shows that "... a preference that satisfies the Gilboa-Schmeider [1989] axioms admits a MEU representation; however, ..., the *same* preference typically admits an α-MEU representation, and the set of priors appearing in the two representations are different. Thus, additional considerations must be invoked in order to determine which of these sets, if any, comprises all possible probabilistic descriptions of the uncertainty, and hence which decision criterion reflects the decision-maker's attitudes towards ambiguity." (p. 92).

appendix. The basic design objective was to have some choice tasks that allow us to estimate and identify attitudes towards objective risk, and another choice task for the same subjects that allows us to estimate subjective beliefs and attitudes towards uncertainty. Thus we exploit the ability in an experiment to bundle several tasks together to allow estimation of a structural model in which all parameters of theoretical interest are identified (e.g., the joint estimation of risk attitudes and discount rates of Andersen, Harrison, Lau and Rutström [2008]).

Figure 5 illustrates the lottery choice that subjects were given. Each subject faced 45 such choices, where prizes spanned the domain $0 up to $100. 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 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 of one of the levels. 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.

We employed two popular types of scoring rules on a between-subjects basis. One is the Quadratic Scoring Rule (QSR) and the other is the Linear Scoring Rule (LSR). Figure 6 shows the interface for the Quadratic Scoring Rule (QSR) as it was presented to subjects on a computer screen and in printed instructions. The interface was explained with these instructions:

> 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.

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.

The subject was then taken through displays of their payoffs if they chose to report 0% or 100%.

The LSR used the same instructions, except for references to payoffs for interior reports. For

example, the payoffs shown at the end of the above instructions were $710 and $290 for the LSR.

We then concluded the main instructions in this manner:

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.

Each subject participated in a training choice, in which they were told the number of orange balls in

the a 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 that this was to be a trainer.

Each subject participated in several 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: subjects were told that there were 60 balls in all 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. We do not consider these events here. The belief elicitation task we examine is based on the 2008 U.S. Presidential Election, which was to be held about one week after the session. We elicited beliefs that the winning share of the popular vote would be 5 percentage points or more greater than the losing share. Our own *a priori* expectations for the subjective probability of this event was around 65%.

The exact phrasing of these events was explained in written instructions, which were also read out loud, and are available in an appendix. Subjects were asked to place bets on the question: "Will the popular vote for the winning candidate be 5 or more percentage points greater than the popular vote for the losing candidate?" For the benefit of subjects that might not know the difference, it was explained that the question was about the popular vote and not the outcome of the Electoral College vote, that the popular vote is just the sum of all votes across the United States, and that we were only referring to the Presidential Election, and not to any other elections that might occur on the same day. We also clarified what we meant by a percentage point: "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." Finally, we told subjects that the final outcome, and hence payoffs, would be determined by a widely respected public source for the outcome, the *New York Times* of Friday, November 7, 2008.[11] Finally, payments were to be made immediately after the election, and subjects were told that they "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."

The experiments were conducted between Monday October 27 and Friday October 31, in the week prior to the election. Including other belief elicitation treatments, a total of 354 subjects participated, earning a total of $32,101 for an average of just over $90 per subject. Each session lasted around 2 hours. There was considerable variation in earnings, with one subject taking home $3 and another subject taking home $205.

Our betting task provides a clean counterpart to the theoretical framework used for decades to operationalize what is meant by subjective beliefs. Consider two recent examples from the literature. Machina [2004; p.2] carefully defines two ways of representing uncertainty. One he calls "objective uncertainty," and involves known probabilities and choices over lotteries. The other he calls "subjective uncertainty," and is represented by mutually exclusive and exhaustive states of nature, and where the objects of choice consist of bets or acts which yield outcomes that depend on the realized state of nature. Similarly, Klibanoff, Marinacci and Mukerji [2005; p.1854] stress the importance of modeling preferences over what they call "second order acts" which assign utility-relevant consequences to the events that the subject is uncertain about. They

> suggest that second order acts are not as strange or unfamiliar as they might first
> appear. Consider any parametric setting, i.e., a finite dimensional parameter space
> [such that the elements of this parameter space define the subjective belief]. Second

---

[11] Subjects were also informed that "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." The experiments were conducted in Florida, after all.

order acts would simply be bets on the value of the parameter. In a parametric
portfolio investment example, these could be bets about the parameter values that
characterize the asset returns, e.g., means, variances, and covariances. Similarly, in
model uncertainty applications, second order acts are bets about the values of the
relevant parameters in the underlying model. Closer to decision theory, for an
Ellsberg urn, second order acts may be viewed as bets on the composition of the
urn.

This is exactly the choice task our subjects faced when one views the scoring rules as an array of

bookies, each with different odds to place bets with. Each report in the scoring rule generated a

virtual bookie willing to take odds on the outcomes, so the choice of a report is a choice of a bet to

place based on the subject's beliefs about "the composition of the urn."[12]

3. Econometric Model

We explain the econometric model to be estimated in five stages. First we restate the earlier

numerical example in terms of the experimental tasks our subjects faced, so we can see how to

translate the logic presented earlier to these tasks. Second, we present the specification of risk

attitudes assuming an EUT model of latent choice, where the focus is entirely on the concavity of

the estimated utility function $v(\cdot)$. Third, we consider the joint estimation of risk attitudes and

subjective probability, using the conventional SEU specification that assumes that the decision

---

[12] It is possible to design artefactual laboratory experiments that can do even more. Imagine an
Ellsberg urn with unknown mixtures of orange and white balls, but where the subject is told that there are
only 10 balls. The urn is presented, but with a thick towel draped over it. The subject is asked to place bets on
each of the 11 possible mixtures: they earn $1 if the mixture is in fact that one, $0 otherwise. Then simply
remove the towel and have the composition of the urn verified. Our scoring rule procedure rewards subjects
based on bets defined over a single *realization* from this urn, not from bets defined directly over the *composition*
of the urn. Of course, beliefs about the realization depend on beliefs about the composition, but they are not
the same thing. The critical problem with this design is that it does not extend to naturally occurring events,
such as the outcome of an election or an economic indicator. In those cases one has only the realization to bet
on; hence we focus on designs that extend beyond the lab. One important intermediate experimental setting
are bets defined over virtual reality simulations, such as in the wild fire simulations of Fiore, Harrison, Hughes
and Rutström [2009]. In this case one can elicit bets over realizations (e.g., "does your virtual house burn
down when we pick one of the 11 possible states of nature?") or over the process (e.g., "will your virtual
house burn down in only 1 of the possible states of nature? In 2 of the possible states of nature? Etc.").

maker is neutral to uncertainty. Fourth, we consider the joint estimation of risk attitudes, uncertainty attitudes, and subjective probability assuming the SDRA model. Finally, we consider the extension to the general UP model, under the identifying assumptions noted earlier.

A. An Example

Assume that the subject has a true subjective probability that Obama would win the election of 0.68, and is facing the LSR payoffs. The key parameters α and β are defined as before (α=0.5 and β=0.9). If the subject reports truthfully the payoffs are $68 or $32 depending on who wins the election. Now assume that the subject considers the merits of three possible reports of 0.68, 0.67 and 0.69. When α=β=1, and the subject is completely risk neutral and uncertainty neutral, the optimal response of these three is to report 0.69: as is well known, with the LSR a risk neutral subject will go to the extreme (feasible) report that matches the side of 0.5 that their true subjective belief is on. With our default parameter values for α and β, the overall evaluations using the SDRA model are W(0.68) = 6.063, W(0.67) = 6.054, and W(0.69) = 6.071; so the subject still gains from reporting higher than 0.68, but this incentive only takes the subject as far as a report of 0.89 if we extend the range of reports to the interval [0, 1]. If the subject is uncertainty neutral, and β=1 while α=0.5, the optimal report is 0.92, so we see that the degree of uncertainty aversion makes an observable difference to behavior here.

The evaluations in our experimental task are simpler than for the Ellsberg example. Replacing events r and b with m and o, the EU of each outcomes is the same as the conventional utility valuations of each payoff since the payoffs are degenerate in the election outcomes. That is, it is as if the roulette wheel in the Anscombe-Aumann device always paid out the same amount if horse Obama wins and the same (but different) amount if horse McCain wins. Specifically, for the

0.67 report, the roulette wheel always pays $67 for Obama, and always pays $33 for McCain. In the Ellsberg example these payouts were conditional on a further roll of the roulette wheel to decide if R or B occurred.

Using the QSR payoffs, we find that the optimal report of the three is 0.67, which is again consistent with our *a priori* knowledge that risk aversion and uncertainty aversion pulls the decision-maker towards a report of 0.5, where all risk and uncertainty is removed. We confirm that reporting 0.68 is optimal under the QSR if the subject is risk neutral and uncertainty neutral ($\alpha=\beta=1$).

B. Risk Attitudes under Expected Utility Theory

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

$$v(y) = [1-\exp(-\alpha y^{1-r})]/\alpha, \tag{10}$$

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 $\alpha \to 0$) and CARA (as $r \to 0$).

The utility function (10) can be estimated using maximum likelihood and a latent EUT structural model of choice. Let there be S possible outcomes in a lottery; in our lottery choice task $S \leq 4$. Under EUT the probabilities for each outcome s in the lottery choice task, p_s, are those that are induced by the experimenter, so expected utility is simply the probability weighted utility of each outcome in each lottery c:

$$EU_c = \sum_{s=1,S} [\, p_s \times v_s \,]. \tag{11}$$

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

$$\nabla EU = eu_R / (eu_L + eu_R) \tag{12}$$

calculated, where

$$eu_R = \exp(EU_R) \tag{13}$$
$$eu_L = \exp(EU_L) \tag{13'}$$

and EU_L is the "left" lottery and EU_R is the "right" lottery, as displayed to the subject and illustrated in Figure 5. This latent index, based on latent preferences, is already in the form of a probability.[13]

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 indifference[14] the log-likelihood is then

$$\ln L(r, \alpha; y, \mathbf{X}) = \sum_i [(\ln \nabla EU \times \mathbf{I}(y_i = 1)) + (\ln (1 - \nabla EU) \times \mathbf{I}(y_i = -1))] \tag{14}$$

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 $\mathbf{X}$ includes data on the characteristics of the choice task (e.g., value of the prizes and probabilities) or the subject (e.g., sex). To allow for subject heterogeneity with respect to risk attitudes, the parameters r and α can each be modeled as linear functions of observed individual characteristics of the subject, as explained in Harrison and Rutström [2008; p72ff.]. That is not essential to our methodological point here, but can provide better estimates.

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

---

[13] It is well known, but useful to note, that (12) is equivalent to $\Lambda(EU_R - EU_L)$ where $\Lambda(\cdot)$ is the logistic cumulative density function. Thus (12) embodies a statistical "link function" between the difference in the EU of the two lotteries and the probability of the observed choice.

[14] 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 by recognizing that such choices implied 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 and adds needlessly to notation.

(12) that the probability of choosing a lottery is not 1 when the EU of that lottery exceeds the EU of the other lottery.[15] By varying the shape of the link function implicit in (12), one can informally imagine subjects that are more sensitive to a given difference in the index ∇EU and subjects that are not so sensitive. We use the contextual error specification proposed by Wilcox [2009]. It posits the latent index

$$eu_i = \exp[(EU_i/\nu)/\mu], \qquad\qquad (13'')$$

instead of (13) and (13'), where ν 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 ν 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 $[(EU_R - EU_L)/\nu]$ remains in the unit interval. As $\mu \to \infty$ this specification collapses ∇EU to 0 for any values of EU_R and EU_L, so the probability of either choice converges to ½. So a larger μ means that the difference in the EU of the two lotteries, conditional on the estimate of r and α, has less predictive effect on choices. Thus μ can be viewed as a parameter that flattens out, or "sharpens," the link functions implicit in (12). 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, α, μ; y, **X**).[16] Additional details of the

---

[15] This assumption is clear in the use of a link function between the latent index ∇EU and the probability of picking one or other lottery; in the case of the normal CDF, this link function is $\Phi(\nabla$EU). If the subject exhibited no errors from the perspective of EUT, this function would be a step function: zero for all values of ∇EU<0, anywhere between 0 and 1 for ∇EU=0, and 1 for all values of ∇EU>0. Harrison [2008; p.326] illustrates the implied CDF, referring to it as the CDF of a "Hardnose Theorist."

[16] The normalizing term ν is given by the values of r, α and the lottery parameters, and the latter parameters are part of **X**.

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].

C. Estimating the Subjective Probability

The responses to the belief elicitation task can be used to estimate the subjective probability that each subject holds if we are willing to assume something about how they make decisions under risk and that they are uncertainty neutral.

