

Inferring Beliefs as Subjectively Uncertain Probabilities

by

Steffen Andersen, John Fountain, Glenn W. Harrison, Arne Risa Hole and E. Elisabet Rutström[†]

September 2010

Abstract. We propose a method for estimating subjective beliefs, viewed as a subjective probability distribution. The key insight is to *characterize beliefs as a parameter to be estimated from observed choices in a well-defined experimental task, and to estimate that parameter as a random coefficient.* The experimental task consists of a series of standard lottery choices in which the subject is assumed to use conventional risk attitudes to select one lottery or the other, and then a series of betting choices in which the subject is presented with a range of bookies offering odds on the outcome of some event that the subject has a belief over. Knowledge of the risk attitudes of subjects conditions the inferences about subjective beliefs. Maximum simulated likelihood methods are used to estimate a structural model in which subjects employ subjective beliefs to make bets. We present evidence that some subjective probabilities are indeed best characterized as probability distributions with non-zero variance.

[†] Department of Economics, Copenhagen Business School, Copenhagen, Denmark (Andersen); Department of Economics, University of Canterbury, Christchurch, New Zealand (Fountain); Department of Risk Management & Insurance and Center for the Economic Analysis of Risk, Robinson College of Business, Georgia State University, USA (Harrison); Department of Economics, University of Sheffield, UK (Hole); and Robinson College of Business and Department of Economics, Andrew Young School of Policy Studies, Georgia State University, USA (Rutström). E-mail contacts: sa.eco@cbs.dk, john.fountain@canterbury.ac.nz, gharrison@gsu.edu, a.r.hole@sheffield.ac.uk and erutstrom@gsu.edu. We thank the U.S. National Science Foundation for research support under grants NSF/HSD 0527675 and NSF/SES 0616746.

Subjective probabilities are widely viewed as being less certain than objective probabilities. This view underlies theories of choice where the lack of precision in subjective probabilities affects preferences directly, due to so-called ambiguity aversion or uncertainty aversion. In part it derives from the intuition that subjectively perceived chances of an event occurring might suffer from some lack of precision, and that this should matter for behavior.

In traditional expected utility models of choice behavior the uncertainty about the subjective probability distribution does not affect the observed choice behavior for an individual as long as the mean is unaffected. In other words, if one person has a mean-preserving change in their subjective beliefs, they would have no reason to change their observed choices when betting on the eventual outcome.

To account for preferences over ambiguity, and to obtain potentially refutable hypotheses about how the degree of uncertainty in beliefs affects behavior, one therefore has to relax the traditional model. Virtually all of the extensions of the expected utility model to date involve allowing for the subjective probability of an event to be characterized as being uncertain, as if it were one draw from a subjective probability distribution. Hence one talks about subjective beliefs, rather than a single subjective probability. The subjective beliefs are therefore conceived of as a non-degenerate probability distribution, and not just a scalar. One can then consider a range of non-traditional models that admit of refutable changes in observable behavior (e.g., Segal [1987], Gilboa and Schmeidler [1989], Klibanoff, Marinacci and Mukerji [2005] and Nau [2006]).¹

We propose a method for estimating subjective beliefs, viewed as a subjective probability distribution. The key insight is to *characterize beliefs as a set of distributional parameters to be estimated from observed choices in a well-defined experimental task, and to estimate those parameters as random coefficients*. The experimental task consists of a series of standard lottery choices defined over objective probabilities, and then a series of betting choices in which the subject is presented with a range of bookies

¹ The design of experiments and econometric methods to test these alternatives is a delicate matter, but considerable progress has been made (e.g., Ahn, Choi, Gale and Kariv [2009], Halevy [2007] and Hey, Lotito and Maffioletti [2007][2008]).

offering odds on some outcome that the subject has a belief over. The event we focus on is a draw from a humble bingo cage, populated with 60 balls that are orange and white. The subject knows that there are 60, and that there are only orange and white balls. But the cage is only visible for about 10 seconds, and is rotating all the time to make it (practically) impossible for the subject to count the number of orange balls. Thus it is unlikely that the subjective belief will match the objective probability, or be an exact subjective probability.

The “random coefficients” approach is commonly used in econometrics to capture unobserved individual heterogeneity, but that is just one interpretation of the estimates. To understand the basic idea, assume that we were to estimate risk attitudes in a population, using a sample of choices over standard lotteries (a common inferential task in experimental research). One way to characterize individual heterogeneity in risk attitudes would be to assume that the risk aversion coefficient was a linear function of observables. This is one way to model the observed individual heterogeneity in behavior. Alternatively, and the approach adopted here, one might assume that the coefficient was normally distributed across the sample.² This is a way of modeling the unobserved individual heterogeneity in behavior. In effect, in the traditional interpretation that views all variance as being caused by observed or unobserved exogenous variables, each subject would be assumed to have some true and precise coefficient value, but these values are viewed as being distributed across the sample in a way that can be characterized by a normal distribution. One then estimates *hyper-parameters* to reflect this population distribution.³

For example, if one assumes a normal distribution, the hyper-parameters to be estimated would be a mean (population) risk aversion coefficient and a standard deviation in the (population) risk aversion coefficient. Each *hyper-parameter* would be estimated by a point estimate and a standard

² There is no necessary tension between the two approaches to characterize heterogeneity, although in practice one tends to see data evaluated using one or the other method. We employ them in a complementary manner, as illustrated later.

³ The terminology across the econometric literature is not standard, so the expression “hyper-parameter” can have other meanings. Here it refers to the parameters that characterize the estimated distribution of the behavioral parameter of interest.

error.⁴ As the sample size gets larger, one would expect the estimated standard *errors* of the *hyper-parameters* to shrink, for the usual reasons, but one would not necessarily expect the standard *deviation* of the population *parameter* to shrink. This approach generalizes naturally to non-normal distributions for the population parameter, and to multi-variate distributions where there are several population parameters.

Our approach then is to estimate subjective beliefs as a normal distribution in the population. There are two, equivalent ways that one could interpret these estimates, which we discuss below. For the moment, to focus on essentials, assume that we have a representative decision maker in the population with beliefs about the event that are normally distributed.

In section 1 we review the “mixed logit” specification of random utility models, in which the latent index is assumed to be a linear function of observable characteristics of the choice and/or individual. One or more of the parameters is then estimated as a random coefficient, by estimating a set of hyper-parameters that characterize the distribution of the coefficient in the population. Our exposition is brief, reflecting the availability of excellent treatments by Train [2003] and others, and sets the stage for an extension to allow for the latent index to be a non-linear function of characteristics. This extension is needed since virtually all interesting functional forms for utility and subjective probabilities involve non-linear functions, but virtually all applications of random utility models assume a linear or linearized approximation of the true utility function.⁵

In section 2 we consider our experimental design and the raw data. The lottery choice task is a replication of the classic study of individual choice under risk due to Hey and Orme [1994]. The betting task we implement is one of the experimental procedures for eliciting subjective beliefs proposed by Andersen, Fountain, Harrison and Rutström [2010], but building on an old and venerable literature such as de Finetti [1937][1970], Savage [1971] and Epstein [1977; p. 298ff.].

