

Subjective Bayesian Beliefs

by

Constantinos Antoniou, Glenn W. Harrison, Morten Lau and Daniel Read[†]

August 2013

Abstract. A large literature suggests that many individuals do not apply Bayes Rule correctly when making decisions that depend on them correctly pooling prior information and sample data. We replicate and extend a classic experimental study of Bayesian updating from psychology, employing the methods of experimental economics, with careful controls for the confounding effects of risk aversion. Our results show that risk aversion significantly alters inferences on deviations from Bayes Rule.

[†] Warwick Business School, University of Warwick (Antoniou and Read); Department of Risk Management & Insurance and Center for the Economic Analysis of Risk, Robinson College of Business, Georgia State University, USA (Harrison); Durham Business School, Durham University, UK, and Copenhagen Business School, Denmark (Lau). Email: Constantinos.Antoniou@wbs.ac.uk, gharrison@gsu.edu, mla.eco@cbs.dk, and Daniel.Read@wbs.ac.uk. Harrison thanks the U.S. National Science Foundation for research support under grants NSF/HSD 0527675 and NSF/SES 0616746, and we thank the referee and seminar and conference participants for valuable comments.

Table of Contents

1. Experimental Design.....	2
<i>A. Characterizing Attitudes Towards Risk</i>	2
<i>B. Eliciting Subjective Beliefs</i>	3
<i>C. General Procedures</i>	5
2. A Subjective Expected Utility Representation	6
<i>A. Bayes Rule</i>	6
<i>B. A Structural Model for Risk Aversion</i>	6
<i>C. A Structural Model for Beliefs</i>	8
<i>D. Results</i>	11
3. Probability Weighting with Subjective Probabilities.....	12
4. Alternative Hypotheses.....	16
5. Previous Literature	17
6. Conclusions	18
References.....	20
Appendix 1: Experimental Instructions (NOT FOR PUBLICATION)	31
<i>A. Risk Attitudes Task</i>	31
<i>B. Betting Task</i>	32
<i>C. Demographic Questionnaire</i>	35
Appendix 2: Derivations (NOT FOR PUBLICATION).....	39
Appendix 3: Detailed Estimation Results (NOT FOR PUBLICATION)	40

Virtually all decision making occurs in the context of information that acts as the basis for action. A pedestrian standing at the sidewalk looks both ways and then crosses. The decision to cross or not will be a function of (at least) how urgently the pedestrian wants to cross, how averse he is to being struck by a car, and his assessments of the probability that this will happen. Similarly, a bettor at the racetrack bets on a particular horse for specific odds. This decision will again be a function of the bettor's inferred probabilities and her utility function. If we observe many such decisions by the same person, and if we make some "structural assumptions," such as the functional form of the pedestrian's or bettor's utility function, and/or their probability weighting function, we can in principle infer their beliefs about the likelihood of the event of interest. We can also ask how these choice probabilities are related to the objective probabilities that "should" be inferred from the information they have. We investigate how individuals infer the probability of an uncertain event from information sufficient to generate a unique posterior probability using Bayes Rule. We apply theory, experiments, and econometric methods to evaluate how subjective beliefs are formed.

Bayes Rule provides us with a method for determining the objective relationship between information (e.g., "how fast was the horse in the last race?") with probabilities ("how likely is the horse to win this race?"). A large literature in experimental economics and psychology suggests that many individuals fail to apply Bayes Rule correctly. This literature typically assumes, explicitly or implicitly, that respondents are risk neutral. But in reality respondents exhibit preferences consistent with non-linear utility and/or probability weighting functions, so it remains unclear from these experiments whether these violations of Bayes Rule are genuine or somehow relate to non-risk neutral preferences. We address this issue by modeling information use structurally, to establish whether apparent deviations from Bayes Rule are eliminated, diminished, or possibly even increased when we account for subjects' actual, revealed preferences toward risk.

We obtained the data on which we conducted our analysis by replicating and extending a

classic experiment from psychology (Griffin and Tversky, 1992), but now using the methods of experimental economics. We employed real monetary consequences for decisions, and used a transparent and physical mechanism to generate the uncertain outcomes. We directly elicited bets about outcomes that we could use to infer the respondent's beliefs, and directly estimated their risk attitudes using data from a separate task.

