

Risk Aversion and Uncertainty Aversion

Glenn W. Harrison

RMI Brownbag Presentation

The brownbag will cover the theoretical models and discussion about estimation that is developed in section 1 of the attached paper, and not deal much with specific estimation issues or experimental issues covered in later sections. In particular, I would like to focus on the potential of the simple “identification strategy” noted in footnote 4.

There are some lengthy appendices to the paper, discussing experimental instructions and previous experimental literature, which are not likely to be of direct interest. I will be happy to send those to anyone interested privately.

Estimating Aversion to Uncertainty

by

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

September 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); Department of Risk Management & Insurance and Center for the Economic Analysis of Risk, Robinson College of Business, Georgia State University, USA (Harrison); and Department of Economics, College of Business Administration, University of Central Florida, USA (Rutström). E-mail: sa.eco@cbs.dk, john.fountain@canterbury.ac.nz, gharrison@gsu.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}, \dots, U_{ij}) and (C_{i1}, \dots, 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 $\boldsymbol{\rho} = (\rho_1, \dots, \rho_I)$. 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 $\boldsymbol{\pi}^u = (\pi_{i1}^u, \dots, \pi_{ij}^u)$. 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 $\boldsymbol{\pi}^c = (\pi_{i1}^c, \dots, \pi_{ik}^c)$ 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 $(\sigma_{ij1}, \dots, \sigma_{ijk})$ denote the subjective conditional probability of events U_{ij} and C_k occurring under prior i, so $\sigma_{ijk} \equiv \pi_{ij}^U \pi_k^C$. The overall evaluation of an act z is then accomplished by evaluating over all I priors in the natural manner:

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

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 $\alpha < 1$ (> 1) implies risk aversion (loving), $\beta < 1$ (> 1) implies uncertainty aversion (loving), $\alpha = 1$ implies risk neutrality, and $\beta = 1$ implies uncertainty neutrality. When there is uncertainty neutrality (1) behaves just like a conventional SEU characterization. Assume $\alpha = 0.5$, consistent with modest risk aversion, and $\beta = 0.9$, 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 $\Lambda_{1jk} = \sum_{j=1,J} \sum_{k=1,K} \sigma_{1jk} v(z_{jk}) = 0.25 \times 100^\alpha + 0.25 \times 100^\alpha + 0.25 \times 0^\alpha + 0.25 \times 0^\alpha = 5$. 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

$$W(x) = 0.6 \times 5^\beta + 0.2 \times 5^\beta + 0.2 \times 5^\beta = 4.26 \quad (4)$$

$$W(x') = 0.6 \times 5^\beta + 0.2 \times 5^\beta + 0.2 \times 5^\beta = 4.26 \quad (4')$$

$$W(y) = 0.6 \times 5^\beta + 0.2 \times 4^\beta + 0.2 \times 6^\beta = 4.25 \quad (5)$$

$$W(y') = 0.6 \times 5^\beta + 0.2 \times 6^\beta + 0.2 \times 4^\beta = 4.25 \quad (5')$$

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 (ρ_1, \dots, ρ_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

$$W(z) = u[\sum_{j=1,J} \sum_{k=1,K} \pi_j^u \pi_k^c v(z_{jk})] \quad (6)$$

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

$$W(z) = \sum_{j=1,J} \sum_{k=1,K} \pi_j^u \pi_k^c v(z_{jk}) \quad (7)$$

Of course, (7) is the same as (6) when the $u(\cdot)$ 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_j^u , so that (1) becomes

$$W(z) = \sum_{i=1,I} \rho_i u[\sum_{k=1,K} \pi_i^u \pi_k^c v(z_{jk})] \quad (8)$$

Then set $\pi_i^u = \{1, 0\}$ for each i , and we get

$$W(z) = \sum_{i=1,I} \rho_i u[\sum_{k=1,K} \pi_k^c v(z_{jk})] \quad (9)$$

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_{ij}^u 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 \mathbf{p} 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 $\rho_1 = \rho_2 = 1/2$. 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_i 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_{ij}^u 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 $\pi_{11}^u=1$, $\pi_{21}^u=0$ 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 $1/2$ of the subjective distribution is attached to each distinct estimated probability, so $\rho_1 = \rho_2 = 1/2$ by assumption, and estimate two subjective probability values π_{11}^u and π_{21}^u .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_{ij}^u and π_{2j}^u 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.

$\pi_{1j}^u=0.42$ and $\pi_{2j}^u=1$. Alternatively, one could assume κ specific subjective probability values π_{ij}^u , 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 ($\alpha=0.5$ and $\beta=0.9$). 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 $\alpha=\beta=1$, 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 $\beta=1$ while $\alpha=0.5$, 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

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

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

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 :

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

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

$$\nabla EU = eu_R / (eu_L + eu_R) \quad (12)$$

calculated, where

$$\begin{aligned} eu_R &= \exp(EU_R) \\ eu_L &= \exp(EU_L) \end{aligned} \quad (13)$$

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 indifference¹⁴ the log-likelihood is then

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

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

$$eu_i = \exp[(EU_i/v)/\mu], \quad (13'')$$

instead of (13) and (13'), where v is a normalizing term for each lottery pair L and R, and $\mu > 0$ is a structural “noise parameter” used to allow some errors from the perspective of the deterministic EUT model. The normalizing term v is defined as the maximum utility over all prizes in this lottery pair minus the minimum utility over all prizes in this lottery pair, and ensures that the *normalized* EU difference $[(EU_R - EU_L)/v]$ remains in the unit interval. As $\mu \rightarrow \infty$ this specification collapses ∇EU to 0 for any values of EU_R and EU_L , so the probability of either choice converges to $1/2$. So a larger μ means that the difference in the EU of the two lotteries, conditional on the estimate of r 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 $\ln L(r, \alpha, \mu; y, \mathbf{X})$ by estimating r , α and μ , given observations on y and \mathbf{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 $\nabla EU < 0$, anywhere between 0 and 1 for $\nabla EU = 0$, and 1 for all values of $\nabla EU > 0$. Harrison [2008; p.326] illustrates the implied CDF, referring to it as the CDF of a “Hardnose Theorist.”

¹⁶ The normalizing term v is given by the values of r , α and the lottery parameters, and the latter parameters are part of \mathbf{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

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

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

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

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

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

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

$$eu_{\theta} = \exp[(EU_{\theta}/v)/\mu] \quad (16)$$

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

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

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 $\{0\%, 1\%, 2\%, \dots, 100\% \}$ 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

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

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

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

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 $\rho_1 = \rho_2 = 1/2$. In this case we would estimate two subjective probability values π_{1j}^u and π_{2j}^u , 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 $1/2$, so ρ_2 is thereby constrained to be $1/2$ since $I=\kappa=2$, and estimate the two subjective probability values. We evaluate the expected utility of a report θ as

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

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

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

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:

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

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 observe 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 ($\pi_{11}^u=1$, $\pi_{21}^u=0$), 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 $\pi_{11}^u=0.19$ and $\pi_{23}^u=0.99$, 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




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




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	<i>p</i> -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
GPA _{low}	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
GPA _{low}	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
GPA _{low}	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
GPA _{low}	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	<i>p</i> -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
	GPA _{low} 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
	GPA _{low} 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.