⁴ To be pedantic, there are then four estimates: the point estimate and standard error of the mean of the population parameter, and the point estimate and standard error of the standard deviation of the population parameter.

⁵ This extension to non-linear mixed logit, while formally modest, is also likely to be of considerable value more generally, since it allows direct estimation of latent structural parameters of virtually any specification. Andersen, Harrison, Hole and Rutström [2010] document general software, developed within *Stata*, to estimate non-linear mixed logit models.

What is important here is that the second task be one that associates beliefs about some event with observable consequences, such as bets with monetary consequences based on the outcome of an event.

In section 3 we generate estimates of the subjective beliefs of our subjects, initially using methods that assume that there is no subjective uncertainty and then methods that allow for some subjective uncertainty. We find evidence that there is some subjective uncertainty in the population.

1. Linear and Non-Linear Mixed Logit

To fix ideas, assume that we initially want to estimate a parameter to reflect risk attitudes in a conventional lottery choice task. Then we extend the analysis to consider the estimation of subjective probabilities, and finally to subjective beliefs.

A. Non-Linear Mixed Logit for Lottery Choices

Assume a sample of N subjects making choices over J lotteries in T experimental tasks.⁶ In all of the applications we consider, $J=2$ since the subjects are making choices over two lotteries, but there are many designs in which the subject is asked to make choices over $J>2$ lotteries (e.g., Binswanger [1981], Eckel and Grossman [2002]). In the traditional mixed logit literature one can view the individual n as deriving utility Δ from alternative j in task t , given by

$$\Delta_{njt} = \beta_n x_{njt} + \epsilon_{njt} \quad (1)$$

where β_n is a vector of coefficients specific to subject n , x_{njt} is a vector of observed attributes of individual n and/or alternative j in task t , and ϵ_{njt} is a random term that is assumed to be an identically and independently distributed extreme value. We use the symbol Δ for utility in (1), since we will need to generalize to allow for non-linear utility functions, and expected utility functionals, and prefer to think of (1) as defining a general, latent index rather than as specifically utility. In our experience, this purely semantic difference avoids some confusions about interpretation. Each β_n is

⁶ It is trivial to allow J and T to vary with the individual, but for ease of notation we omit that generality.

traditionally estimated as a distribution based on 2 or more hyper-parameters. In the case of a normal distribution they would be the mean and standard deviation of the population distribution.

Specifically, for our purposes we need to extend (1) to allow for non-linear functions H defined over β and the values of x , such as

$$\Delta_{njt} = H(\beta_n, x_{njt}) + \varepsilon_{njt} \quad (2)$$

For example, x might consist of the vector of monetary prizes m_k and probabilities p_k , for outcome k of K in a given lottery, and we might assume an Expo-Power (EP) utility function originally proposed by Saha [1993]. Following Holt and Laury [2002], the EP function is defined as

$$U(m_k) = [1 - \exp(-\alpha m_k^{1-\rho})] / \alpha, \quad (3)$$

where α and ρ are parameters to be jointly estimated. The EP function can exhibit increasing or decreasing relative risk aversion (RRA), depending on the parameter α : RRA is defined by $\rho + \alpha(1 - \rho)m_k^{1-\rho}$, so RRA varies with income if $\alpha \neq 0$ and the estimate of ρ defines RRA at a zero income.

This function nests CRRA (as $\alpha \rightarrow 0$) and CARA (as $\rho \rightarrow 0$). Under expected utility theory (EUT) the probabilities for each outcome are those that are induced by the experimenter.⁷ Hence expected utility is simply the probability weighted utility of each outcome in each lottery j :

$$EU_j = \sum_k [p_k \times U(m_k)] \quad (4)$$

If we let $\beta = \{\alpha, \rho\}$ here, so that we would estimate the hyper-parameters of the distribution over α and ρ , our two risk preferences, we will want to let $H(\beta_n, x_{njt})$ be defined by $H(\alpha_n, \rho_n, m_{njt}, p_{njt})$ using (3) and (4). It is then possible to evaluate the latent index Δ in (2).⁸

⁷ This approach generalizes immediately to non-EUT models in which there are more parameters, say to account for probability weighting and loss aversion. It also generalizes to non-CRRA specifications within EUT models that allow for more flexible specifications of risk attitudes that might vary with the level of the prizes.

⁸ The extension from a linear mixed logit specification, assuming (1), to a non-linear mixed logit specification, assuming (2), has an attractive side-benefit when it comes to identifying the effects of demographic variables such as sex. In the usual specification with linear latent indices of utility the effects of attribute-invariant effects drop out, and one can only consider them by considering interactions with attributes. In effect, the non-linearity of (2) builds this interaction in at a structural theoretical level. The intuition derives from the fact that only differences in (expected) utility matter for choice. Thus one can re-normalize (expected) utility more or less at will, with minor mathematical constraints, as long as the (expected) utility numbers have the same ordering. Then, as Train [2003; p.25] notes: "The same issue affects the way that socio-demographic variables enter a model. Attributes of the alternatives, such as the time and cost of travel on different modes, generally vary over alternatives. However, attributes of the decision maker do not vary over alternatives. They can only enter the model if they are specified in ways that create differences in utility over alternatives." If the sex of the agent affects the risk attitude, then this characteristics will affect the (expected) utility evaluation of a given lottery, since each lottery will typically have attributes given by probabilities and outcomes that vary across the two

The population density for β is denoted $f(\beta | \theta)$, where θ is a vector defining what we refer to as the hyper-parameters of the distribution of β . Thus individual realizations of β , such as β_n , are distributed according to some density function f . For example, if f is a Normal density then θ_1 would be the mean of that density and θ_2 the standard deviation of that density, and we would estimate the hyper-parameters θ_1 and θ_2 . Or f could be a Uniform density and θ_1 would be the lower bound and θ_2 would be the upper bound. If β consisted of more than two parameters, as it does in the case of an EP utility function (3), then θ might also include terms representing the covariance of those parameters.

Conditional on β_n , the probability that the subject n chooses alternative i in task t is then given by the conditional logit formula, modestly extended to allow our non-linear index

$$L_{nit}(\beta_n) = \exp\{H(\beta_n, x_{nit})\} / \sum_j \exp\{H(\beta_n, x_{njt})\} \quad (5)$$

The probability of the observed choices by subject n , over all tasks T , again conditional on knowing β_n , is given by

$$P_n(\beta_n) = \prod_t L_{n i(n,t)t}(\beta_n) \quad (6)$$

where $i(n,t)$ denotes the lottery chosen by subject n in task t , following the notation of Revelt and Train [1998]. The unconditional probability involves integrating over the distribution of β :

$$P_n(\theta) = \int P_n(\beta_n) f(\beta | \theta) d\beta \quad (7)$$

and is therefore the weighted average of a product of logit formulas evaluated at different values of β , with the weights given by the density f .