In our analysis we start by making the conventional structural assumption of subjective expected utility (SEU). With this specification, average behavior strongly violates Bayes Rule. Interestingly, this model shows respondents in a better light, in the sense that their inferred probabilities are closer to the correct one, if we impose a linear utility function on the model. However, the data reject the assumption of risk neutrality. We then allow for violations of the independence axiom by means of probability weighting, in a rank dependent utility (RDU) framework, but find that probability weighting cannot explain the observed deviations from Bayes Rule. In summary, our evidence suggests that SEU or RDU preferences influence but do not eliminate violations from Bayes Rule.

1. Experimental Design

We employed two choice tasks. The first was aimed at identifying risk attitudes in settings with known probabilities, and the second at eliciting subjective beliefs for events that differed in terms of the priors and sample stimuli presented to subjects. Both tasks are needed to fully evaluate the extent to which individuals correctly apply Bayes Rule.

A. Characterizing Attitudes Towards Risk

We presented subjects with 20 binary choices between lotteries, patterned on those used by Hey and Orme (1994). Each lottery consisted of 1, 2 or 3 monetary prizes, with four possible monetary values of £0, £5, £10 or £15. Figure 1 displays a typical lottery pair. We created 60

distinct lottery pairs, and then divided them into 3 groups of 20. Each subject made choices for one of these groups. At the end of the task one of the 20 choices of each subject was selected at random and played out for money. Appendix 1 provides the complete instructions.

These data allow us to identify the risk attitudes of subjects, using the econometric methods described in Harrison and Rutström (2008). The estimation approach can include RDU specifications as well as traditional SEU.

B. Eliciting Subjective Beliefs

In the second task subjects chose between bets concerning events with unknown probabilities. They were provided with enough information to infer these probabilities using Bayes Rule. From the choices made, in combination with some assumptions about risk attitudes, we could estimate the subjective probabilities that best described their choices. Table 1 shows the “betting sheet” with which respondents made their choices.

The task involved a white box and a blue box, each containing ten-sided dice. The white box contained N 10-sided dice that each had 6 white and 4 blue sides, and the blue box contained similar dice that each had 6 blue and 4 white sides. The number of dice in each box varied across choice tasks ($N=3, 5, 9$ or 17), but the two boxes always contained the same number of dice.

At the beginning of each round we rolled a 6-sided die, with 3 blue and 3 white sides, and then selected either the white or blue box depending on the outcome of this roll. We then rolled the N dice in the selected box and announced the outcome. Hence the *prior probability* of the box being white or blue was 0.5, and the subject was given some *sample information* from the selected box with which to make more informed inferences about the color of the selected box. For example, suppose the two boxes each contained $N=3$ dice, and the six-sided die showed a white face. We would then roll the 3 dice in the white box and announce the number of white and blue sides that came up. The

color of the selected box was not announced.

After announcing the outcome from rolling the dice, subjects were asked to place their bets. This was framed as placing a £3 bet in each of 19 different betting houses that offer different odds on which box was selected. The £3 stake from one bookie was not transferable to another. The general form of this betting procedure for eliciting choices that depend on subjective probabilities is well known (e.g., Savage, 1971).

For each bookie, subjects placed their £3 stake on either the blue or the white box for that bookie. One would expect that subjects would be inclined to bet on the white box when the odds were generous enough for that box, and to switch to betting on the blue box for less generous odds, where what is “generous” depends on the subjects’ beliefs. For example, the odds offered by Bookie 4 imply a probability of 1/5 for the white box, and a probability of 4/5 for the blue. A *risk-neutral* subject would then bet on White if and only if they believed the probability the selected box was white was 1/5 or greater. The switch point therefore corresponds to an interval of betting house probabilities that the white box was chosen. Again, if the same risk-neutral subject chose White for Bookie 4 and Blue for Bookie 5, this means his subjective probability for the White box would lie in the interval $[1/5, 1/4]$. In this way that subjective probability is “trapped” in the classic revealed preference manner.