If they are assumed to be risk neutral, then we can directly infer the subjective probability from the report of the subject.[17] 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%. In that case the behavioral error story has a lot of explaining to do, if one wants to be formal. 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 report "risk neutral subjective beliefs" for the LSR.

Moving to the models that allow for general 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 EU

$$EU_\theta = \quad \pi_A \times v(\text{payout if A occurs} \mid \text{report } \theta) + \\ (1-\pi_A) \times v(\text{payout if B occurs} \mid \text{report } \theta) \tag{15}$$

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 the specific report θ, and span the interval [$0, $100]. For the

---

[17] 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.

QSR and a report of 75%, for example, we have

$$EU_{75\%} = \pi_A \times v(\$93.75) + (1-\pi_A) \times v(\$43.75) \qquad (15')$$

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

$$EU_{75\%} = \pi_A \times v(\$75) + (1-\pi_A) \times v(\$25) \qquad (15'')$$

and so on for other possible reports. We observe the report made by the subject, and we know that they had 101 possible reports defined over percentage points, so we can calculate the likelihood of that choice given values of r, α, π_A and μ. In this case the likelihood is the multinomial analogue of the binary logit specification used for lottery choices. We define

$$eu_\Theta = \exp[(EU_\Theta/v)/\mu] \qquad (16)$$

for any report Θ, analogously to $(13')$, and then

$$\nabla EU = eu_\theta/(eu_{0\%} + eu_{1\%} + ... + eu_{100\%}) \qquad (17)$$

for the specific report θ observed, analogously to (12).

We need r and α to evaluate the utility function in (10) and then (15), we need π_A to calculate the EU_θ in (15) for each *possible* report Θ in {0%, 1%, 2%, ..., 100% } once we know the utility values, and we need μ to calculate the latent indices (16) and (17) 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 parameters that best explain observed choices in the belief elicitation tasks as well as observed choices in the lottery tasks. In effect, the lottery task allows us to identify r and α under EUT, since π_A plays no direct role in explaining the choices in that task. Observed individual heterogeneity can be allowed for by estimating the π parameter as a linear functions of demographic characteristics, as well as allowing risk attitudes to vary with demographic characteristics. We focus on allowing heterogeneity with respect to risk attitudes.

D. Estimating Uncertainty Aversion

Recall that when we only allow for risk aversion, which is the same thing in the SDRA and

UP models as assuming that the decision-maker is uncertainty-neutral, we evaluate the EU using

(15). When we allow for uncertainty aversion we need to specify a functional form for uncertainty

aversion, and then show how (15) is modified to embed it in the evaluation of lotteries implied by

different scoring rule reports. We employ a simple Constant Relative Uncertainty Aversion (CRUA)

specification

$$u(z) = z^{1-\phi}/(1-\phi) \tag{18}$$

for uncertainty aversion parameter ϕ. Hence $\phi>0$ implies uncertainty aversion, $\phi=0$ implies

uncertainty neutrality, and $\phi<0$ implies uncertainty loving.

For the SDRA model we evaluate the expected utility of a report θ as

$$EU_\theta = \quad \rho_A \times u\,[v(\text{payout if A occurs} \mid \text{report } \theta)] + \\ (1-\rho_A) \times u\,[v(\text{payout if B occurs} \mid \text{report } \theta)] \tag{19}$$

instead of (15). Because of the latent assumptions about the two uncertain processes in the SDRA

model, estimates of ρ_A are the weighted-average estimates of the subjective belief that A will occur

(because ρ_A weights the degenerate subjective probability 1 and $1-\rho_A$ weights the degenerate

subjective probability 0). Even though the argument of $u(\cdot)$ is *not* a lottery, we need $v(\cdot)$ in order to

evaluate the payouts.[18] This point is stressed in the theoretical literature on the comparison of

uncertainty aversion or "ambiguity attitudes" across decision-makers (e.g., Klibanoff, Marinacci and

Mukerji [2005; §3.2]). The upshot is that one *must estimate risk aversion and uncertainty aversion jointly*,

absent assumptions that rule one or the other out.

---

[18] To repeat, in the original Ellsberg two-urn example, the argument of $u(\cdot)$ *was* a lottery, and one would need to evaluate that argument as the EU of the lottery in that case. It is not a lottery in our application to scoring rules, but the formal "recursive EU" representation is essentially the same.

For the UP model assume $\kappa=2$ and that we impose the identifying assumption on the latent subjective prior probabilities that $\rho_1 = \rho_2 = \frac{1}{2}$. In this case we would estimate two subjective probability values π^U_{1j} and π^U_{2j}, where j=1 is when Obama wins by more than 5%, j=2 is when this event does not occur, and J=2. It is worth noting how this compares to the SDRA model: in the SDRA model we constrained the two subjective probability values to be 0 and 1 by assumption, and estimated ρ_1. Here we constrain ρ_1 to be $\frac{1}{2}$, so ρ_2 is thereby constrained to be $\frac{1}{2}$ since I=κ=2, and estimate the two subjective probability values. We evaluate the expected utility of a report θ as

$$EU_\theta = \frac{1}{2}\, u[\, \pi^U_{1A}\, v(\text{payout if A occurs}\,|\,\text{report } \theta) + (1-\pi^U_{1A})\, v(\text{payout if B occurs}\,|\,\text{report } \theta)\,]$$
$$+ \frac{1}{2}\, u[\, \pi^U_{2A}\, v(\text{payout if A occurs}\,|\,\text{report } \theta) + (1-\pi^U_{2A})\, v(\text{payout if B occurs}\,|\,\text{report } \theta)\,]. \quad (20)$$

The generalization to $\kappa>2$ is obvious, if demanding from an identification and estimation perspective. The special case in which $\kappa=1$, noting that I=κ, would imply that we simply evaluate

$$EU_\theta = u[\, \pi^U_{1A}\, v(\text{payout if A occurs}\,|\,\text{report } \theta) + (1-\pi^U_{1A})\, v(\text{payout if B occurs}\,|\,\text{report } \theta)\,] \quad (21)$$

So stated, this special case is actually quite odd. Unlike the $\kappa=2$ version (20), or the SDRA model (19), it posits no "uncertainty" at all in the latent subjective beliefs.

Let Π be the prior-weighted average of the subjective probability distributions:

$$\Pi = \sum_{i=1,I} \rho_i \sum_{j=1,J} \pi^U_{ij} \quad (22)$$

This will turn out to be a useful statistic to be able to compare estimates across the SDRA and UP models.

4. Results

We first examine the raw data elicited from the scoring rule tasks, then the estimated risk attitudes from the lottery tasks over objective probabilities, then the estimated subjective probabilities assuming an SEU model in which uncertainty aversion is assumed away but where choices in the betting task are conditioned on risk aversion, and finally the estimated subjective

probabilities and uncertainty aversion assuming an SDRA and UP model. We stress again that we do not estimate the most general form of the UP model, nor do we believe that is possible.

A. Raw Elicited Beliefs

The average report from the QSR was 0.59, with a standard deviation of 0.23 and a median of 0.60. Virtually identical statistics are observed for the LSR (mean of 0.58, median 0.6, standard deviation 0.26). The fact that a majority appeared to think it likely that Barack Obama would win by 5 percentage points or more accords with our priors at the time, although the probability of this size victory was much lower than the probability of his outright victory. As it happened, the *New York Times* did report that the number of votes for Obama was 64.5 million and the number of votes for McCain was 56.7 million, so the popular votes are 52.5% and 46.2%, resulting in a difference of 6.3%. Hence the popular vote did have a winning margin of 5 percentage points or more.

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

B. 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 $r>0$, and r defines RRA at $y=0$), and evidence of increasing relative risk aversion as the prizes climb to \$100 (since $\alpha>0$). Specifically, we estimate $r=0.297$ and $\alpha=0.0286$, with each being statistically significantly different from zero, with *p*-values below 0.001. Given these parameter estimates we can calculate RRA at various prize levels: at \$25, \$50, \$75 and \$100 the RRA is estimated to be 0.58, 0.67, 0.75 and 0.82, respectively. Thus subjects

exhibit greater risk aversion for the higher stakes in the belief task than they do for the lower stakes. We do not report detailed estimates for these data, because we do report estimates for the more interesting cases below in which we jointly estimate risk attitudes, subjective probabilities, and then uncertainty aversion.

C. Estimating Subjective Probabilities

Given that we find evidence of risk aversion in our subjects over the domain of prizes on offer in their belief elicitation tasks, our estimated subjective probability, assuming no uncertainty aversion, is a translation of the raw responses *away* from the 50% response. The reason, again, is that risk averse subjects are drawn to respond toward 50% 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 qualitative logic, to be able to quantify the extent of the required translation and the precision of the resulting inference about the latent subjective probability.

Table 1 presents detailed estimation results, for a model that assumes homogeneous risk preferences and then for an extension to allow for heterogeneous risk preferences that are reflected in observable demographic characteristics.[19] We estimate a marked translation from the raw response average of 59% to 73%, although the 95% confidence interval on this estimate, between 54% and 92%, does include 59%. Although the estimates of r and α are jointly estimated with the data from the objective lottery tasks and the scoring rule tasks, they are effectively "pinned down" recursively by the former. Allowing for heterogeneous risk attitudes, we estimate a slightly lower latent

---

[19] The log-likelihood values show that adding demographics improve fit, and this improvement is always statistically significant.

subjective probability of 70% with a tighter 95% confidence interval between 56% and 83%. This is close to the representation illustrated in Figure 2.[20]

D. *Estimating Subjective Probabilities and Uncertainty Aversion*

When we extend the estimation to allow attitudes towards uncertainty to play a role, we find evidence of *uncertainty loving*, and a sharp *reduction* in the estimated subjective probability. Table 2 shows results for the SDRA model. We estimate ϕ to be -0.66, with a standard error of 0.282, a *p*-value of 0.019, and a 95% confidence interval between -1.21 and -0.11. The resulting estimate of ρ_1 is 63%, with a 95% confidence interval between 52% and 74%. Because of the assumptions that the two underlying subjective probability distributions are degenerate at values of 1 and 0 ($\pi^U_{11}=1$, $\pi^U_{21}=0$), the estimate of ρ_1 is the estimate of the *average of the two subjective probability distributions*, Π. In terms of Figure 1, the actual estimates reflect a density of 0.63 at the subjective probability 1, and a density of 0.37 at the subjective probability 0. The standard error on the estimate of the subjective probability drops from 9.7 percentage points to 5.4 percentage points when we allow uncertainty aversion. Given the *p*-value on the estimated coefficient for ϕ, we can reject the hypothesis that the estimated SDRA model is statistically the same as the estimated SEU model. Extending the estimation to allow for heterogeneous risk attitudes, the estimate of ϕ drops to -0.80 and the estimate of Π drops to 60%.

We therefore obtain the striking result that subjects behave in an entirely different qualitative way towards risk (objective risk) as towards uncertainty (subjective risks). They are risk *averse*, and yet at the same time uncertainty *loving*. At an anecdotal level, our "field" observations of the subjects are

---

[20] Note that the *sampling* error on our estimate is different from any uncertainty in the *population* parameter being estimated: so the probability mass in Figure 2 is all concentrated at 0.71, even if our sample estimate of it reflects some statistical uncertainty in that value.

consistent with this finding. The objective lottery tasks were something of a familiar academic chore, although ones that have some serious prizes of up to $100 dangling in front of them. We were frankly astonished to see how excited the subjects were about the draws from the bingo cages used to introduce the scoring rule task: the nightlife in Orlando is not that bad. This interest persisted throughout the belief elicitation task we focus on here, given the general public attention on the election at that time. Of course, none of these observations necessarily translate into uncertainty loving behavior, but they are consistent with the subjects having different attitudes towards the objective and subjective stochastic processes, which is the central assumption of the SDRA model. Moreover, betting on an imminent, historical election outcome is a far cry from betting on an urn in which you are being "deliberately" left in the dark.

Table 3 extends the analysis to the 2-point version of the UP model. Estimates of ϕ drop to -1.33 when heterogeneous risk attitudes are assumed, with a 95% confidence interval that reflects uncertainty loving, and the estimate of Π drops to 59%. The 95% confidence interval for Π is quite tight, between 52% and 66%, reflecting a further drop in the standard error compared to the SDRA model. In terms of the general characterization of Figure 4, this estimate of Π reflects underlying estimates of $\pi^{U}_{1j}=0.19$ and $\pi^{U}_{2j}=0.99$, since $(0.19+0.99)/2=0.59$.

It is striking that we are exactly back to where the "raw data" started, with a latent subjective belief of 59%. Of course, there is no reason for this to have occurred, and it reflects offsetting effects from risk aversion and uncertainty loving, whereas any inference about subjective beliefs from the raw data alone required the strong assumptions of risk neutrality and uncertainty neutrality (so it is not so "raw" as it might seem).