We can then define the log-likelihood by

$$LL(\theta) = \sum_n \ln P_n(\theta) \quad (8)$$

and approximate it numerically using simulation methods, since it cannot be solved analytically.

Using the methods of Maximum Simulated Likelihood (MSL) reviewed in Train [2003; §6.6, ch.10]

alternatives presented to the subject in any given choice. In effect, the non-linear specification (2) naturally builds in the effect that characteristics have on utility differences. Andersen, Harrison, Hole and Rutström [2010] provide a more formal statement of this result.

and Cameron and Trivedi [2005; ch.12], we define the simulated log-likelihood by taking $r=1, R$ replications β^r from the density $f(\beta | \theta)$:

$$SLL(\theta) = \sum_n \ln \{ \sum_r P_n(\beta^r) / R \} \quad (9)$$

The core insight of MSL is to evaluate the likelihood conditional on a randomly drawn β^r , do that R times, and then simply take the unweighted average over all R likelihoods so evaluated. The average is unweighted since each replication r is equally likely, by design. If R is “large enough,” then MSL converges, under modest assumptions, to the Maximum Likelihood (ML) estimator.⁹

It is a simple matter to extend this approach to include non-random coefficients as well. In this case the hyper-parameter of the coefficient consists of just one scalar. To use the earlier example, if f is assumed to be a Normal density then θ_1 would be the mean of that density and θ_2 the standard deviation of that density, and we would estimate θ_1 and simply constrain $\theta_2=0$.

B. Estimating a Subjective Probability

Having demonstrated the method of MSL based on estimating risk preference parameters, we continue with a discussion of the parameters of core interest to us, the subjective beliefs. We start by assuming away uncertainty and discuss the estimation of subjective probabilities.

Assume that there are two outcomes, A and B, that exhaust the set of outcomes. In our case A would be an orange ball being selected from our experimental bingo cage, and B would be a white ball being selected. The subject that selects event or outcome A from a given bookie b receives EU

$$EU_A = \pi_A \times U(\text{payout if A occurs} \mid \text{bet on A}) + (1-\pi_A) \times U(\text{payout if B occurs} \mid \text{bet on A}) \quad (10)$$

where π_A is the subjective probability that A will occur. The payouts that enter the utility function are defined by the odds that each bookie offers, and are set by the experimenter. We discuss them in

⁹ An important practical consideration with MSL is the manner in which replicates are drawn, and the size of R that is practically needed. We employ Halton draws to provide better coverage of the density than typical uniform number generators: see Train [2003; ch.9] for an exposition, and Drukker and Gates [2006] for the numerical implementation we employ. Our computational implementation generalizes the linear mixed logit program developed for *Stata* by Hole [2007], and is documented in Andersen, Harrison, Hole and Rutström [2010].

the next section, but for now assume that they vary from bookie to bookie. For the bet offered by a bookie offering 10:1 odds on event A, for example, and assuming no “house take” by the bookie, these payouts are \$10 and \$0 for every 1\$ that is bet, so we have

$$EU_A = \pi_A \times U(\$10) + (1-\pi_A) \times U(\$0) \quad (10')$$

We similarly define the EU received from a bet on event B as:

$$EU_B = \pi_A \times U(\text{payout if A occurs} \mid \text{bet on B}) + (1-\pi_A) \times U(\text{payout if B occurs} \mid \text{bet on B}). \quad (11)$$

and this translates for the first bookie, offering 10:1 odds on A and (10/9):1 odds on B, into payouts of \$0 and \$1.11, so we have

$$EU_B = \pi_A \times U(\$0) + (1-\pi_A) \times U(\$1.11) \quad (11')$$

for this particular bookie and bet. We observe the bet made by the subject, so we can calculate the likelihood of that choice given values of ρ , α and π_A .

We need ρ and α to evaluate the utility function in (10) and (11), and we need π_A to calculate the EU in (10) and (11) once we know the utility values, and hence the latent indices (5) that generate the likelihood of observing the choice of event A or event B. 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 allow us to identify ρ and α under EUT, since π_A plays no direct role in explaining the choices in that task.¹⁰

In practical terms the only difference between the likelihood contribution from the lottery tasks and the betting tasks is how the probability of the outcome enters each problem. For the lottery tasks this is given as data, and for the betting tasks this is estimated as a parameter. Either way, once the EU is defined, the likelihood of the observed choices are evaluated identically.

¹⁰ The value of joint estimation, particularly when paired with an experimental design with multiple tasks to help identification, is discussed in more general terms in Harrison and Rutström [2008], Andersen, Harrison, Lau and Rutström [2008], and Andersen, Fountain, Harrison and Rutström [2010].

C. Estimating a Subjective Belief

The extension to estimate a subjective belief is immediate, and involves setting $\beta = (\alpha, \rho, \pi_A)$ and treating the risk aversion coefficients α and ρ as non-random while the subjective belief π_A is treated as a random coefficient. We eventually treat both as random coefficients, but prefer to focus attention here on the subjective belief alone.

There are two possible interpretations of the estimates obtained when one allows the subjective probability to be estimated as a random coefficient. These interpretations are fundamental to the validity of our approach.

One interpretation is that each subject picks one specific subjective *probability* value from that *distribution* when they place their bets: it is as if they draw a specific subjective probability from the population distribution “urn” and then use that probability draw to place their bets. In effect, one subject sees the turning bingo cage and thinks he sees 15% orange balls, another subject sees the same cage and thinks he sees 20%, and another subject sees the same cage and thinks he sees 25%; in each case the subjects can be viewed as making draws from a single population distribution. Each subject then acts as if there is no uncertainty around those individual perceptions when making their bets. This interpretation is exactly the one that the traditional econometric literature makes when estimating random coefficients, adapted for our task of estimating subjective beliefs. From a behavioral viewpoint, at the time of placing the bet there is no uncertainty from the perspective of the subject.

Another interpretation is that our subjects each carry around with them a personal subjective (posterior) probability distribution after they see the spinning bingo cage, and access that distribution when making their bets. This distribution can be assumed to be a replica of the population distribution. Under this interpretation the estimated population distribution is in fact the distribution used by *each* subject. Behaviorally, each subject is viewed as facing uncertainty when placing the bet.