The recovery of subjective probabilities and beliefs requires formal theoretical and parametric assumptions, which we detail below. The essential logic is that this decision sheet is a “multiple price list,” just like the decision sheet used to infer discount rates by Harrison, Lau and Williams (2002), the decision sheet used to infer risk attitudes by Holt and Laury (2002), and the decision sheet used to infer valuations for goods by Andersen, Harrison, Lau and Rutström (2007). The general “multiple price list” betting interface employed here was first used by Fiore, Harrison, Hughes and Rutström (2009).

Each subject participated for 30 rounds of this betting task, 4 rounds with $N=3$ dice, 14

rounds with $N=5$ dice, 6 rounds with $N=9$ dice, and 6 rounds with $N=17$ dice. This distribution of sample sizes was chosen to ensure that we observe roughly the same number of “extreme” samples.¹

To control for possible order effects, in half of the 12 sessions we counterbalanced the order of the risk and betting tasks. In the betting task we varied the presentation of sample sizes in ascending order (i.e., first 4 rounds of 3, then 14 rounds of 5, etc.) and in descending order (6 rounds of 17, 6 rounds of 9, etc). Therefore we have an overall 2×2 design, with 4 treatments in total.

Our instructions illustrated the factors that affect betting in field betting markets, such as betting on a horse race with different bookies, and drew parallels between such naturally occurring events and our task. Subjects first read these instructions quietly; we then read them aloud and allowed time for questions. The instructions are reproduced in Appendix 1. We had 3 practice rounds with boxes containing 4 dice and hypothetical bets, so that subjects could become accustomed to the process. At the end of the 30 rounds we randomly choose one bet for each subject, and played that bet out for real consequences.²

C. General Procedures

We recruited 111 subjects from the University of Durham. Subjects were recruited using a computerized interface, after being solicited in general terms to register for paid experiments. All subjects received a £5 show up fee. Apart from the tasks described above, each subject completed a survey of demographic characteristics shown in Appendix 1. Payments for the experiment totaled £2,692, for an average payment of £24.26 per subject.

There were 12 experimental sessions, each having approximately 10 subjects. To ensure

¹ The most unlikely pattern is (5,0) with probability of occurring equal to 0.044, followed by the pattern (7,2) with probability 0.091, and the pattern (11,6) with probability 0.104. Based on these probabilities we chose the frequency of each sample size to roughly equalize the expected frequency of the most unlikely patterns in each session.

² For example, each subject filled out 30 betting sheets, such as the one shown by Table 1. At the end of the experiment we first randomly chose, for each subject, one of these betting sheets, and then we chose 1 of the 19 bookies within that selected sheet. If the subject placed the allocated £3 on the box that was actually chosen, he was paid the amount that corresponds to the odds offered by that bookie.

credibility, in each session a randomly selected subject acted as a monitor for the belief task, rolling the dice, counting the number of blue and white sides, and announcing the outcome. The monitor was paid a flat fee of £10 for the belief task, and participated with everyone else in the risk task.

2. A Subjective Expected Utility Representation

A. Bayes Rule

With no additional information from the rolling of dice from the chosen box, the probability of either box being selected is 0.5. By announcing the outcome from rolling the N dice, this prior can be updated in accordance with Bayes Rule.³ For example, if we roll N=5 dice and get 4 dice with a white face and 1 die with a blue face, the posterior probability of the white box is 0.77.

B. A Structural Model for Risk Aversion

We assume a Constant Relative Risk Aversion (CRRA) utility function, defined as

$$U(y) = y^{(1-r)} / (1-r), \quad (1)$$

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

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

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

$$\nabla EU = EU_R - EU_L \quad (3)$$

³ We provide a complete derivation of the application of Bayes Rule for this process in Appendix 2.

calculated, and where EU_L is the “left” lottery and EU_R is the “right” lottery, as displayed to the subject and illustrated in Figure 1. This latent index, based on latent preferences, is then linked to the observed choices using a standard cumulative normal distribution function $\Phi(\nabla EU)$. This “probit” function takes any argument between $\pm\infty$ and transforms it into a number between 0 and 1. Thus we have the probit link function,

$$\text{prob}(\text{choose lottery R}) = \Phi(\nabla EU) \quad (4)$$

Thus the likelihood of the observed responses, conditional on the EUT and CRRA specifications being true, depends on the estimates of r given the above statistical specification and the observed choices. If we ignore responses that reflect indifference⁴ the log-likelihood is then