7. Conclusions

We demonstrate that one can *jointly* estimate attitudes towards uncertainty, attitudes towards risk, and subjective probabilities in a rigorous manner. Our structural econometric model constructively demonstrates the theoretical claims by Ergin and Gul [2009], Klibanoff, Marinacci and Mukerji [2005], Nau [2001][2006][2007] and Neilsen [1993][2008] that it is possible to define uncertainty aversion in a parsimonious and empirically tractable manner. Our results show that attitudes towards risk and uncertainty can be different, qualitatively and quantitatively, and that allowing for these differences can have significant effects on inferences about subjective probabilities.

Much work remains to explore the empirical value of the most general characterizations of uncertainty aversion, the implications of this decomposition for understanding behavior in other settings (e.g., whether "strategic uncertainty" in games is better characterized as risk or uncertainty), and to continue the comparative evaluation of alternative approaches in the theoretical literature. It would also be valuable to undertake a controlled comparison of complementary experimental designs and statistical procedures in the existing literature.

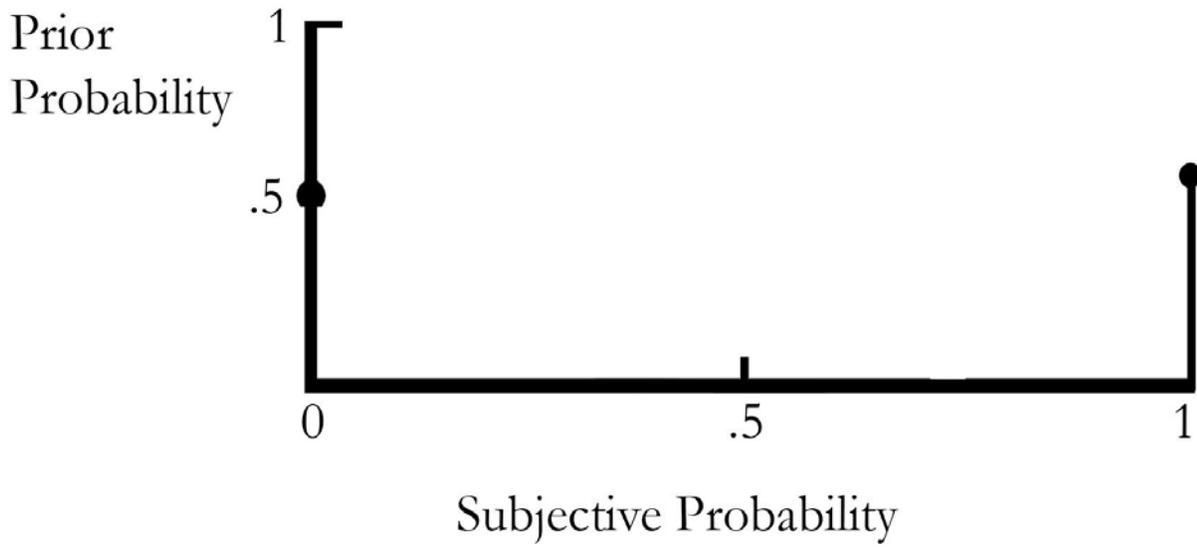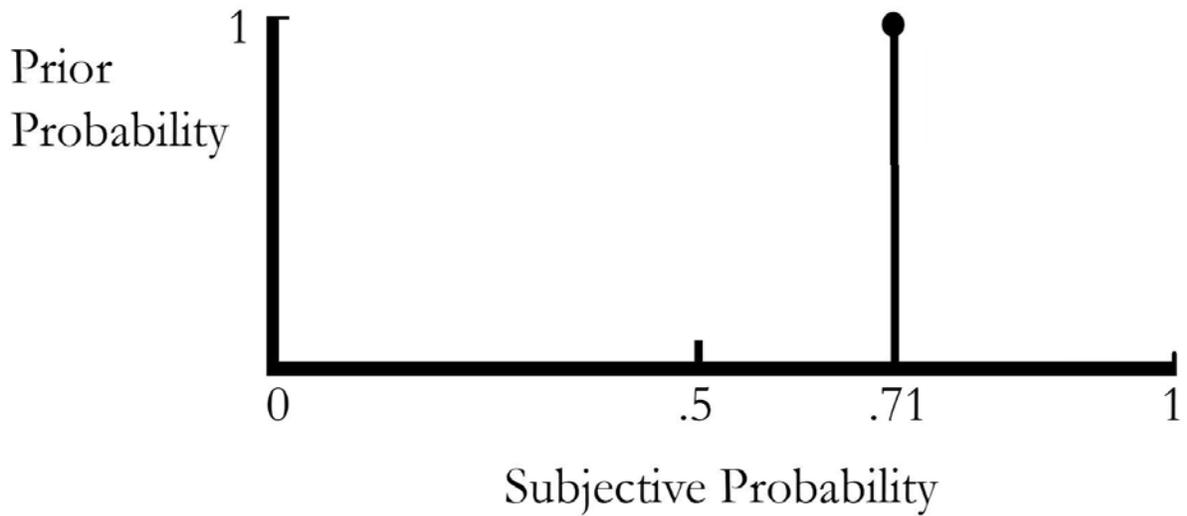## Figure 1: Source-Dependent Risk Attitudes Model

Figure 2: Subjective Expected Utility Model

**Figure 3: The Uncertain Priors Model with
Normally Distributed Subjective Beliefs**

Prior
Probability

1

0   .5   .6   1

Subjective Probability
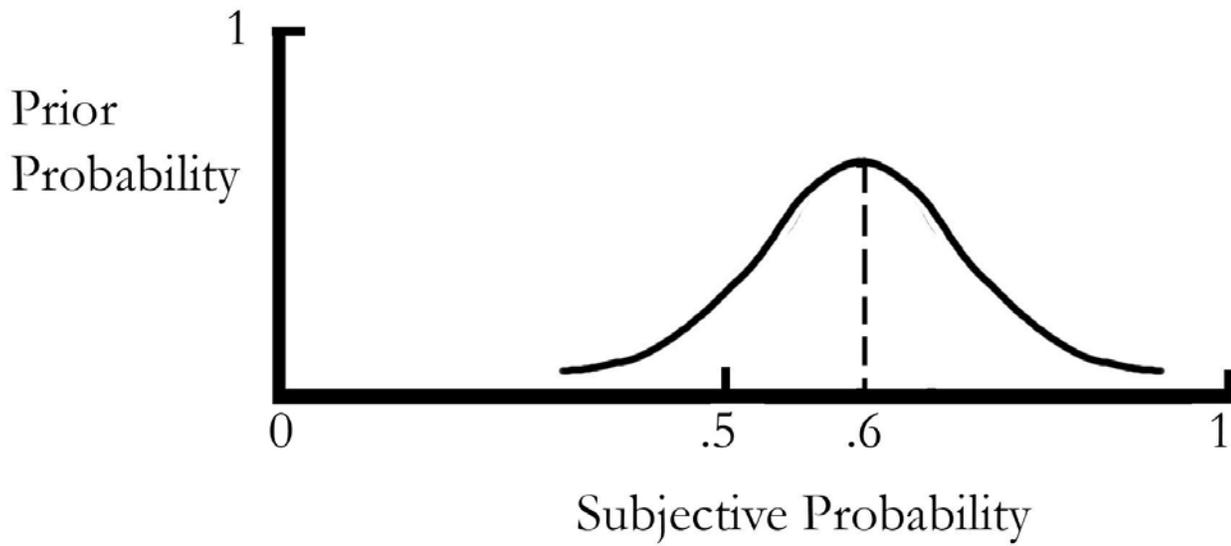
**Figure 4: The Uncertain Priors Model with
a 2-Point Discrete Uniform Distribution of Subjective Beliefs**