If we adopt the first interpretation, then the belief π_A that we recover is actually consistent

with many supporting belief distributions *at the level of the individual*, and is the mean of those distributions. To see this well-known result, initially assume that we elicit the belief $\pi_A=0.65$, and that the subject has a degenerate belief distribution with all mass at that point. Under EUT the subject has the EU for a bet on A occurring given by (10), and the EU for a bet on B occurring given by (11). Now assume that the subject actually had a 2-point distribution with density $f(\pi_A^1)=1/2$ at 0.60 and $f(\pi_A^2)=1/2$ at 0.70. Define EU_A^1 by substituting π_A^1 for π_A in (10), define EU_A^2 by substituting π_A^2 for π_A in (10), and similarly define EU_B^1 and EU_B^2 by corresponding substitutions in (11). Then the EU for a bet on A is now the compound lottery

$$EU_A = [f(\pi_A^1) \times EU_A^1] + [f(\pi_A^2) \times EU_A^2] \quad (12)$$

and the EU for a bet on B is similarly defined. Since the outcomes in the conditional lotteries are the same, one can collect terms and see that (12) is identical to

$$\mu(\pi_A) \times U(\text{payout if A occurs} \mid \text{bet on A}) + (1-\mu(\pi_A)) \times U(\text{payout if B occurs} \mid \text{bet on A}) \quad (13)$$

where $\mu(\pi_A) \equiv [f(\pi_A^1) \times \pi_A^1] + [f(\pi_A^2) \times \pi_A^2]$, the mean of the 2-point distribution. Thus the choice behavior of an individual with a 2-point probability distribution is identical to the choice behavior of an individual with a 1-point probability distribution where that one point is given by the mean of the 2-point distribution for that individual. This does not mean that the individual *has* a 1-point distribution, just that we cannot use his choice behavior to say whether he has a 1-point or 2-point distribution.

Thus we can only claim, at least under EUT, that we elicit the mean belief for an individual. We cannot identify the underlying distribution of beliefs for an individual *based solely on the choice of that individual*, nor can we rule out the possibility that the individual has a non-degenerate probability distribution. It is immediate that this result generalizes to asymmetric distributions with more than 2 mass points, and indeed to continuous distributions. The same result generalizes to RDU¹¹ if one

¹¹ Of course, in the case of RDU it is the average *weighted* probability that is elicited, and an additional step is needed to recover the subjective probability itself.

maintains the Reduction of Compound Lottery axiom, as is common.¹²

The upshot is that one has to look to choices *across individuals* to be able to identify the underlying distribution of beliefs. Ideally this would involve different individuals making comparable choices in response to the same physical stimuli, and where there is no reason to expect different individuals to have different priors about the true outcome. This is exactly the domain in which controlled laboratory experiments can play a critical role, and the design we propose implements this ideal environment.

D. Bounded Beliefs

One essential characteristic of beliefs, if they are to have the interpretation of uncertain probabilities, is that the distribution have a domain that is bounded on the unit interval.¹³ Unfortunately, Normal distributions do not have this property: even if the mean is constrained to be in the unit interval (e.g., using some common non-linear transformation), there is no easy way to constrain the standard deviation to ensure this property. Of course, one might be “lucky” and generate estimates that are practically bounded in the unit interval in terms of non-negligible densities, but in general one wants a more elegant solution to this problem.

One attractive option is to employ a transformation of the Normal distribution known as the Logit-Normal (L-N) distribution. Originally proposed by Aitchison and Begg [1976; p.3] as an excellent, tractable approximation to the Beta distribution, it has been resurrected by Lasaffre, Rizopoulos and Tsonaka [2007]. One nice property of the L-N distribution is that MSL algorithms developed for univariate or multivariate Normal distributions can be applied directly, providing one allows non-linear transformations of the structural parameters, which is exactly what we need to do

¹² The axiom of EUT that is typically relaxed is the Independence Axiom, and the Reduction of Compound Lotteries axiom is almost always retained (e.g., Quiggin [1993; p. 19, 134, 154]). It also has considerable normative appeal in a-temporal settings, such as we have in mind here. However, Segal [1987][1990] illustrates some implications of relaxing the Reduction of Compound Lotteries axiom for RDU. To reiterate, we are focusing solely on probabilistically sophisticated versions of RDU in which the subjective belief can be recovered.

¹³ Smith [1969] provides a particularly eloquent explanation of this elementary point, which was initially forgotten in the debates over how to account for the Ellsberg paradox. Of course, subjective probabilities are not the same as decision weights and “Choquet capacities,” which were applied to this problem well after Smith [1969].

anyway to model subjective beliefs as bounded in the unit interval.

Figures 1 and 2 illustrate the wide array of distributional forms that are accommodated by the L-N distribution. The bi-modal distributions in the top right of Figure 1 and top left of Figure 2 are particularly attractive, since they reflect the maintained assumptions of several of the most tractable models of uncertainty aversion.

In the context of the betting task in our experiment, where there is uncertainty about the mix of orange and white balls in a bingo cage, the continuous distributions of Figures 1 and 2 might be used to approximate different discrete subjective belief distributions over the subjective probability π_A on the event A of an orange ball being drawn. The earlier discussion in §C provided an example of 1 and 2 point belief distributions over the subjective probability π_A , but one can imagine other belief distributions with positive probability support for any of the discrete points in the range of possible proportions of orange balls in the bingo cage. The distributions in Figure 1 can be used to approximate discrete belief distributions that reflect a symmetry in beliefs about which color is likely, while the distributions in Figure 2 can be used to approximate belief distributions expressing uncertainty over all possible subjective probabilities π_A , but where the chances of a ball being one color can be favored. L-N distributions can also be used to approximate extreme 2-point distributions about the proportion of orange balls in the cage, such as that there are either all orange or all white balls in the bingo cage.

Similarly, more extreme versions of the distributions illustrated in Figure 2 capture the “pessimism in priors” of the “maximin” characterization of Gilboa and Schmeidler [1989]. And the top right panel of Figure 1 captures the characterization of Ghirardoto, Maccheroni and Marinacci [2004] in which there is some weight on the most pessimistic prior and the most optimistic prior. The interim cases in Figures 1 and 2 allow one to capture more general characterizations, such as proposed by Klibanoff, Marinacci and Mukerji [2005].¹⁴

¹⁴ The expression “prior” in these models refers to one specific subjective probability for an event π_A from a set of such subjective probabilities for an event (e.g., Gilboa and Schmeidler [1989; p.142-3]).

We stress the flexibility of the L-N family because provides an attractive specification for efforts to estimate general models of attitudes towards uncertainty.

E. A Final Extension

One final extension is to allow for a “behavioral noise” term to affect choices. We adopt the contextual utility specification advocated by Wilcox [2008][2010] for this purpose, which adds one additional non-random coefficient μ to estimate. When the constraint $\mu=1$ is imposed there is no behavioral noise term in the specification.

2. Experimental Design and Sample

We recruited 97 subjects from the student population of the University of Central Florida in late October 2008 to participate in these experiments.

Figure 3 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 flipping a coin 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].

Figure 4 shows the betting interface presented to subjects, which was shown to subjects on a computer screen and in hand instructions. Prior to this interface, the subjects had been introduced to naturally occurring instances of this type of betting interface, on events ranging from the United States Presidential Election, the winner of the *American Idol* television program, and outcomes of the NBA basketball season. The interface in Figure 4 was then explained with these instructions:

First, for each betting task we will provide you with odds from 9 betting houses. Think of these as just 9 different physical or online locations where you can make your bet, each offering different odds. You need to decide what bet to place for each betting house. You do not need to make the same bet for each betting house.