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

where $\mathbf{I}(\cdot)$ is the indicator function, $y_i = 1(-1)$ denoting the choice of the Option R (L) lottery in risk aversion task i , and $\mathbf{X}$ is a vector of covariates (e.g., individual characteristics reflecting age, sex, race, and so on, or treatment dummies). The structural parameter r is modeled as a linear function of the covariates in $\mathbf{X}$.

An important extension of the core model is to allow for subjects to make some errors. The notion of error is one that has already been encountered in the form of the statistical assumption (4) that the probability of choosing a lottery is not 1 when the EU of that lottery exceeds the EU of the other lottery. By varying the shape of the link function implicit in (4), one can informally imagine subjects that are more sensitive to a given difference in the index ∇EU and subjects that are not so sensitive. We use the contextual error specification proposed by Wilcox (2011). It posits the latent index:

⁴ In our lottery experiments the subjects are told at the outset that any expression of indifference would mean that the experimenter would toss a fair coin to make the decision for them if that choice was selected to be played out. Hence one can modify the likelihood to take these responses into account either by recognizing this is a third option, the compound lottery of the two lotteries, or alternatively that such choices imply a 50:50 mixture of the likelihood of choosing either lottery, as illustrated by Harrison and Rutström (2008; p.71). We do not consider indifference here because it was an extremely rare event.

$$\text{prob}(\text{choose lottery } R) = \Phi [(\nabla EU)/v)/\mu] \quad (4')$$

instead of (4), where v is a normalizing term for each lottery pair L and R , and $\mu > 0$ is a structural “noise parameter” used to allow some errors from the perspective of the deterministic EUT model.⁵

Thus we extend the likelihood specification to include the noise parameter μ and maximize $\ln L(r, \mu; y, \mathbf{X})$ by estimating r and μ , given observations on y and $\mathbf{X}$.⁶ Additional details of the estimation methods used, including corrections for “clustered” errors when we pool choices over subjects and tasks, are provided by Harrison and Rutström (2008; p.69ff).

C. A Structural Model for Beliefs

The responses to the belief elicitation task can be used to draw estimates about the belief that each subject holds if we are willing to assume something about how they make decisions under risk. To allow for the general case in which we have risk aversion, we jointly estimate the subjective probability and the parameters of the utility function, following Andersen, Fountain, Harrison and Rutström (2013). Using the schema in Table 1, the subject that selects event W from a given bookie b receives EU

$$EU_w = \pi_w \times U(\text{payout if } W \text{ occurs} \mid \text{bet on } W) + (1-\pi_w) \times U(\text{payout if } B \text{ occurs} \mid \text{bet on } W) \quad (6)$$

where π_w is the subjective probability that W will occur. The payouts that enter the utility function are defined by the odds that each bookie offers, and are shown in Table 1. For the bet offered by the first bookie, for example, these payouts are £60 and £0, so we have

$$EU_w = \pi_w \times U(£60) + (1-\pi_w) \times U(£0) \quad (6')$$

⁵ The normalizing term v is defined as the maximum utility over all prizes in this lottery pair minus the minimum utility over all prizes in this lottery pair, and ensures that the normalized EU difference $[(EU_R - EU_L)/v]$ remains in the unit interval. As $\mu \rightarrow \infty$ this specification collapses ∇EU to 0 for any values of EU_R and EU_L , so the probability of either choice converges to $1/2$. So a larger μ means that the difference in the EU of the two lotteries, conditional on the estimate of r , has less predictive effect on choices. Thus μ can be viewed as a parameter that flattens out, or “sharpens,” the link functions implicit in (4). This is just one of several different types of error story that could be used, and Wilcox (2008) provides a masterful review of the implications of the strengths and weaknesses of the major alternatives.

⁶ The normalizing term v is given by the value of r and the lottery parameters, which are part of $\mathbf{X}$.