Prior
Probability

1

.5

0   .42   .5   1
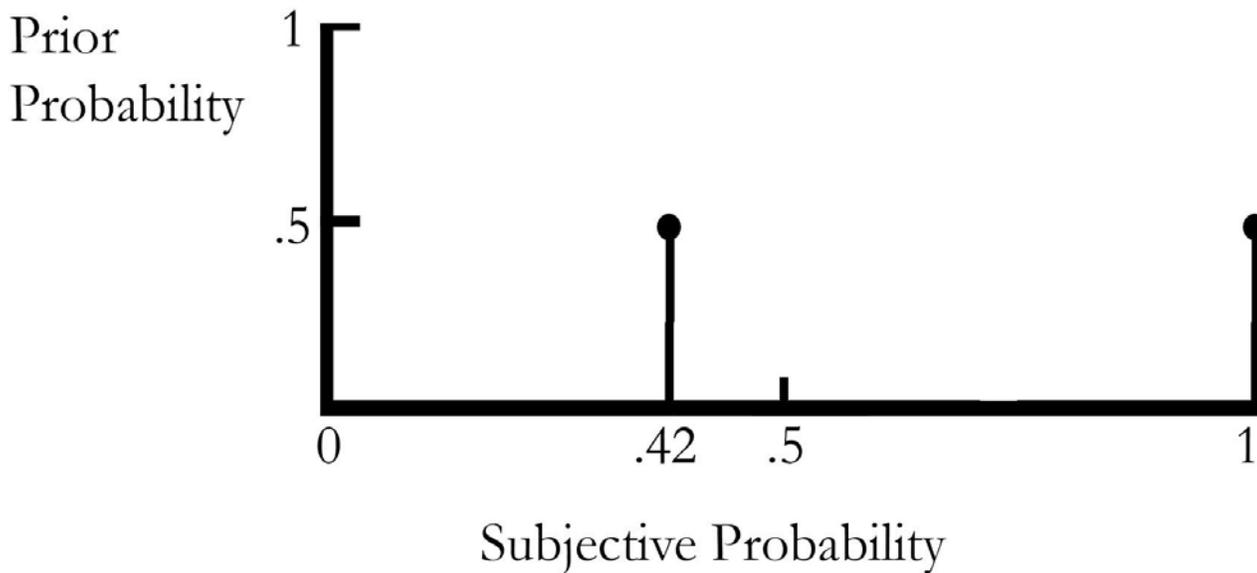
Subjective Probability

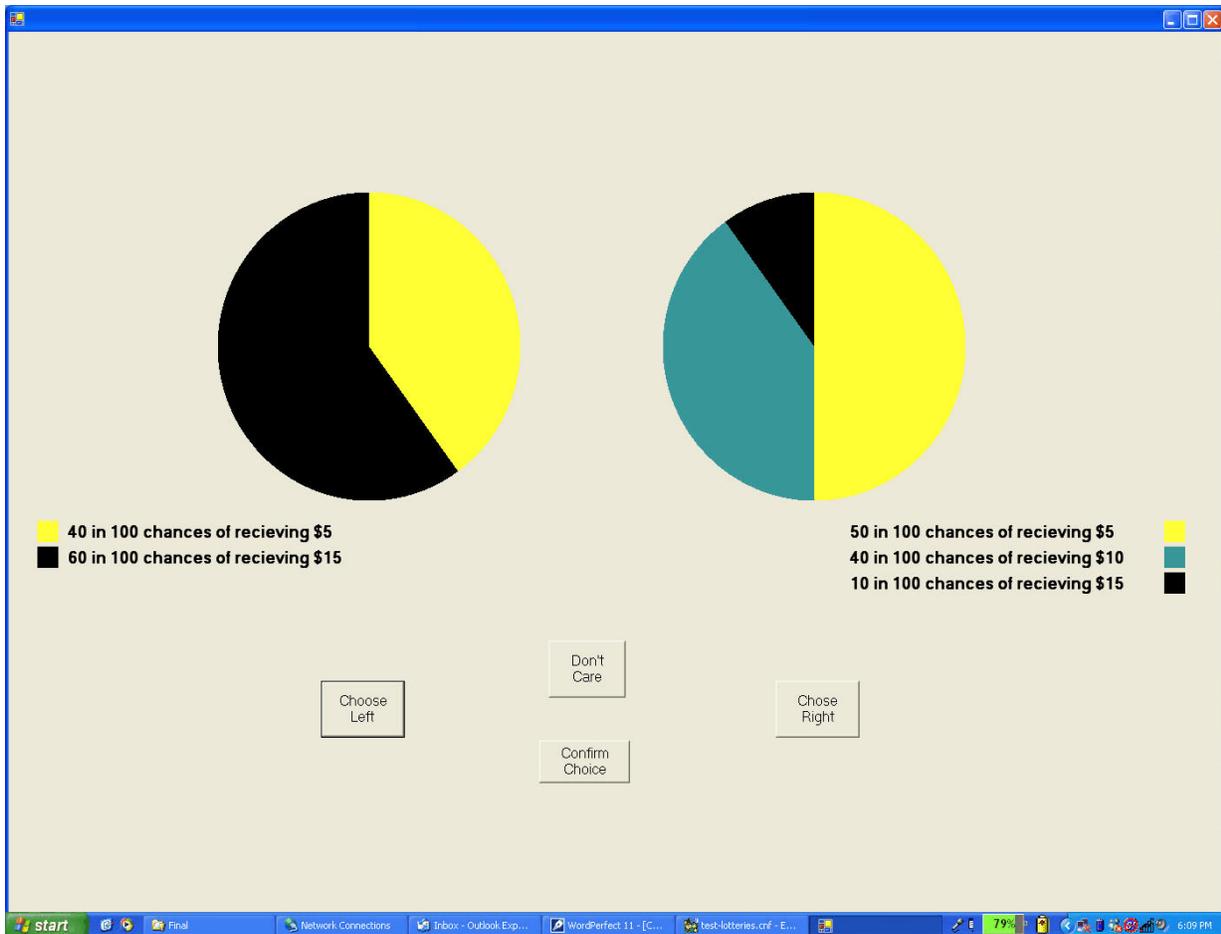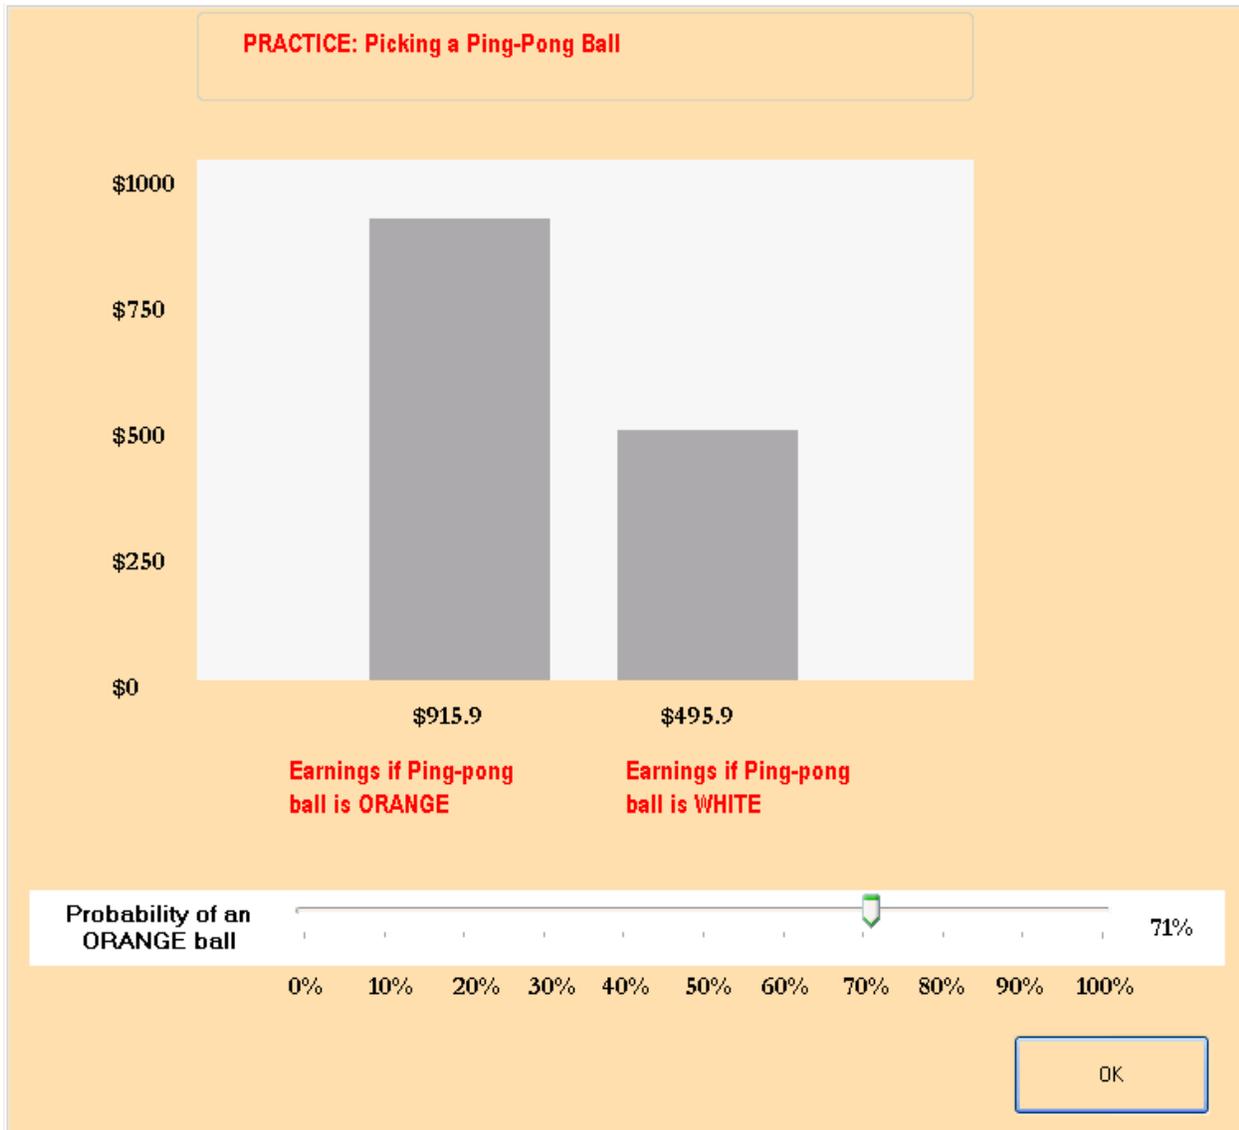**Figure 5: Illustrative Lottery Choice with Objective Probabilities**

Figure 6: Illustrative Quadratic Scoring Rule Interface

Table 1: Estimated Subjective Probabilities
Assuming Subjective Expected Utility

Variable	Description	Estimate	Standard Error	*p*-value	Lower 95% Confidence Interval	Upper 95% Confidence Interval
A. Homogenous Risk Attitudes (LL = -4579.32)						
r	RRA at zero income	0.296	0.034	<0.001	0.229	0.363
α	Normalized increment in RRA over domain	0.029	0.006	<0.001	0.018	0.039
μ_{RA}	Fechner error on risk choices	0.075	0.004	<0.001	0.066	0.083
π	**Subjective probability of a 5% win**	**0.730**	**0.097**	**<0.001**	**0.540**	**0.919**
μ_{BE}	Fechner error on belief choices	0.260	0.109	0.017	0.046	0.475
B. Heterogeneous Risk Preferences (LL = -4523.75)						
r						
Constant		0.464	0.484	0.338	-0.485	1.413
Female	Female	0.078	0.086	0.366	-0.091	0.246
Asian	Asian	0.120	0.137	0.380	-0.148	0.388
Hispanic	Hispanic	0.148	0.103	0.151	-0.054	0.349
Citizen	U.S. citizen	-0.002	0.117	0.990	-0.231	0.228
Age	Age in years	-0.011	0.020	0.583	-0.051	0.029
Business	College of Business Administration major	-0.122	0.075	0.105	-0.269	0.025
Graduate	Graduate student	0.018	0.155	0.904	-0.285	0.322
GPAlow	GPA less than 3.25	0.071	0.084	0.403	-0.095	0.237
Republican	Political affiliation	0.092	0.099	0.352	-0.101	0.284
α						
Constant		-0.052	0.058	0.379	-0.167	0.049
Female	Female	0.018	0.155	0.240	-0.012	0.043
Asian	Asian	-0.001	0.022	0.961	-0.450	0.029
Hispanic	Hispanic	-0.021	0.025	0.404	-0.071	0.088
Citizen	U.S. citizen	0.039	0.025	0.127	-0.011	0.005
Age	Age in years	0.001	0.002	0.461	-0.002	0.044
Business	College of Business Administration major	0.008	0.018	0.650	-0.028	0.047
Graduate	Graduate student	0.009	0.019	0.640	-0.029	0.020
GPAlow	GPA less than 3.25	-0.013	0.017	0.442	-0.047	0.035
Republican	Political affiliation	0.005	0.015	0.714	-0.024	0.064
μ_{RA}	Fechner error on risk choices	0.077	0.006	<0.001	0.064	0.089
π	**Subjective probability of a 5% win**	**0.698**	**0.070**	**<0.001**	**0.561**	**0.835**
μ_{BE}	Fechner error on belief choices	0.206	0.073	0.005	0.063	0.350

Table 2: Estimated Subjective Probabilities and Uncertainty Aversion Assuming Source-Dependant Risk Attitudes Model

Variable	Description	Estimate	Standard Error	p-value	Lower 95% Confidence Interval	Upper 95% Confidence Interval
	A. Homogenous Risk Attitudes (LL = -4577.88)					
r	RRA at zero income	0.296	0.034	<0.001	0.229	0.363
α	Normalized increment in RRA over domain	0.029	0.006	<0.001	0.018	0.040
μ_{RA}	Fechner error on risk choices	0.074	0.004	<0.001	0.066	0.083
ϕ	**Relative Uncertainty Aversion**	**-0.662**	**0.282**	**0.019**	**-1.215**	**-0.109**
π	**Subjective probability of a 5% win**	**0.630**	**0.054**	**<0.001**	**0.524**	**0.736**
μ_{BE}	Fechner error on belief choices	0.189	0.065	0.003	0.063	0.316
	B. Heterogeneous Risk Preferences (LL = -4520.80)					
r						
Constant		0.446	0.521	0.392	-0.575	1.466
Female	Female	0.076	0.086	0.371	-0.144	0.244
Asian	Asian	0.121	0.105	0.160	-0.058	0.386
Hispanic	Hispanic	0.147	0.105	0.160	-0.058	0.353
Citizen	U.S. citizen	0.001	0.116	0.997	-0.228	0.228
Age	Age in years	-0.010	0.022	0.637	-0.054	0.033
Business	College of Business Administration major	-0.124	0.075	0.100	-0.271	0.024
Graduate	Graduate student	0.013	0.158	0.936	-0.297	0.323
GPAlow	GPA less than 3.25	0.073	0.085	0.395	-0.095	0.240
Republica	Political affiliation	0.091	0.097	0.348	-0.099	0.281
α						
Constant		-0.049	0.057	0.396	-0.161	0.063
Female	Female	0.019	0.016	0.215	-0.045	0.046
Asian	Asian	0.001	0.023	0.992	-0.045	0.045
Hispanic	Hispanic	-0.021	0.027	0.426	-0.073	0.031
Citizen	U.S. citizen	0.040	0.024	0.099	-0.008	0.087
Age	Age in years	0.001	0.002	0.698	-0.030	0.045
Business	College of Business Administration major	0.007	0.019	0.698	-0.030	0.045
Graduate	Graduate student	0.010	0.019	0.599	-0.028	0.049
GPAlow	GPA less than 3.25	-0.013	0.017	0.443	-0.046	0.020
Republica	Political affiliation	0.005	0.015	0.740	-0.024	0.034
μ_{RA}	Fechner error on risk choices	0.077	0.006	<0.001	0.064	0.089
ϕ	**Relative Uncertainty Aversion**	**-0.803**	**0.288**	**0.005**	**-1.368**	**-0.239**
Π	**Average subjective probability of a 5% win**	**0.596**	**0.036**	**<0.001**	**0.525**	**0.667**
μ_{BE}	Fechner error on belief choices	0.137	0.041	0.001	0.057	0.218

Table 3: Estimated Subjective Probabilities and Uncertainty Aversion Assuming Uncertain Priors Model