Second, in this hypothetical example we are giving you a \$100 stake to bet with for each betting house in each of the betting tasks you will be presented with. The rules here are that you only get to place one bet in each betting house, and your stake for one house cannot be applied to another betting house. In the actual tasks today we will not be having stakes as high as \$100, but the stakes we provide are real.

Third, after you have placed your bets on all 9 betting houses, we will randomly choose one of the 9 betting houses to determine your actual payment. After you have finished placing all of your bets, we will come over and let you roll a die to decide which betting house we will actually use. So you are placing 9 bets in all, will play out one of them, picked at random, and you could win or lose on that bet.

The subjects were then taken through a worked example in which they were told how many orange and white balls were in the cage, and shown how to place bets. They were then taken through each possible outcome. The concept of “house probabilities” was also explained, as follows:

Beside each set of **house odds**, we also display the **house probability** that each event will occur. This is just another way of thinking about the odds offered by each betting house. Some people understand **house odds** better, and some people understand **house probabilities** better.

The **house probabilities** are just the inverse of the odds. So if the odds for a particular betting house say that you will be paid \$5 for every \$1 bet if the ball is Orange, as in betting House 2, then this implies a **house probability** for the Orange ball of $\$1 \div \$5 = 0.2$. This is the same thing as saying that betting House 2 believes there is a 20% chance of the ball being Orange. We have rounded some of the probabilities to make the screen easier to read, and you will be paid according to the odds.

The **house probabilities** might help you work out what is the best bet for you. For example, if you personally think the probability of an Orange outcome is lower than the house probability offered by a betting house, you might be inclined to bet against Orange with that house. (You bet “against” Orange by betting “for” White, of course). But if your **personal probability** of an Orange outcome is higher than the house probability offered by a betting house, you might be inclined to bet on Orange with that house. It is just as if you were placing a bet with a friend, because you disagree on the chances of something happening.

All subjects were able to place a series of bets in this practice round before the actual bets for real payment.

The actual betting tasks involved more draws from bingo cages with ping-pong balls. They were explained as follows:

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

We have a number of ping pong balls in each of three bingo cages, which we have labeled Cage A, Cage B and Cage C. Some of the ping pong balls are Orange

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

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

Figure 3 displays the sequence that was followed here. The cage was initially draped in an opaque towel, bearing the *Golden Knight* logo of the University of Central Florida. The cage was placed atop a tall table, so that all subjects had a clear view of it. A tall research assistant then took off the towel and turned the cage for roughly 10 seconds, timed by another assistant. Then the cage was covered again and subjects asked to place their bets. Once the bets for Cage A were completed, a similar sequence was followed for Cages B and C. At that point we selected a ball from each of Cages A, B and C. The final event for payment was decided at the end of the experimental session, and there was a 3-in-7 chance that one of these cages would be selected for payment.¹⁵

The distribution of orange balls within each session spanned a “low,” “medium” and “high” value, in random order across sessions. There were 6 physical stimuli used across all sessions. These stimuli used 6, 12, 30, 33, 45 and 48 orange balls out of a total of 60, implying true objective probabilities of 0.1, 0.2, 0.5, 0.55, 0.75 and 0.8, respectively.

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

¹⁵ Beliefs for completely unrelated events were elicited after the beliefs about the ping-pong ball draws. One event was the outcome of scores on a psychology test of the relative “empathy” of men and women, and three events related to the 2008 Presidential Election outcomes. Those events did not provide the control that the ping-pong ball events had, in terms of every subject being provided with exactly the same physical signals. We note, as anecdote, that we were astonished to see how excited the subjects were about the draws of ping-pong balls from these bingo cages. The nightlife in Orlando is not that bad.

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 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 considers the array of bookies, each with different odds, they had to place bets with. The set of bet choices is based on the subject’s beliefs about “the composition of the urn.”¹⁶

An additional feature of our design was that it allowed us to estimate the conventional risk attitudes of the sample, as we were jointly estimating subjective beliefs. Because those beliefs can only be identified in the context of some betting task of the kind we implemented,¹⁷ one has to know risk attitudes in order to separate out the effects of subjective beliefs and preferences over risk.

¹⁶ 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.

¹⁷ Virtually all other experimental procedures for evaluating choice under uncertainty, such as quadratic scoring rules and auctions for eliciting certainty-equivalents, amount to bets in one form or another. So this point is not specific to our design, of course.

3. Results

We initially estimate the subjective probability that a representative agent would have for each of the stimuli using maximum likelihood, and thereby assuming away any variance in the population distribution. This will allow us to see the “value added” from having the richer characterization provided by a random coefficients specification using maximum simulated likelihood. We then consider the interim random coefficients case in which we constrain the population mean to be equal to the objective probability for each stimulus, and only estimate the population standard deviation.¹⁸ This allows us to see the pure effect of allowing for some population standard deviation in the estimated distribution. We finally consider the random coefficients case in which we estimate the mean *and* standard deviation of the population process for each stimulus.¹⁹

A. Maximum Likelihood Estimates

Table 1 shows the estimated subjective probabilities obtained using maximum likelihood methods *and* the assumption that there is no population standard deviation in subjective beliefs. Underlying these estimates is evidence of a concave utility function, reflecting risk aversion for these subjects. The estimates are familiar from the literature with student subjects of this type (see Holt and Laury [2002] and Harrison and Rutström [2008]). The coefficient ρ is estimated to be 0.48 with a 95% confidence interval of [0.44, 0.52], and the coefficient α is estimated to be 0.083 with a 95% confidence interval of [0.06, 0.11]. So there is evidence of risk aversion at low levels of prizes, since $\rho > 0$, and there is evidence of increasing relative risk aversion as prizes increase up to \$100, since $\alpha > 0$. The same qualitative pattern emerges with random coefficients estimates of these parameters, reported in the next sections.

¹⁸ Even though we refer to estimating a mean and/or standard deviation, the resulting population is not Gaussian, since we employ the Logistic-Normal distribution.

¹⁹ Detailed estimation results, and all software and data to replicate our results using *Stata*, are available on request.

Turning to the estimates of the subjective probabilities, we find that there are some significant deviations in subjective probabilities from objective probabilities for the case in which there are 6 or 12 orange balls out of 60. However, for all other cases the estimated subjective probabilities are relatively close to the objective probability. In the case of 6 orange balls, our estimates “collapse” to suggest that subjects behaved as if there were none. In the case of 12 orange balls, our estimates suggest that the subjective probability is significantly *greater* than the objective probability.

These estimates reflect a further assumption that there is a single representative agent whose subjective probability we are estimating. It is a simple matter to extend this analysis to allow the core structural parameters to reflect observable covariates.²⁰ Thus we could estimate the subjective probability that women hold when we have 6 orange balls and contrast it to the subjective probability that men hold. Even if the same bets were placed by men and women in this case, any differences in their estimated utility functions could lead to differences in inferred subjective probabilities.