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

$$EU_B = \pi_w \times U(\text{payout if W occurs} \mid \text{bet on B}) + (1-\pi_w) \times U(\text{payout if B occurs} \mid \text{bet on B}). \quad (7)$$

and this translates for the first bookie in Table 1 into payouts of £0 and £3.15, so we have

$$EU_B = \pi_w \times U(£0) + (1-\pi_w) \times U(£3.15) \quad (7')$$

for this particular bookie and bet. We observe the bet made by the subject for a range of odds, so we can calculate the likelihood of that choice given values of r , π_w and μ .

The rest of the structural specification is exactly the same as for the choices over lotteries with objective probabilities. Given (6) and (7), we can define the latent index that is the counterpart to (3) as

$$\nabla EU = EU_W - EU_B \quad (3')$$

for each of the W and B bets from Table 1. The counterpart to (4') is then

$$\text{prob}(\text{choose lottery R}) = \Phi [(\nabla EU)/\nu] / \omega \quad (4'')$$

where ν is a normalizing term for each lottery pair W and B, and calculated with the same logic as before, and $\omega > 0$ is a structural “noise parameter” for the belief choices that is used to allow some errors from the perspective of the deterministic SEU model.

Writing out the complete likelihood function, we have

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

for the observed choices in the task defined over objective probabilities, and

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

for the observed choices in the task defined over subjective probabilities. The *joint estimation* problem is to find values for r , π_w , μ and ω that maximize the *sum* of (5') and (5''). One can think of each binary choice in the two tasks as rows in a matrix, and then (5') as being the likelihood of the choices in the top part of the matrix and (5'') as being the likelihood of the choices in the bottom part of the matrix.

The overall likelihood is then just the sum of the likelihoods for all choices made, whether the

probabilities are objective or subjective.

It is useful to see how the estimation procedure maps back to the economics of the SEU model. Ignoring the behavioral error terms μ and ω , we need r to evaluate the utility function in (6) and (7), we need π_w to calculate the EU in (6) and (7) once we know the utility values, and we need both of them to calculate the latent index in (4'') that generate the probability of observing the choice of bet W or bet B. The *joint* maximum likelihood problem, again, 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 under EUT, since π_w plays no direct role in explaining the choices in that task.

This formal analysis assumes that we are estimating one subjective probability π_w . There are two simple extensions that allow us to consider our complete design, which involves the 6 posterior probabilities shown in Table 2. The first is to assume symmetry, in the sense that the estimate of π_w is treated as $(1-\pi_w)$ when evaluating the choices made for the corresponding task. In other words, if we only use the choices for the betting task that has posterior probability 0.60, we could directly apply the existing formal analysis. But we can also include the choices for the betting task that has posterior probability 0.40, and *assume* that the subjective probability for those choices is one minus the subjective probability for the choices with posterior probability 0.60. Therefore we only need to estimate one subjective probability. This seems to be an innocuous assumption, and is directly testable. The second extension is then to introduce two extra subjective probability parameters, so we have one for the betting task with posterior probability 0.6 (and 0.4 by symmetry), one for the betting task with posterior probability 0.77 (and 0.23), and one for the betting task with posterior probability 0.88 (and 0.12). Hence we have to estimate one risk attitude parameter and three subjective probabilities, along with the two behavioral error terms.

We include binary dummies in the vector $\mathbf{X}$ to control for procedural checks. One is whether

the risk aversion task came before or after the betting task, and the other is whether the number of dice (N) was presented to subjects in an ascending or descending order in the betting task. We also include a series of individual demographic characteristics defined in a binary fashion: female (1 if the subject is female), teenager (1 if subject is less than 20 years old), white (1 if the subject described himself as white), British (1 if the subject is a British citizen), high income (1 if the subjects earn more than £10,000 per year), graduate (1 if the subject is a graduate student) and math (1 if the subject studies Economics, Finance, Engineering, Mathematics, Computer or Physical Sciences).