Variable	Description	Estimate	Standard Error	*p*-value	Lower 95% Confidence Interval	Upper 95% Confidence Interval
A. Homogenous Risk Attitudes (LL = -4578.52)						
r	RRA at zero income	0.296	0.034	<0.001	0.229	0.363
α	Normalized increment in RRA over domain	0.029	0.006	<0.001	0.018	0.040
μ_{RA}	Fechner error on risk choices	0.074	0.004	<0.001	0.066	0.083
ϕ	**Relative Uncertainty Aversion**	**-1.262**	**0.607**	**0.038**	**-2.452**	**-0.072**
π	**Subjective probability of a 5% win**	**0.635**	**0.060**	**<0.001**	**0.517**	**0.752**
μ_{BE}	Fechner error on belief choices	0.224	0.085	0.009	0.056	0.393
B. Heterogeneous Risk Preferences (LL = -4521.42)						
r						
Constant		0.441	0.530	0.405	-0.597	1.479
Female	Female	0.076	0.086	0.372	-0.091	0.244
Asian	Asian	0.121	0.136	0.372	-0.144	0.386
Hispanic	Hispanic	0.145	0.105	0.168	-0.061	0.352
Citizen	U.S. citizen	-0.001	0.117	0.991	-0.231	0.229
Age	Age in years	-0.010	0.022	0.650	-0.054	0.033
Business	College of Business Administration major	-0.123	0.075	0.103	-0.270	0.024
Graduate	Graduate student	0.012	0.160	0.939	-0.301	0.325
GPAlow	GPA less than 3.25	0.073	0.086	0.394	-0.095	0.242
Republica	Political affiliation	0.091	0.098	0.349	-0.099	0.283
α						
Constant		-0.047	0.057	0.412	-0.159	0.065
Female	Female	0.019	0.016	0.217	-0.011	0.050
Asian	Asian	-0.001	0.024	0.978	-0.048	0.046
Hispanic	Hispanic	-0.020	0.026	0.437	-0.072	0.031
Citizen	U.S. citizen	0.040	0.024	0.101	-0.008	0.087
Age	Age in years	0.001	0.002	0.534	-0.002	0.005
Business	College of Business Administration major	0.007	0.019	0.705	-0.031	0.045
Graduate	Graduate student	0.011	0.020	0.592	-0.028	0.049
GPAlow	GPA less than 3.25	-0.013	0.017	0.444	-0.046	0.020
Republica	Political affiliation	0.005	0.015	0.731	-0.159	0.065
μ_{RA}	Fechner error on risk choices	0.077	0.006	<0.001	0.064	0.089
ϕ	**Relative Uncertainty Aversion**	**-1.378**	**0.462**	**0.003**	**-2.283**	**-0.472**
Π	**Average subjective probability of a 5% win**	**0.594**	**0.037**	**<0.001**	**0.521**	**0.666**
μ_{BE}	Fechner error on belief choices	0.153	0.049	0.002	0.057	0.250

References

Abdellaoui, Mohammed; Baillon, Aurélien; Placido, Laetitia, and Wakker, Peter P., "The Rich Domain of Uncertainty," *Working Paper*, Econometric Institute, Erasmus University, July 2008.

Ahn, David; Choi, Syngjoo; Gale, Douglas, and Kariv, Shachar, "Estimating Ambiguity Aversion in a Portfolio Choice Experiment," *Working Paper*, Department of Economics, University of California at Berkeley, February 2009.

Andersen, Steffen; Fountain, John; Harrison, Glenn W.; Hole, Arne Risa, and Rutström, E. Elisabet, "Inferring Subjectively Uncertain Beliefs," *Draft Manuscript*, Department of Economics, College of Business Administration, University of Central Florida, 2009.

Andersen, Steffen; Harrison, Glenn W.; Lau, Morten Igel, and Rutström, E. Elisabet, "Eliciting Risk and Time Preferences," *Econometrica*, 76(3), May 2008, 583-618.

Anscombe, F.J., and Aumann, R.J., "A Definition of Subjective Probability," *Annals of Mathematical Statistics*, 34, 1963, 199-205.

Ellsberg, Daniel, "Risk, Ambiguity, and the Savage Axioms," *Quarterly Journal of Economics*, 75, 1961, 643-669.

Ergin, Haluk, and Gul, Faruk, "A Theory of Subjective Compound Lotteries," *Journal of Economic Theory*, 2009 forthcoming.

Fiore, Stephen M.; Harrison, Glenn W., Hughes, Charles E., and Rutström, E. Elisabet, "Virtual Experiments and Environmental Policy," *Journal of Environmental Economics and Management*, 57, 2009, 65-86.

Gärdenfors, Peter, and Sahlin, Nils-Eric, "Unreliable Probabilities, Risk Taking, and Decision Making," *Synthese*, 53, 1982, 361-386.

Gärdenfors, Peter, and Sahlin, Nils-Eric, "Decision Making with Unreliable Probabilities," *British Journal of Mathematical and Statistical Psychology*, 36, 1983, 240-251.

Ghirardoto, Paolo; Maccheroni, Fabio, and Marinacci, Massimo, "Differentiating Ambiguity and Ambiguity Attitude," *Journal of Economic Theory*, 118, 2004, 133-173.

Ghirardoto, Paolo; Maccheroni, Fabio; Marinacci, Massimo, and Siniscalchi, Marciano, "A Subjective Spin on Roulette Wheels," *Econometrica*, 71(6), November 2003, 1897-1908.

Gilboa, Itzhak, and Schmeidler, David, "Maxmin Expected Utility with a Non-Unique Prior," *Journal of Mathematical Economics*, 18, 1989, 141-153.

Harrison, Glenn W., "Neuroeconomics: A Critical Reconsideration," *Economics & Philosophy*, 24(3),

2008, 303-344.

Harrison, Glenn W., and Rutström, E. Elisabet, "Risk Aversion in the Laboratory," in J.C. Cox and G.W. Harrison (eds.), *Risk Aversion in Experiments* (Bingley, UK: Emerald, Research in Experimental Economics, Volume 12, 2008).

Hey, John D.; Lotito, Gianna, and Maffioletti, Anna, "Choquet OK?" *Discussion Paper No. 2007/12*, Department of Economics and Related Studies, University of York, 2007.

Hey, John D.; Lotito, Gianna, and Maffioletti, Anna, "The Descriptive and Predictive Adequacy of Theories of Decision Making Under Uncertainty/Ambiguity," *Discussion Paper No. 2008/04*, Department of Economics and Related Studies, University of York, 2008.

Hey, John D., and Orme, Chris, "Investigating Generalizations of Expected Utility Theory Using Experimental Data," *Econometrica*, 62(6), November 1994, 1291-1326.

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

Klibanoff, Peter; Marinacci, Massimo, and Mukerji, Sujoy, "A Smooth Model of Decision Making Under Ambiguity," *Econometrica*, 73(6), November 2005, 1849-1892.

Machina, Mark J., "Almost-objective uncertainty," *Economic Theory*, 24, 2004, 1-54.

Nau, Robert F., "Uncertainty Aversion With Second-Order Probabilities and Utilities," *2nd International Symposium on Imprecise Probabilities and Their Applications*, Ithaca, New York, 2001; available at ftp://decsai.ugr.es/pub/utai/other/smc/isipta01/s063.pdf.

Nau, Robert F., "Uncertainty Aversion with Second-Order Utilities and Probabilities," *Management Science*, 52, 2006, 136-156.

Nau, Robert F., "Extensions of the Subjective Expected Utility Model," in Ward Edwards, Ralph Miles, Jr., and Detlof von Winterfeldt (eds.), *Advances in Decision Analysis: From Foundations to Applications* (New York: Cambridge University Press, 2007).

Neilson, William S., "Ambiguity Aversion: An Axiomatic Approach Using Second Order Probabilities," *Working Paper*, Department of Economics, Texas & M University, April 1993.

Neilson, William S., "Ambiguity Aversion: An Axiomatic Approach Using Second Order Probabilities," *Working Paper*, Department of Economics, University of Tennessee, January 2008.

Saha, Atanu, "Expo-Power Utility: A Flexible Form for Absolute and Relative Risk Aversion," *American Journal of Agricultural Economics*, 75(4), November 1993, 905-913.

Savage, Leonard J., *The Foundations of Statistics* (New York: Dover Publications, 1972; Second

Edition).

Segal, Uzi, "The Ellsberg Paradox and Risk Aversion: An Anticipated Utility Approach," *International Economic Review*, 28, 1987, 175-202.

Siniscalchi, Marciano, "A Behavioral Characterization of Plausible Priors," *Journal of Economic Theory*, 128, 2006, 91-135.

Smith, Vernon L., "Measuring Nonmonetary Utilities in Uncertain Choices: the Ellsberg Urn," *Quarterly Journal of Economics*, 83(2), May 1969, 324-329.

Wilcox, Nathaniel T., "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, UK: Emerald, Research in Experimental Economics, Volume 12, 2008).

Wilcox, Nathaniel T., "'Stochastically More Risk Averse:' A Contextual Theory of Stochastic Discrete Choice Under Risk," *Journal of Econometrics*, 142, 2009 forthcoming.

Appendix A: Instructions

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.

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.
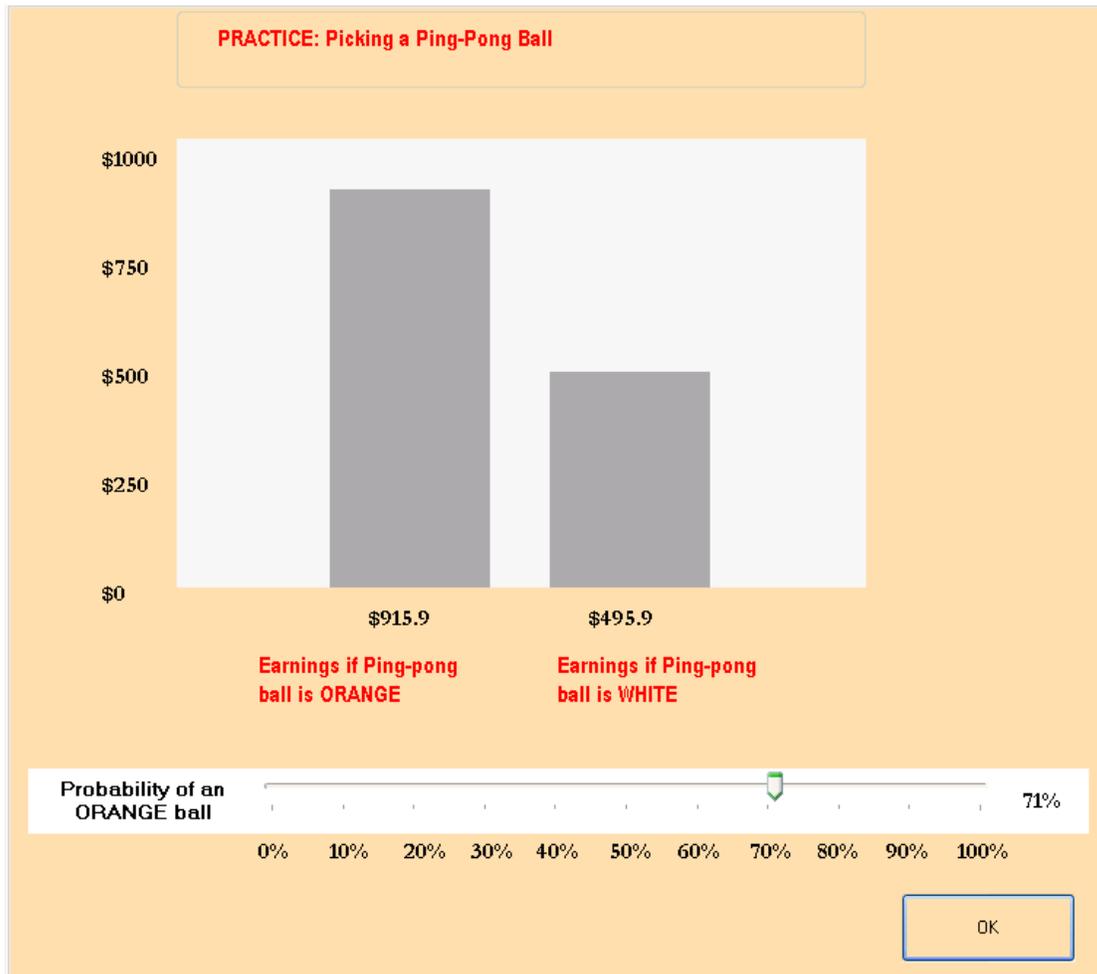
<p align="center">YOU MAY NOW PROCEED WITH THE FIRST STAGE.</p>

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.

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.