To illustrate, consider the 12-ball case. If we allow for the effects of sex and having a self-reported GPA greater than 3.75. The marginal effect on inferred subjective probabilities is for women to have a higher probability of 5.4 percentage points (95% confidence interval of [0.8, 10] percentage points) and those with a high GPA to have a higher probability of 0.1 percentage points (95% confidence interval of [-4.9, 5.1] percentage points). Since a male that does not have a high GPA has an estimated subjective probability of 0.36 (95% confidence interval of [0.33, 0.40]), we can infer four subjective probabilities:

- women with a high GPA have a subjective probability of 0.420;
- women without a high GPA have a subjective probability of 0.4219;
- men a high GPA have a subjective probability of 0.366; and
- men without a high GPA have a subjective probability of 0.365.

One could obviously extend this analysis to consider more covariates, and more interactions of

²⁰ All standard errors are corrected for “clustering” at the level of the individual, which captures heterogeneity to some extent.

covariates, as in Andersen, Fountain, Harrison and Rutström [2010].

It is worth noting that women have a significantly higher estimate of ρ , and those with a high GPA have a significantly higher estimate of α . These differences in utility functions need to be taken into account when inferring subjective probabilities. That is, one might see the same pattern of bookie choices for men and women, but if they have different utility functions one would draw different inferences about their implied subjective probabilities. The same principle of inference holds with the random coefficients specifications.

B. Random Coefficients Estimates, with the Mean Constrained to Equal the True Mean

Figure 6 collects results assuming that the average subjective belief is equal to the true probability of the urn, but that allows for there to be some subjective uncertainty around that average. This assumption maintains some degree of rational expectations, in the sense that there is some “wisdom of crowds” with respect to the average belief. It also illustrates our approach in stark form, since all that we estimate here is the population standard deviation in beliefs. All estimates were obtained using maximum simulated likelihood methods, using 100 Halton draws.

We find a familiar pattern of risk aversion, with modest levels of RRA for small prize levels and increasing RRA for higher prize levels, using the Expo-Power utility function (3). The estimate of $1-\rho$ is distributed as $N(0.551, 0.180)$, with the estimate of the population mean and population standard deviation each being statistically significantly greater than zero. Hence we reject the hypothesis that all subjects have the same risk attitude parameter ρ , consistent with the conventional view that risk preferences are heterogeneous across subjects. There is not so much heterogeneity in the parameter α , which is distributed as $N(0.022, 0.004)$, and where the estimate of the population standard deviation is not statistically different from zero. But the point estimate of the population mean of α is positive, and statistically significant, implying that RRA increases by 0.22 ($= 0.022 \times 100$, since the highest prize was \$100) over the prize domain [$\$0, \100]. Joint estimation of these risk

aversion parameters and the subjective belief distributions takes this increasing RRA into account.²¹

The estimates in Figure 6 are derived from 10,000 simulations of the L-N distribution parameter estimates. To illustrate, consider the case of 30 orange balls first. Here the L-N distribution is $\Lambda(N(0, 0.11))$. The Normal distribution inside this composite function has a mean of 0 and standard deviation of 0.11, and each of those point estimates has a standard error (indicating, incidentally, that each point estimate is significantly greater than 0). So one would take 10,000 draws from a normal distribution centered around 0 with that standard deviation, and then transform each such draw with the logistic transform. Since the logistic transforms values of 0 to be 0.5, it is not surprising that this distribution is centered in Figure 6 at 0.5. The L-N distribution has a standard deviation of 0.028.

Now consider the case of 33 balls, where the L-N distribution is $\Lambda(N(-0.2, 0.18))$. The Normal distribution inside this composite function has a lower mean than the 30-ball distribution, and a larger standard deviation. The logistic transforms negative values to be greater than 0.5, hence we see the L-N distribution for the 33-ball case to be centered around 0.55 with a standard deviation of 0.044. Thus it is the same as the 30-ball distribution apart from being shifted to the right and with a slightly larger standard deviation. Although the transformations are more extreme, of course, a similar logic explains the shape of the other L-N distributions in Figure 6 given the core parameter estimates.

Turning to the substance of the estimates, we see predictable symmetry for the stimuli centered near 0.5, and skewness for the stimuli closer to 0 or 1. Given the lack of skewness for the middle stimuli and the skewness for the extreme stimuli, the modes are either right on the objective probabilities or to one side of them. There is more uncertainty about the extreme stimuli, not surprisingly.

The distributions shown in Figure 6 could also be conditioned on observable covariates,

²¹ In other words, when bets with bookies offering one or two “low” stakes are evaluated, a smaller RRA is used than when evaluating bookies offering one or more higher stakes.

such as sex and GPA level. In this case the underlying coefficients would be linear functions of these covariates, so one would generate four distinct population distributions rather than just one for a given stimulus.

C. Unconstrained Random Coefficients Estimates

Figure 7 takes one more step in the estimation compared to Figure 6, and allows the mean of the population distribution to be unconstrained. The differences are striking.

For the 30-ball case we see an almost degenerate L-N distribution, with a mean of 0.54 and a standard deviation of 0.003, much smaller than the standard deviation for the corresponding distribution in Figure 6 where the mean was constrained to be 0.5. For the 33-ball case we see virtually the same estimates as in the constrained case. Thus, shifting the physical stimuli from exactly 0.5 to 0.55 added some degree of uncertainty to the subjective belief distribution.

For the 45-ball and 48-ball cases we infer virtually the same subjective probability distributions as in Figure 6. But for the 6-ball and 12-ball cases we get very different results. In each case the mode and mean of the subjective distribution are significantly greater than the objective probability. This contrasts with the distributions for the 45-ball and 48-ball cases. Now the 45-ball case corresponds to an objective probability of 0.75, and is therefore not directly symmetric to the 12-ball (0.2) or 6-ball (0.10) cases, but the 48-ball and 12-ball cases are exactly symmetric. Thus one might have expected to see “mirror image” subjective belief distributions in these two cases, and we do not.²²

4. Conclusions

In some sense, our approach to estimating subjectively uncertain beliefs is “too simple.” That simplicity derives from some strong assumptions, which we want to be clear about. We are not

²² Nor did we have a subject pool dominated by Dutch football fans or pro-British Protestants from Northern Ireland.

particularly concerned here about the specific parametric forms assumed, since they can be relaxed easily enough.²³ Nor does the artefactual exploitation of simple “bingo cage” stimuli, in which we know the true population process, concern us immediately, since if we cannot handle this controlled domain of inference then we have no business wandering blind into the field of natural processes. There are some deeper conceptual assumptions.