D. Results

Detailed estimates of all models are given in Appendix 3. We found evidence for modest risk aversion on average, consistent with the general finding for populations of this kind. The estimate for r was around 0.5, and highly significant ($p < .001$), but there were also significant order effects. Subjects were more risk averse when the risk aversion came first than when it came second (see Table A1). Such order effects are common in the literature (e.g., Harrison, Johnson, McInnes and Rutström, 2005). With respect to demographic characteristics, women were significantly more risk averse than men, as were high income subjects.

Figure 2 displays the estimated subjective probabilities, pooling the symmetric cases, with reference lines showing the corresponding posterior probability using Bayes Rule. The dispersion in these distributions reflects the standard errors in the parameter estimates of the subjective probabilities, as well as variations across subjects due to differences in demographic characteristics. There is a marked *underestimation* of the true Bayesian posterior probability, which becomes larger as the posterior increases from 0.6 to 0.77 and 0.88. In addition, the precision of the subjective probability estimate also declines as the posterior gets larger. Given that we pooled the symmetric cases, Figure 2 also implies a systematic *overestimation* of the true Bayesian posterior probability when

the posterior is *less* than $\frac{1}{2}$, becoming larger as the posterior decreases from 0.4 to 0.23 and 0.12. We return to this pattern in the next section.

Figure 3 shows the corresponding estimates of subjective probability when we impose risk neutrality. The striking result is that Bayes Rule does much *better* than when we allow the data to determine the risk attitudes of subjects. In the risk neutral case the subjective estimate of the 0.6 posterior probability is exactly correct, and the estimates for the 0.77 and 0.88 posterior probabilities are much closer to the Bayesian posterior. Figure 4 superimposes the results from Figures 2 and 3. The qualitative effect of risk aversion follows immediately from theory: the more risk averse the subject, the more likely she is to bet as if her subjective probability is 0.5, since this reduces the dispersion in final outcomes from all bets.⁷ The log-likelihood of the risk neutral model (-92280.4) is significantly worse than the log-likelihood for the general model (-83967.5), as one would expect from the rejection of risk neutrality when just looking at the risk choices.

These are striking results. If one *fails* to correct for the non-linearity of the utility function evident in the data, as analyzed here, then one finds *stronger* behavioral support for the use of Bayes Rule. But if one makes these corrections, this support melts away.

3. Probability Weighting with Subjective Probabilities

The stylized finding from our estimates is that subjective beliefs underestimate the true Bayesian posterior probability when the posterior is greater than 0.5, and overestimate it when the posterior is less than 0.5. Of course, that sounds a lot like the standard type of probability weighting that plays a role in RDU and Prospect Theory models of decision-making under risk.⁸ Since our choices involve prizes framed entirely as gains, we examined a conventional RDU model to assess if probability weighting accounts for our results on subjective beliefs. The RDU model relaxes the

⁷ This qualitative result is exactly the same as one finds using a Quadratic or Linear scoring rule (providing, in the latter case, that one does not go to the extreme of exact risk neutrality): see Andersen, Fountain, Harrison and Rutström (2013).

⁸ Holt and Smith (2009; §5) make exactly the same observation in a comparable study that we discuss further in section 5.

Independence axiom of and extends the SEU model by allowing for non-additive decision weights on lottery outcomes.⁹

To calculate decision weights under RDU one first rank-orders outcomes from worst to best, and then replaces the expected utility defined by (2),

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

with the RDU

$$RDU_i = \sum_{k=1, K} [w_k \times u_k] \quad (12)$$

where

$$w_i = \omega(p_i + \dots + p_n) - \omega(p_{i+1} + \dots + p_n) \quad (13a)$$

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

$$w_i = \omega(p_i) \quad (13b)$$

for $i=n$, the subscript indicating outcomes ranked from worst to best, and where $\omega(p)$ is some probability weighting function.

The following probability weighting function popularized by Tversky and Kahneman (1992) has been widely used:

$$\omega(p) = p^\gamma / [p^\gamma + (1-p)^\gamma]^{1/\gamma} \quad (14)$$

for $0 < p < 1$. The normal assumption, backed by some evidence reviewed by Gonzalez and Wu (1999), is that $0 < \gamma < 1$. This gives the weighting function an “inverse S-shape,” characterized by a concave section signifying the overweighting of small probabilities up to a crossover-point where $\omega(p)=p$, beyond which there is then a convex section signifying underweighting. If $\gamma > 1$ the function takes the less conventional “S-shape,” with convexity for smaller probabilities and concavity for larger probabilities.