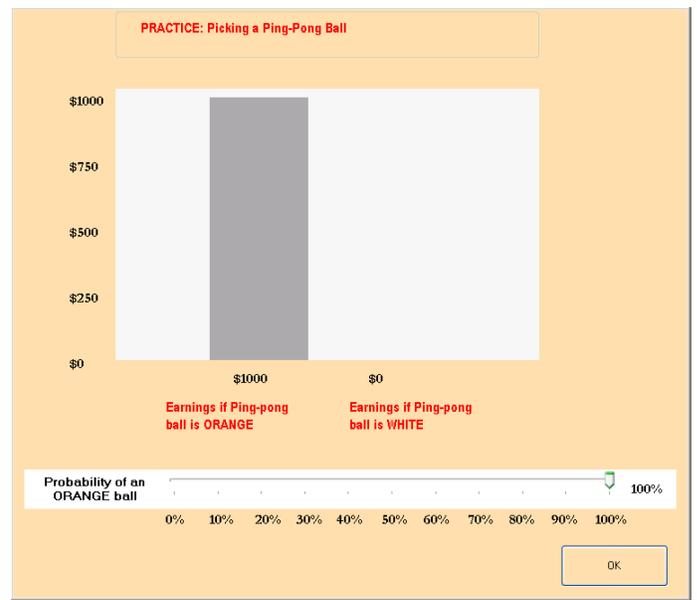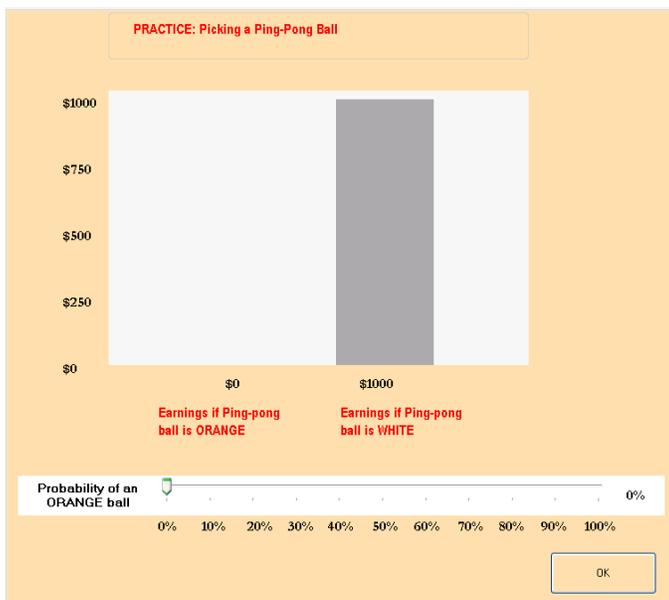ENTER THE PASSWORD THAT IS BEING ANNOUNCED NOW.

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:

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%.

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.

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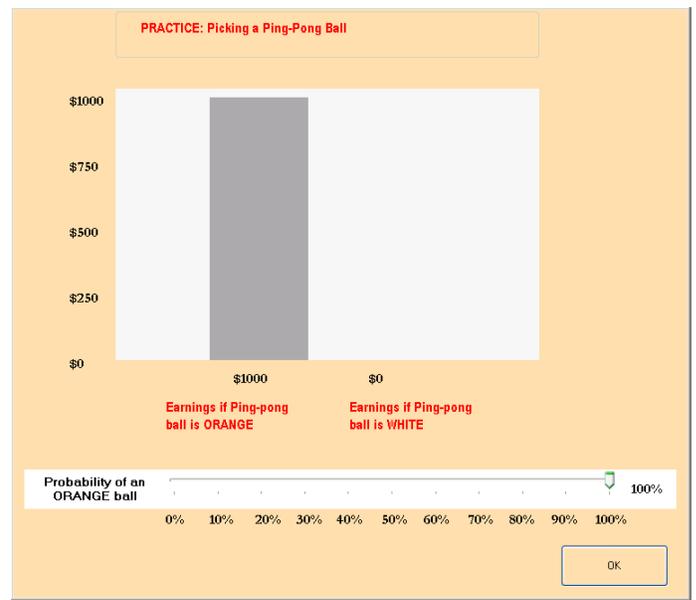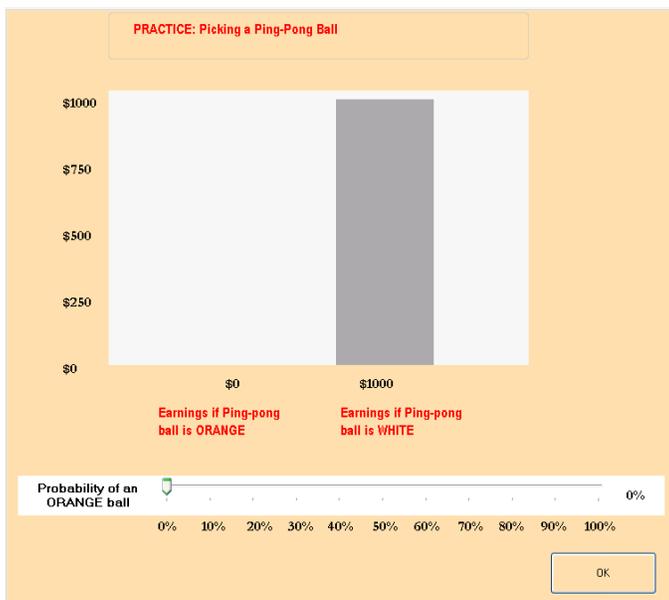
These screens are a little small, but you can see that these two reports lead to extreme payoffs. The "good news" is the possible $1,000 payoff, but the "bad news" is the possible $0 payoff. In between the reports of 0% and 100% you will have some positive payoff no matter what happens, but it will vary,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

2008 Presidential Elections

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

1.   Will the next President of the United States be a Democrat or a Republican?

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?

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.

40 in 100 chances of recieving $5
60 in 100 chances of recieving $15

50 in 100 chances of recieving $5
40 in 100 chances of recieving $10
10 in 100 chances of recieving $15

Don't Care

Choose Left

Chose Right

Confirm Choice

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: Related Literature

There are several complementary approaches to estimating uncertainty (or ambiguity) aversion in the literature, each with strengths and weaknesses. Given the theoretical and inferential delicacy of the task, this is to be welcomed. Many of these studies are designed to compare alternative specifications of uncertainty aversion, whereas our objective is to operationalize one that is closest to received EUT to generate estimates of the degree of uncertainty aversion.[21]

A. Individual Portfolio Managers

Ahn, Choi, Gale and Kariv [2009] develop an interface for portfolio choice that allows them to study how an individual subject allocates a budget across three securities with state-dependent payoffs. Each security pays off \$1, and the subject is confronted with an endowment point and a rate at which they can re-allocate the endowment. One security pays off \$1 with known probability ⅓, but the division of the residual ⅔ probability between the other two securities is not known to the subject. So allocations between the security with the known payoff probability and the other two reflect risk attitudes, and the allocations between the two securities with unknown payoff probabilities reflect risk *and* ambiguity attitudes. Their design has the great advantage of generating enough responses for each individual to allow them to estimate parameters for each individual, following in the tradition of Hey and Orme [1994], and this is obviously an attractive way to handle

---

[21] For example, Halevy [2007] generates ordinal predictions of certainty-equivalents of alternative prospects for each theory, and evaluates the ability of different theories to explain those ordinal patterns. Measures of the degree of uncertainty aversion were not the inferential focus of the experimental design. Bossaerts, Ghirardato, Guarneschelli and Zame [2007] estimate risk attitudes and attitudes towards ambiguity in the context of experimental asset markets, which go beyond the individual choice setting of immediate interest here. Chen, Katuščák and Ozdenorem [2007] estimate a version of the α-MEU representation of ambiguity aversion in the context of first and second price sealed bid auctions. They also explicitly recognize (p.525) the need for identifying assumptions in order to generate structural estimates of parameters.

individual heterogeneity.[22]

A key empirical finding (their Figure 1) is that a relatively large fraction of choices sought to avoid ambiguity by equally splitting the allocations between the two securities with unknown payoffs. Thus, no matter how the residual ⅔ was resolved between these two, the expected payoff would be the same *ex ante*. The key finding is that this "ambiguity and risk avoiding" allocation appeared to occur more often than familiar "risk avoiding" allocations. However, 20% of the choices overall reflected (more or less) perfectly balanced portfolios in which the subject chose the same fraction of lotteries, and these choices were "screened out" of the analysis for some reason (p.8). If they are spread evenly across all budget planes, these balanced-portfolio choices reflect *extreme* risk aversion and *some* uncertainty aversion. If they are only for budget planes in which the relative price of the security with a known probability is close to 1, then they reflect *some* modest risk aversion; if they are only for budget planes in which the relative price of the security with a known probability is greater than 1, then they reflect *some* risk loving (aversion); and if they are only for budget planes in which the relative price of the security with a known probability is less than 1, then they reflect *considerable* risk aversion. Similarly, one cannot ascertain the extent of the atitudes towards uncertainty from these balanced-portfolio choices without knowing all relative prices across securities. So there is no obvious reason why these choices should be screened out, since they are likely to be informative about preferences when evaluated conditional on the budget plane. Nor can

---

[22] Ahn et al. [2009; p.3] suggest that estimation "... at the individual level is *crucial* because of the possibility of individual heterogeneity" (our emphasis). There are alternative ways of addressing individual heterogeneity from an econometric perspective, and we would suggest that individual estimation is not absolutely necessary, even if it is certainly attractive when feasible. Careful comparisons of the ability of different econometric tools for addressing individual heterogeneity is an ongoing topic for future research with rich experimental data of this kind. That is, when one has the luxury of having enough data for individual estimation, do pooled estimators, controlling for observable covariates and/or allowing for random coefficients, capture the true heterogeneity? We simply do not know the answer to this in general.

one make inferences about the extent of ambiguity aversion without conditioning on the budget plane.

Their estimation approach is parametric, and considers "kinked" and "smooth" models of uncertainty aversion; our focus is solely the latter, and specifically the Recursive Expected Utility (REU) specifications. Again, they estimate at the level of the individual, which allows a rich characterization of individual heterogeneity. The functional forms assumed for the REU estimation are Constant Absolute Risk Aversion utility functions for risk, and a Constant Absolute Uncertainty Aversion utility function for the second-order transformation. They correctly stress that the measures of uncertainty aversion depend on the measures of risk aversion, as one would expect from the formal REU representations. But they also stress that interpersonal comparisons of estimates of uncertainty aversion depend on the ability to interpersonally compare estimates of risk aversion as well, and that this implies that the first-order utility function can only be unique up to a cardinal scale. This concern is particularly important for their analysis, because they estimate at the level of the individual, but it is no less a concern, even if it is implicit, when one pools data across subjects as in our analysis.[23]

Estimates are obtained with Non-Linear Least Squares (NLS) rather than Maximum

---

[23] It is not obvious that a post-estimation correction to normalize cardinal utility to have the same cardinal scale across individuals is statistically correct (p.21). First, this normalization could affect the estimates for each individual, and would therefore have inferential implications. This correction would seem to call for something akin to a "seemingly unrelated" non-linear regression approach, or could be directly handled using maximum likelihood. The key is to allow the effects of the normalization on estimated standard errors for one subject to affect the inferences about comparative risk/uncertainty aversion for other subjects. Second, to generate hypothesis tests on the estimates one must normalize the point estimate *and the estimate of the standard error*, and the normalization appears to adjust the former but not the latter. So if one wants to make statements about how many subjects exhibited uncertainty neutrality, for example, one would need the (raw or normalized) point estimate *and* standard error for that subject. It is important to note, following Ahn et al. [2009; p.17], that the need to make joint inferences about risk attitudes and attitudes towards uncertainty is theoretically motivated in the REU specification, and fundamental to it. The issue here is ensuring that the statistical specification correctly honor that requirement for the inference of interest.

Likelihood (ML), but if the assumed error terms are well-behaved the two generate the same results.[24] The fascinating result reported is that 26% of subjects behave as if ambiguity-neutral, in the sense that their estimated uncertainty aversion parameter is not statistically different from zero. Hence these subjects behave as if consistent with SEU. This fraction is misleading, however, because it apparently employs a normalized point estimate and a raw standard error, and these two cannot be combined in this manner. Since this is an inference *for each individual* at a given significance level, and we are then just counting up the number of individuals that pass or fail the test, there is no need for normalization. In this case the detailed individual estimates of Ahn et al. [2009; Appendix IV] show that 43% of subjects had point estimates for the uncertainty aversion coefficient that were smaller than the estimated standard error for that coefficient and individual. This fraction rises quickly to 60%, 62% and 72% if one applies two-sided 10%, 5% and 1% levels of significance, respectively.[25] So these data suggest that a significant fraction of subjects in this task behaved as if exhibiting no uncertainty aversion.

B. *Testing the Source-Dependent Risk Aversion Model Directly*

Chakravarty and Ray [2009] conduct experiments using a multiple price list (MPL) design popularized by Holt and Laury [2002], and modified to provide a direct operationalization of the SDRA model. Their subjects completed a standard MPL task, allowing one to identify their risk

---

[24] Ahn et al. [2009; p.20] argue for NLS "because it provides a good fit and offers flexibility, tractability and straightforward interpretation." Under the maintained error structure, the fit must be the same as ML (e.g., Davidson and MacKinnon [1993; §8.10]). And while "flexibility and tractability" are software-dependent, modern implementations of ML and NLS are extremely easy to program (e.g., see Harrison and Rutström [2008; Appendix F] for a primer on the use of *Stata* for ML estimation of the kind employed here). Finally, the interpretation of NLS and ML should be identical, if we are just concerned with simple inferences such as point estimates and standard errors.