We assume that the subject evaluates lotteries defined over a stochastic process that has uncertain probabilities with the same risk attitudes that are used to evaluate lotteries defined over a stochastic process that has certain probabilities. In effect, we assume that aversion to uncertainty²⁴ is the same as aversion to risk. Starting with Ellsberg [1961], the validity of this assumption has been questioned. Working within an EUT framework, or very close to it, Ergin and Gul [2009], Klibanoff, Marinacci and Mukerji [2005], Nau [2006][2007] and Neilsen [2008] offer representations of preferences over uncertain and risk choices that differentiate uncertainty aversion and risk aversion.²⁵ Each representation collapses to the one we employ as a special case in which the two types of aversion are assumed to be the same, so in an important formal sense they each provide “smooth” characterizations of uncertainty aversion that differentiate it from risk aversion. Our approach provides the basis for estimating representations of this type.

Our results do show that beliefs for simple, physical stimuli do not appear to be statistically degenerate probability distributions. Identifying the theoretical structure of that non-degeneracy remains an open area of research.

²³ For example, it is a simple matter to more flexible specifications that allow for the experimental prizes to be integrated with some baseline income level (e.g., Andersen, Harrison, Lau and Rutström [2008]). Or non-EUT specifications that allow for probability weighting and/or loss aversion (e.g., Harrison and Rutström [2008; p.85ff.]). Our objective is not to find the best parametric specification for these data.

²⁴ The literature uses different terminology for concepts that are very similar, if not identical. The expression “uncertainty aversion” is used in Nau [2006], and is closely related to the concept of “ambiguity aversion” in Klibanoff, Marinacci and Mukerji [2005], the concept of “second-order risk aversion” in Ergin and Gul [2009], and even the concept of “information aversion” in Grant, Kajii and Polak [1998].

²⁵ Non-EUT characterizations are also popular: see Gilboa and Schmeidler [1989], Ghirardoto, Maccheroni and Marinacci [2004] and Gilboa, Postlewaite and Schmeidler [2008].

Figure 1: Illustrative Symmetric Logit-Normal Distributions





Figure 2: Illustrative Asymmetric Logit-Normal Distributions



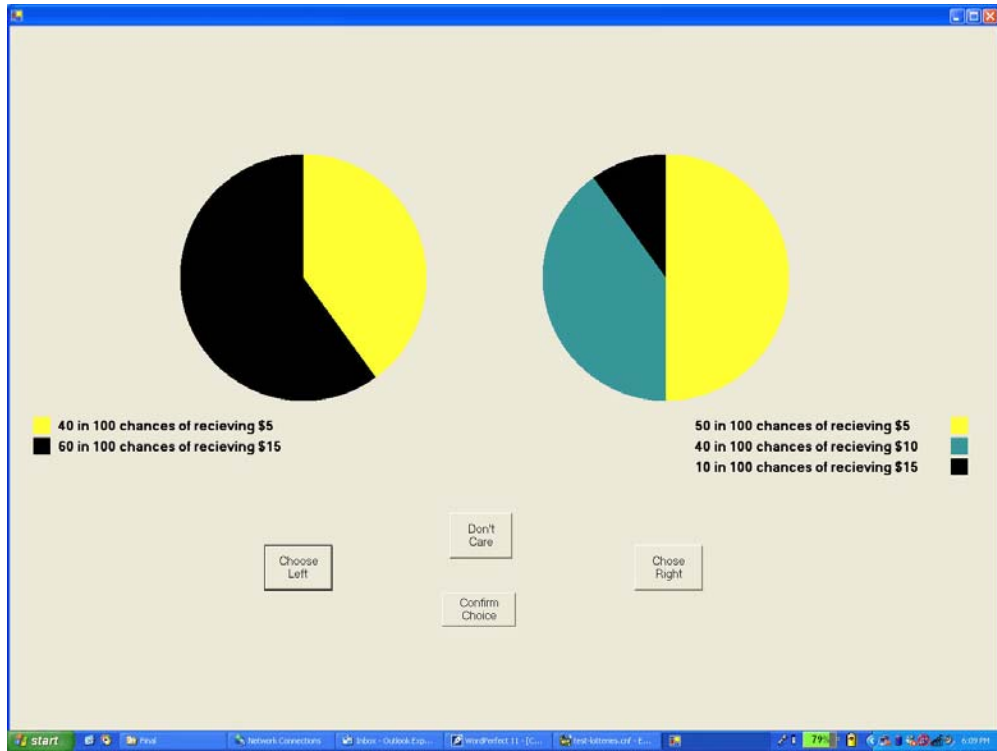


Figure 3: Illustrative Lottery Choice

PRACTICE: Picking a Ping-Pong Ball
Orange or White?

	House 1	House 2	House 3	House 4	House 5	House 6	House 7	House 8	House 9	
Stake you have to bet with	\$100	\$100	\$100	\$100	\$100	\$100	\$100	\$100	\$100	<input type="button" value="Submit Bets"/>
Orange	<input type="radio"/> \$10 0.1	<input type="radio"/> \$5 0.2	<input type="radio"/> \$3.33 0.3	<input type="radio"/> \$2.5 0.4	<input type="radio"/> \$2 0.5	<input type="radio"/> \$1.67 0.6	<input type="radio"/> \$1.43 0.7	<input type="radio"/> \$1.25 0.8	<input type="radio"/> \$1.11 0.9	Odds by house Probability by house
White	<input type="radio"/> \$1.11 0.9	<input type="radio"/> \$1.25 0.8	<input type="radio"/> \$1.43 0.7	<input type="radio"/> \$1.67 0.6	<input type="radio"/> \$2 0.5	<input type="radio"/> \$2.5 0.4	<input type="radio"/> \$3.33 0.3	<input type="radio"/> \$5 0.2	<input type="radio"/> \$10 0.1	Odds by house Probability by house
Earnings from your bet, if Orange										
Earnings from your bet, if White										

Figure 4: Illustrative Betting Choices




Figure 5: The Event

Table 1: Estimated Subjective Probabilities

There are 60 balls in all. So when there are 6, 12, 30, 33, 45 and 48 orange balls, the true probability is 0.1, 0.2, 0.5, 0.55, 0.75 and 0.8, respectively.

Parameter	Definition	Point Estimate	Standard Error	95% Confidence	
π_6	Subjective probability when 6 balls	0.001	0.014	0.000 [‡]	0.029
π_{12}	Subjective probability when 12 balls	0.369	0.017	0.335	0.403
π_{30}	Subjective probability when 30 balls	0.556	0.012	0.533	0.579
π_{33}	Subjective probability when 33 balls	0.548	0.007	0.534	0.562
π_{45}	Subjective probability when 45 balls	0.710	0.044	0.624	0.795
π_{48}	Subjective probability when 48 balls	0.718	0.025	0.668	0.767

Note: [‡] These estimates are calculated using the delta method, and approximation error with a point estimate so close to the lower boundary of 0 results in this lower bound being calculated to be negative (-0.028). This value of 0 is imposed *a priori*.

Figure 6: Estimated Belief Distributions
With Means Constrained to True Probabilities





Figure 7: Estimated Belief Distributions



References

- 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.
- Aitchison, J., and Begg, C.B., "Statistical Diagnosis when Basic Cases are Not Classified with Certainty," *Biometrika*, 63(1), 1-12.
- Andersen, Steffen; Fountain, John; Harrison, Glenn W., and Rutström, E. Elisabet, "Estimating Subjective Probabilities," *Working Paper 2010-06*, Center for the Economic Analysis of Risk, Robinson College of Business, Georgia State University, 2010.
- Andersen, Steffen; Harrison, Glenn W.; Hole, Arne Rise, and Rutström, E. Elisabet, "Non-Linear Mixed Logit," *Working Paper 2010-07*, Center for the Economic Analysis of Risk, Robinson College of Business, Georgia State University, 2010.
- 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.
- Binswanger, Hans P., "Attitudes Toward Risk: Theoretical Implications of an Experiment in Rural India," *Economic Journal*, 91, December 1981, 867-890.
- de Finetti, Bruno, "La Prévision: Ses Lois Logiques, Ses Sources Subjectives," *Annales de l'Institut Henri Poincaré*, 7, 1937, 1-68; English translation as "Foresight: Its Logical Laws, Its Subjective Sources," in H.E. Kyburg and H.E. Smokler (eds.), *Studies in Subjective Probability* (Huntington, NY: Robert E. Krieger, 1980, Second Edition).
- de Finetti, Bruno, "Logical Foundations and Measurement of Subjective Probability," *Acta Psychologica*, 34, December 1970, 129-145.
- Drukker, David M., and Gates, Richard, "Generating Halton sequences using Mata," *Stata Journal*, 6(2), 214-228.
- Eckel, Catherine C., and Grossman, Philip J., "Sex Differences and Statistical Stereotyping in Attitudes Toward Financial Risk," *Evolution and Human Behavior*, 23(4), 2002, 281-295.
- Ellsberg, Daniel, "Risk, Ambiguity, and the Savage axioms," *Quarterly Journal of Economics*, 75, 1961, 643-669.
- Epstein, Richard A., *The Theory of Gambling and Statistical Logic* (San Diego: Academic Press, 1977).
- Ergin, Haluk, and Gul, Faruk, "A Theory of Subjective Compound Lotteries," *Journal of Economic Theory*, 144(3), May 2009, 899-929.
- Gates, Richard, "A Mata Geweke-Hajivassiliou-Keane Multivariate Normal Simulator," *Stata Journal*, 6(2), 2006, 190-213.
- Ghirardoto, Paolo; Maccheroni, Fabio, and Marinacci, Massimo, "Differentiating Ambiguity and Ambiguity Attitude," *Journal of Economic Theory*, 118, 2004, 133-173.

- Gilboa, Itzhak; Postlewaite, Andrew P., and Schmeidler, David, "Probability and Uncertainty in Economic Modeling," *Journal of Economic Perspectives*, 22(3), Summer 2008, 173-188.
- Gilboa, Itzhak, and Schmeidler, David, "Maxmin Expected Utility with a Non-Unique Prior," *Journal of Mathematical Economics*, 18, 1989, 141-153.
- Gould, William; Pitblado, Jeffrey, and Sribney, William, *Maximum Likelihood Estimation With Stata* (College Station, TX: Stata Press, Third Edition, 2006).
- Grant, Simon; Kajii, Atsushi, and Polak, Ben, "Intrinsic Preference for Information," *Journal of Economic Theory*, 83, 1998, 233-259.
- Halevy, Yoram, "Ellsberg Revisited: An Experimental Study," *Econometrica*, 75, 2007, 503-536.
- Harrison, Glenn W.; Lau, Morten I., and Rutström, E. Elisabet, "Estimating Risk Attitudes in Denmark: A Field Experiment," *Scandinavian Journal of Economics*, 109(2), 2007, 341-368.
- 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).
- Hensher, David A., and Greene, William H., "The Mixed Logit Model: The State of Practice," *Transportation*, 30, 2003, 133-176.
- 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.
- Hole, Arne Risa, "Fitting Mixed Logit Models by Using Maximum Simulated Likelihood," *Stata Journal*, 7(3), 2007, 388-401.
- Holt, Charles A., and Laury, Susan K., "Risk Aversion and Incentive Effects," *American Economic Review*, 92(5), December 2002, 1644-1655.
- Huber, Joel, and Train, Kenneth, "On the Similarity of Classical and Bayesian Estimates of Individual Mean Partworths," *Marketing Letters*, 12(3), 2001, 259-269.
- Klibanoff, Peter; Marinacci, Massimo, and Mukerji, Sujoy, "A Smooth Model of Decision Making Under Ambiguity," *Econometrica*, 73(6), November 2005, 1849-1892.
- Lesaffre, Emmanuel; Rizopoulos, Dimitris, and Tsonaka, Roula, "The Logistic Transform for Bounded Outcome Scores," *Biostatistics*, 8(1), 2007, 72-85.
- Long, J. Scott, and Freese, Jeremy, *Regression Models for Categorical Dependent Variables Using Stata* (College Station, TX: Stata Press, Second Edition, 2006).
- Loomes, Graham, and Sugden, Robert, "Testing Different Stochastic Specifications of Risky Choice," *Economica*, 65, 1998, 581-598.

- Machina, Mark J., "Almost-objective uncertainty," *Economic Theory*, 24, 2004, 1-54.
- McFadden, Daniel, and Train, Kenneth, "Mixed MNL Models for Discrete Response," *Journal of Applied Econometrics*, 15, 2000, 447-470.
- 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, University of Tennessee, January 2008.
- Oehlert, Gary W., "A Note on the Delta Method," *The American Statistician*, 46(1), February 1992, 27-29.
- Quiggin, John, *Generalized Expected Utility Theory: The Rank-Dependent Model* (Norwell, MA: Kluwer Academic, 1993).
- Revelt, David, and Train, Kenneth, "Mixed Logit With Repeated Choices: Households' Choices of Appliance Efficiency Levels," *Review of Economics & Statistics*, 80, 1998, 647-657.
- 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., "Elicitation of Personal Probabilities and Expectations," *Journal of American Statistical Association*, 66, December 1971, 783-801.
- Segal, Uzi, "The Ellsberg Paradox and Risk Aversion: An Anticipated Utility Approach," *International Economic Review*, 28(1), February 1987, 175-202.
- Segal, Uzi, "Two-Stage Lotteries Without the Independence Axiom," *Econometrica*, 58(2), March 1990, 349-377.
- Smith, Vernon L., "Measuring Nonmonetary Utilities in Uncertain Choices: the Ellsberg Urn," *Quarterly Journal of Economics*, 83(2), May 1969, 324-329.
- Train, Kenneth E., *Discrete Choice Methods with Simulation* (New York: Cambridge University Press, 2003).
- 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, 2010 forthcoming.