[25] The comparable fractions for the "kinked" parametric representation are 70%, 77% and 86%, respectively, making an even stronger case for SEU in this instance and for this representation.

attitudes using familiar methods. In the main task, the same subjects always faced binary lotteries, and the low payoff was always zero. The subjects had to make 10 choices between an urn in which the high payoff would occur with a known 50% chance, or an urn in which the high payoff would always occur or it would never occur, and where the subjects were not told the chance of either state of the world. The high payoff for the first, risky urn was an amount varying between Indian Rs 140 in the first row down to Rs 20 in the last row. The high payoff in the second, uncertain urn was always Rs 100. We focus entirely on their experiments in a gain frame; they also conduct experiments with a "reflected" loss frame, raising other issues of modeling.

Because the low payoff is always zero, and we can normalize the (first and second order) utility of zero to be zero, the expressions for REU collapse to be extremely tractable. This is a very nice design feature, albeit one that restricts analysis to artefactual lotteries. A power specification is assumed, with r as the coefficient for the usual utility function $v(x)=x^r$ and a as the coefficient for the second-order utility function $u(z)=z^a$. Estimates of $r \approx 0.62$ are obtained using ML over the pooled risk-choice data, and assuming homogenous preferences across individuals.[26] Estimates of the uncertainty aversion parameter a need some assumption about the probabilities of the "lucky urn" and the "unlucky urn," and Chakravarty and Ray [2009; p.211] impose the strong assumption that they are equally likely: so $\rho_1=\rho_2=\frac{1}{2}$ in our earlier notation. They also impose the point estimate of r as a fixed parameter, which is sensible as an initial numerical strategy but ignores the theoretical and statistical interdependence of risk aversion and uncertainty aversion. With ML methods it is a simple matter to estimate r and a jointly, of course, and with the maintained assumption that

---

[26] Chakravarty and Ray [2009] allow for observable demographic characteristics, as well as treatment effects (e.g., order of presentation), but the essentials are captured with this simple case. Our estimates employ their "logit" specification and, following their analysis, contain no "behavioral error" (i.e., let $\mu=\nu=1$ in our (13′)). Nothing essential depends on these simplifications.

$\rho_1=\rho_2=½$ one estimates r=0.73 (standard error 0.026) and a=0.979 (standard error 0.0045), with a log-likelihood value of -227.47. This estimate for a is close to those reported by Chakravarty and Ray [2009; p.211], although the joint estimation results in a slightly higher estimate of r and a (very) slightly lower estimate of a. Under these assumptions there is no evidence for uncertainty aversion on average, since a=1 is consistent with uncertainty neutrality.

Unfortunately, there is an identification problem here. Each subject gets to pick their own "lucky color" first. If we assume that luck is something that some subjects see as metaphysically coherent, then the prior probability for the higher payoff, which we can denote ρ_2, is plausibly greater than ½. If we assume $\rho_2 = ¾$, for example, we estimate r=0.723 (0.026) and a=0.857 (0.0071), which reflects some uncertainty aversion since a<1. But the *log-likelihood for this estimate is identical to the baseline case*: with this experimental design and econometric specification one simply cannot independently estimate uncertainty aversion and the subjective prior probability ρ_2.

C. The Source Be With You!

Abdellaoui, Baillon, Placido and Wakker [2008] adopt an innovative theoretical and empirical strategy to avoid committing to any specific model of ambiguity or decision-making under risk. They start with the concept of "sources of uncertainty," which are defined (p. 3) as "groups of events that are generated by the same mechanism of uncertainty, which implies that they have similar characteristics." In our experiments, for example, one source is the mechanism generating realizations of objective probabilities in the standard lottery choice task, and another source is the mechanism generating realizations of the event that the subjects are placing bets over in the scoring rule task. This is a natural and appealing notion, made explicit in the SDRA representation, and with clear linkages to older concepts in cognitive psychology.

The next step is to allow subjects to be probabilistically sophisticated in a local sense. The concept of "probabilistic sophistication," due to Machina and Schmeidler [1992; p.747], reflects the assumption that decision makers behave as if they employ subjective probabilities of events and utilities of outcomes that are independent of the assignment of outcomes to events.[27] The extension employed by Abdellaoui et al. [2008], following Chew and Sagi [2008] and Ergin and Gul [2009], is to allow probabilistic sophistication "within each source" but to not assume it applies "across sources." Once one defines subjective probabilities coherently for each source, and assuming (reasonably in the lab) that one can *a priori* spot different sources, a simple probability weighting function applied to these probabilities allows for ambiguity aversion and risk aversion. Since Abdellaoui et al. [2008] restrict attention throughout to binary lotteries, the rank-dependence that normally factors into the transformation of subjective probabilities into decision weights does not arise. But the weighting function is, critically, allowed to vary across sources. The SEU approach can then be viewed as assuming global probabilistic sophistication across sources, whether or not one then transforms these probabilities into decision weights (e.g., Offerman, Sonnemans, van de Kuilen and Wakker [2009; §6] and Andersen, Fountain, Harrison and Rutström [2009]).

This "smooth" approach sharply contrasts to the "smooth" approach embodied in REU, which puts the entire modeling burden for uncertainty aversion on to the utility functions. Abdellaoui et al. [2008; p.8] perceptively note that Smith [1969] was the first to argue for "leaving the probabilities alone" in the modeling of Ellsberg behavior, and for explaining that behavior with something loosely akin to what we now see as the SDRA model. The contribution of Abdellaoui et al. [2008] is to leave the probabilities themselves alone, but to allow them to be transformed into

---

[27] One important contribution of non-EUT models of decision making under risk is to relax that independence, although that is not the only contribution.

decision weights using some smooth function that depends on the source.

Two experiments are conducted, one using artefactual stochastic processes and one using natural stochastic processes. The procedures are explained in detail for the artefactual process, and are similar for the natural process. Considerable detail on the experimental procedures for the natural process is provided by Baillon [2008].

Virtual Bingo Ball Processes

The artefactual process was a virtual, computer-generated urn with balls in 8 colors. One urn, K, had known, equal probabilities for each color, and the other urn, U, had unknown probabilities. It is assumed that the subjects trusted that the process on the computer screen would indeed be fair, because the stochastic process generating outcomes could not be verified. A certainty equivalent was elicited for each of 26 lotteries defined over one or other of these urns. For example, one lottery might be defined as paying €25 if the color green came out, and €0 otherwise; another lottery might be defined as paying €25 if the colors green or black came out, and so forth. One critical feature of this artefactual design is that is permits a direct test of the assumption of local exchangeability with respect to events generated by a given source, since those events with the same number of colors should have the same decision weight and hence the same certainty-equivalent.

The procedures for eliciting certainty-equivalents are based on the Iterative MPL developed by Andersen, Harrison, Lau and Rutström [2006]. In the first stage certainty equivalents for 7 gambles were elicited in which the event to be paid consisted of one of 4 colors appearing, so that the objective probability for urn K is ½. Once these 7 values were elicited, NLS estimates of a power utility function were generated, resulting in estimates of the power coefficient r_s, where s indicates the source K or U, and the weighting functions $w_s(p)$ to be applied to the utility of the

higher payoff and probability p. For the certainty-equivalents generated from urn K it is appropriate to view this weighting function as being $w_K(\frac{1}{2})$, the decision weight corresponding to an objective probability of $\frac{1}{2}$. But for urn U we do not know what the subjective probability is: perhaps the subjects though they were unlucky, or that the computer was programmed to be unfair to them. It *appears* that Abdellaoui et al. [2008; p.12] just assume that the decision weight estimated for the responses to urn U reflects a subjective probability of $\frac{1}{2}$. In fact, however, this is where the test of exchangeability is critical. A hypothesis test is conducted to test for exchangeability, and that hypothesis cannot be rejected at conventional significance levels, so they "assume the uniform subjective probability distribution." (p.13). But these hypothesis tests do not show that uniformity holds exactly: they show that one cannot reject the hypothesis that they hold on average. That is, there is *some* statistical uncertainty around the claim that the subjects had a uniform subjective probability.[28]

Of course, the estimates of r_K, r_U, $w_K(\frac{1}{2})$ and $w_U(\frac{1}{2})$ each have standard errors, even if estimated for each individual. And this is true even if the $\frac{1}{2}$ is incorrectly assumed to be known to be $\frac{1}{2}$ with certainty.

---

[28] The data here consist of certainty-equivalents provided by subjects. For the one-color events, there were 3 comparable responses per subject, for the two-color events there were 4 comparable responses per subject, and for the three-color events there were 2 comparable responses per subject. Each type of event was evaluated separately, of course. Take the one-color events only. In this case there were 3 responses for each of 67 subjects. The statistical procedure used is an Analysis of Variance (ANOVA) with repeated measures, which is ANOVA but with a correction factor to reflect the expected correlation of observations for a given subject. The results, pooled over all subjects and responses to the one-color events, imply that exchangeability cannot be rejected with a *p*-value of 0.335 (p.13), so one cannot reject the hypothesis that the means are different across the tested events. This does not say that the responses are literally identical, since in that instance the *p*-value would be much closer to 1. It is not apparent how one applies this test at the level of the individual, although some test was applied: "Because of the central role of this assumption in our analysis, we inspected it also at the individual level, rejecting it only for subject 52 who was accordingly removed from the analysis." (p.13). The problem is that one only has 4 observations for an individual responses to the one-color event, and only 3 or 2 for the other types of events, so it is not clear what statistical test (or power) one can marshal at that level.

The next step is to estimate the weighting functions $w_s(p)$ for $p \neq \frac{1}{2}$. To do this the events defining the lotteries were varied to include all combinations of colors other than one that had 4 colors: so one event might be if green came out, and the objective probability would then be $\frac{1}{8}$. Given values of r_s, and the nice design assumption that the lower payoff is zero, the decision weight is assumed to be inferred directly from the elicited certainty-equivalent after a simple transformation. Again, this is clear for urn K, since the objective probability inside weighting function $w_K(p)$ is known to be p (e.g., $\frac{1}{8}$ in the example above). But just as we do not know that the subjective probability for urn U was exactly $\frac{1}{2}$ in the first stage, we do not know what subjective probability π is being used by the subject in this stage. Again, it appears that Abdellaoui et al. [2008; p.12] assume that the decision weight estimated for the responses to urn U reflect a subjective probability generated by the subjective prior that each color occurs with *exactly* equal probability. So in this instance the inferred decision weight would be assumed to reflect $w_U(\frac{1}{8})$.

One statistical problem with the inferences in the second step is that they do not appear to reflect the uncertainty over the estimate of r_s in the first step. Presumably the point estimate from the first step was used in the second step, which therefore overstates the precision of the inferred decision weights in the second step. Once the decision weights were generated in the second step, they are fit to a flexible probability weighting function using some Mean Square Error statistical specification.

Abdellaoui et al. [2008; p.41] note that they implemented ML estimation procedures based on Harrison and Rutström [2008; Appendix F], as a check on the *ad hoc* procedures outlined above. Although the specification is not provided, this presumably allowed for the error in the estimation of the r_s parameters to affect inferences about the parameters defining the probability weighting function, which would be a significant statistical improvement over the *ad hoc* procedures. Strikingly,

they report linear utility for both sources. They report one parameter of the probability weighting function to be the same across sources, but another parameter to be different: no standard errors of these point estimates are presented, although if a statistical test was done to support these claims then it would have been two-sided with a 5% level of significance (p. 12). The subsequent analysis builds on the estimates obtained from the *ad hoc* procedures explained above, and while it would be useful to see the same analysis with the ML estimates, these are not available.

Accepting the statistical procedures, then, the upshot is that there is virtually no probability weighting for the certain stochastic process generated from urn K, but that there is some significant probability weighting for the uncertain stochastic process generated from urn U. In particular, the function is consistent with "pessimism," in the sense that the value of $w_U(p) < p$ for $p > ¼$, and $w(p) \approx p$ for $p \leq ¼$. Unfortunately, a perfect confound is the assumed value of π for the calculations with respect to urn U: here we use the notation π for the subjective probability that the *subject* is actually using, and p for the probability that the *analysts* is assuming that the subject is actually using. When we have draws from the known urn K then $\pi \equiv p$ is a plausible assumption, even if we agree that there can always be certain subjective cognitive processes that make these known probabilities actually "almost-objectively uncertain" (Machina [2008]). If the subjective probability π was systematically less than the assumed value, and it makes sense that subjects might suspect this increasingly as the odds of the high payoff get larger and larger, then the estimated probability weighting function *could* be generated by these subjective probability and a degenerate EUT-consistent probability weighting function $w(\pi) = \pi$. In general, the deviation of π from the analyst-assumed p, and the extent of probability weighting, appear to be confounded. Even if the hypothesis that $\pi = p$ cannot be rejected based on a statistical test (such as the ANOVA with repeated measures, applied to certainty equivalents of comparable events), the sampling error in assuming that $\pi = p$

exactly need to be accounted for.

Of course, this analysis uses pooled data over all 67 subjects (p.13/14), and considerable

individual heterogeneity is present in these data when evaluated at the level of the individual (p. 15).

But this variability might in turn just reflect variability in the applicability of the same general

confound: maybe different subjects had different subjective beliefs about what the event

probabilities were in urn U. Moreover, it is possible that some subjects exhibit probability weighting

over risky outcomes, and this is subsumed in the estimates of $w_U(p)$ for that subject.

Natural Born Processes

The other experiment used similar procedures but for naturally occurring events to occur on

a specific day 3 months in the future. These experiments, and the detailed procedures they

employed, are presented in Baillon [2008]. One event was how much the French Stock Index would

change, another event was what the Paris temperature would be, and the final event was what the

temperature would be in some remote country (presumably not Antarctica or Saudi Arabia).

One fundamental issue, again, is how these events were "partitioned." A key feature of the

elicitation of the probability weighting function, for a given source, is to be able to know what p is to

be inferred from some response, so one can claim that this is actually $w(p)$ and not $w(p')$ for $p \neq p'$.

This problem already arose with the assumptions apparently required in the artefactual process with

respect to the partition of events derived from the uncertain urn U. But it is even more challenging

here. The objective is to elicit $w(p)$ for an even grid of *subjective* probability values p between ⅛ and

⅞ in increments of ⅛. But this is a subjective probability of ⅛, not an objective probability, so it is

not at all clear how one would elicit the correct subjective probability.

Each subject was asked to state boundary values for each event beyond which there was

almost no chance of the event occurring (p.38). How these values were elicited is not stated, nor is it clear how they could be easily incentivized; in any case, they play no role in the final analysis, but were only used to orient the critical questions to a domain of some relevance for the subject.

The procedures used in this stage are explained by Bailon [2008; p.86]. Consider a subject that says that it is "very unlikely" that the temperature in Paris would be lower than 5°C or higher than 25°C. The subject is first asked to make a choice between betting that the temperature would be below 15°C or betting that it would be above 15°C, for a prize of €140 if correct. So the 15°C here is the midpoint of the initial bounds, but that is not essential to the validity of the subsequent elicitation. Assume the subject picks the first option. Then the second question would be to bet on whether the temperature would be below 10°C or greater than 10°C, where the 10°C comes from the mid-point of the lower bound 5°C and the elicited 15°C from the first question. Assume the subject picks the second option here. Then the third question would be to bet on whether the temperature would be below or above 12.5°C, and so on until the bets would be about a temperature that was less than 0.5°C different from the temperature used in the previous question. This series of questions might lead to the temperature 14°C being selected as the point at which the event is partitioned into two equally likely partitions with subjective probability ½ (e.g., Bailon [2008; Table A3]).[29]

The next series of questions were used to pinpoint the temperature for this subject that identifies the partition corresponding to the subjective probability ¼. The first question in this sequence is for the subject to bet on whether the temperature will be below 9.5°C or will be between 9.5°C and 14°C, where the 9.5°C here comes from the midpoint of 5°C and the elicited 14°C from

---

[29] For some reason the prize for the first two questions in this sequence is €140, then it becomes €130, and finally it is €150 for the last question. This variation should have no effect on behavior.

the first sequence. A comparable sequence might lead this subject to finally select 12.5°C as the temperature that partitions the event into events that will occur with subjective probability ¼ and ¾. To state the obvious, the elicited value 12.5°C is one-quarter of the distance between 5°C and 25°C, the initial extreme bounds elicited for presentation ease only; nor is it mid-way between 5°C and the elicited 14°C.

The end result for this subject is elicited temperatures of 11°C, 12.5°C, 13°C, 14°C, 14.5°C, 16°C and 17.5°C that reflect subjective probabilities of ⅛, ¼, ⅜, ½, ⅝, ¾ and ⅞, respectively.

For the subject illustrated by Baillon [2008; p.86] there were 9 questions needed to elicit the two temperatures discussed above. Conservatively assume that there were 30 questions overall for this subject to elicit all 7 temperatures. Each subject then did this for three sources, for a total of 90 questions under this assumption. Each subject in Baillon [2008] was told that one subject would be selected at random four weeks later, when the realization from the event was known, and one of their questions played out for real. The subject might or might not win, but the experimenter told them that subjects would continue to be randomly selected for payment until 4 subjects had won, for a total payout of up to €600. All experiments were conducted individually and sequentially, so the subjects presumably had no idea how many other subjects they were "competing with" to get paid. In Abdellaoui et al. [2008; p.12] there were 31 subjects, paid about 3 months after the experiment when the uncertainty was resolved. In this case one of the subjects was to be selected at random and one choice played out for real; the subjects appear (p.38) to have participated simultaneously in one group at a time, rather than sequentially, so they would have known that there was a 1-in-31 chance that they would be paid for one of their choices. The upper prize in this case was a significant €1000, offsetting somewhat the low chance that any particular choice would be actually played out.

Tests of the exchangeability hypothesis in Abdellaoui et al. [2008; p.39-40] involve *t*-tests, binomial tests and correlations on bets that should entail indifference. Apparently all tests were conducted for the pooled sample, and for each source, using the elicited certainty equivalents of these bets. For the incentivized responses, there is general support for exchangeability reported as a statistical hypothesis. But there is no claim that it holds exactly; thus there would be some standard errors around the elicited certainty equivalents which should play a role in the final inferences.

Certainty-equivalents for the utility values of the events defined over the uncertain sources were not elicited. Building on the finding from the study of artefactual stochastic processes, that the utility functions each exhibited approximate linearity, the estimated values from the utility function defined over monetary outcomes and objective probabilities was *assumed* by interpolation for the uncertain source tasks. Of course, the approximate linearity in question was based on a statistical evaluation, and was not exact. Thus the assumed, exact linearity ignores the sampling error in the determination of statistical, approximate linearity.

The probability weighting functions were then inferred in more or less the same manner as in the artefactual process study.

The ML procedures that were used as a check on the *ad hoc* statistical procedures in the artefactual process study are not applied here, so one does not have estimates that reflect all sampling error from the estimates at each step.

D. Blowing Air

Hey, Lotito and Maffioletti [2007] developed an elegant experiment built around an ageing British Bingo Blower. This is a device that blows ping poing balls around in a glass cage, and at such a furious pace that nobody can possibly count the number that are each color unless there are very

few balls. The subjects were asked to choose between two lotteries, with each lottery defined over prizes of a loss of £10, a gain of £10, and a gain of £110, all to be subtracted or added to an endowment of £10. So final earnings were £0, £20 and £100, with the intent that the utility of the lowest and highest would be normalized and only one utility for £20 need be estimated.

In one treatment there were very few balls, with the intention that this be more like a standard risky process. There were 2 pink balls, 5 blue balls, and 3 yellow balls, and this was deemed countable in the time available. In another treatment the process was pushed to the edge, such that one could probably count the number of pink and yellow balls, but not the number of blue balls: 4 pink, 10 blue and 6 yellow. Finally, adding more balls of each color made it practically impossible for subjects to determine the exact probability by counting: 8 pint, 20 blue, and 12 yellow. These treatments were applied on a between-subjects basis.

Each subject responded to 162 paired lotteries, reflecting an exhaustive (and exhausting) evaluation of all possible events defined over 3 colors and 3 prizes that survive elimination by first-order stochastic dominance.[30]

A variety of preference functionals are estimated, including SEU and a Choquet Expected Utility (CEU) specification which is identical in this setting to the Rank-Dependent Utility (RDU) model of Quiggin [1982]. The approach is also non-parametric, in the sense that no functional forms are assumed for the utility function or the probability weighting functions; the approach is stochastically parametric. To illustrate, under SEU one estimates a value for u(£20), and values for the subjective probabilities of the blue and yellow balls emerging. The implied constraint on these probabilities, of course, is that the subjective probability of a pink ball emerging is one minus the

---

[30] Although subjective, our experience in our lab is that 60 binary choices on a computer is about the maximum we can ask without considerable boredom setting in. Of course, we do not have a British Bingo Blower in the background to liven things up!

two estimated probabilities.

Under CEU one estimates the decision weights on each outcome that are generated by each possible event. Since there are 6 possible events, there are 6 weights to be estimated, along with u(£20). Critically, for the estimation of ambiguity aversion, there is no presumption that the decision weights sum to one in any sense, and actually there are 5 ways that they could sum to one (p. 15). Under SEU these decision weights should sum to 1, no matter how one defines the summation (i.e., in any of the 5 ways identified).

This is where the three treatments enter. One can agree that the first treatment entails no ambiguity or uncertainty, and that the final treatment entails considerable ambiguity and uncertainty. The middle treatment is in between, in an attractive sense. So non-additivity reflects attitudes to risk in the first treatment, and attitudes towards risk and ambiguity in the final treatment. If there is a difference in the estimated measure of non-additivity between the two treatments, one can make inferences about the degree of ambiguity aversion or ambiguity loving.

Unfortunately, the estimates (Table 6, p. 25) only classify subjects as being ambiguity-averse or ambiguity-loving, noting that none were exactly ambiguity-neutral. But the critical hypothesis test here must account for standard errors on these estimates, or else we cannot say that these are statistically significant violations of additivity. Strangely, given the care of the experimental design and statistical specification, these statistics are absent. In addition, one would want to know if there are more violators of additivity in the final treatment compared to the first, but without standard errors one cannot say. Using the head counts based on point estimates, and with the prior that subjects would exhibit some non-neutral attitude to ambiguity, there is actually a disappointing similarity in the indices of non-additivity across these two treatments; but this is hard to interpret meaningfully without the standard errors.

Hey, Lotito and Maffioletti [2008] expand the analysis of the same experiment to include additional preference functionals, but do not arrive at easily interpreted measures of ambiguity or uncertainty aversion. Nor was this their inferential objective.

Additional References

Andersen, Steffen; Harrison, Glenn W.; Lau, Morten Igel, and Rutström, E. Elisabet, "Elicitation Using Multiple Price Lists," *Experimental Economics*, 9(4), December 2006, 383-405.

Andersen, Steffen; Fountain, John; Harrison, Glenn W., and Rutström, E. Elisabet, "Estimating Subjective Probabilities," *Working Paper 09-01*, Department of Economics, College of Business Administration, University of Central Florida, 2009.

Baillon, Aurélien, "Eliciting Subjective Probabilities Through Exchangeable Events: An Advantage and a Limitation," *Decision Analysis*, 5(2), June 2008, 76-87.

Bossaerts, Peter; Ghirardato, Paolo; Guarneschelli, Serena, and Zame, William R., "Ambiguity in Asset Markets: Theory and Experiment," *Working Paper*, California Institute of Technology, July 2007.

Chakravarty, Sujoy, and Roy, Jaideep, "Recursive Expected Utility and the Separation of Attitudes Towards Risk and Ambiguity: An Experimental Study," *Theory and Decision*, 66, 2009, 199-228.

Chen, Yan; Katuščák, Peter, and Ozdenorem, Emre, "Sealed Bid Auctions with Ambiguity: Theory and Experiments," *Journal of Economic Theory*, 136, 2007, 513-535.

Chew, Soo Hong, and Sagi, Jacob, "Small Worlds: Modeling Attitudes Toward Sources of Uncertainty," *Journal of Economic Theory*, 139, 2008, 1-24.

Davidson, Russell, and MacKinnon, James G., *Estimation and Inference in Econometrics* (New York: Oxford University Press, 1993).

Halevy, Yoram, "Ellsberg Revisited: An Experimental Study," *Econometrica*, 75, 2007, 503-536.

Machina, Mark J., and Schmeidler, David, "A More Robust Definition of Subjective Probability," *Econometrica*, 60(4), July 1992, 745-780.

Offerman, Theo; Sonnemans, Joep; van de Kuilen, Gijs, and Wakker, Peter P., "A Truth-Serum for Non-Bayesians: Correcting Proper Scoring Rules for Risk Attitudes," *Review of Economic Studies*, 76 2009, forthcoming.

Quiggin, John, "A Theory of Anticipated Utility," *Journal of Economic Behavior & Organization*, 3(4), 1982, 323-343.