⁹ In its traditional form the RDU model assumes reduction of compound lotteries, a point we return to below. Our application here is unconventional in the application of concepts of “probabilistic sophistication” to a particular non-EUT model, in the spirit of Machina and Schmeidler (1992, 1995) and Grant (1995). It is unconventional in the sense of assuming that individuals behave as if they distort (weight) objective probabilities, but they are otherwise probabilistically sophisticated.

There are some limitations of this probability weighting function. It does not allow independent specification of location and curvature, and is not even increasing in p for small values of γ . The first of these restrictions is particularly problematic in the present case, since $1/2$ is a natural fixed point for the belief task.

Prelec (1998) offers a two-parameter probability weighting function that exhibits more flexibility than (14):

$$\omega(p) = \exp\{-\eta(-\ln p^\varphi)\}, \quad (15)$$

which is defined for $0 < p < 1$, $\eta > 0$ and $0 < \varphi < 1$. Rieger and Wang (2006; Proposition 2) offer a two-parameter polynomial of 3rd degree which is defined for $0 \leq p \leq 1$, unlike (15):

$$\omega(p) = p + [(3-3b)/(a^2-a+1)] [p^3-(a+1)p^2+ap], \quad (16)$$

where $0 < a < 1$ and $0 < b < 1$. The parameter restrictions on a and b ensure that the function is concave for lower values of p and then convex for larger values of p . Values of b larger than 1 would allow convex and then concave shapes, which we want to allow *a priori*.

To illustrate the basic logic of accounting for probability weighting, Figure 5 shows how the parameter γ in (14) characterizes the probability weighting function and the decision weights used to evaluate lottery choices. We employ the value $\gamma=0.7$, to illustrate how this function might account for the subjective probabilities we infer; in fact, the estimated value using this function and our data is 0.88, which is much closer to the 45° line. For simplicity here we assume lotteries with 2, 3 or 4 prizes that are equally likely when we generate the decision weights. So for the case of 2 prizes, each prize has $p=1/2$; with 3 prizes, each prize has $p=1/3$; and with 4 prizes, each prize has $p=1/4$. For the 3-prize and 4-prize lottery we see the standard result: the decision weights on the largest and smallest prizes are relatively greater than the true probability, and the decision weights on the intermediate prizes are relatively smaller than the true probability. In the belief elicitation task there were only 2 prizes per

lottery, so this value of the parameter a puts greater decision weight on the smaller prize.

Each panel in Figure 5 is important for our analysis. To estimate γ from the observed lottery choices with known probabilities we only need the decision weights in the right panel of Figure 5. But to recover a subjective probability π subject to probability weighting, we only need the probability weighting function. Or rather, we need its inverse function, since it is the π in the $\omega(\pi)$ function that we are seeking to recover. We do not directly observe $\omega(p)$ or $\omega(\pi)$, but we can estimate $\omega(\cdot)$ as part of the latent structure generating the observed choices in the two types of tasks, implicitly assuming that $\omega(p) = \omega(\pi)$. Once we have $\omega(\cdot)$ we can then recover π by directly applying the estimated probability weighting function, such as the one shown, for a typical γ , in the left panel of Figure 5.

As noted, and shown in Figure 5, there is one problem with this function: it imposes a fixed point at $p=1/3$ or $p=2/3$, depending on the value of a . So if it explains the underweighting of posterior probabilities greater than $1/2$, since $\omega(p) < p$ for $p > 1/2$, it fails to account for the consistent overweighting of smaller posterior probabilities $1/3 \leq p < 1/2$. It might be argued that, from a practical perspective, this does not matter. As can be seen in Figure 5 the estimated subjective probability for the posterior of 0.6 is only slightly less than the true Bayesian posterior, implying by symmetry that the estimated subjective probability for the posterior of 0.4 is only slightly greater than the true Bayesian posterior. So probability weighting can account for the overweighting of the smaller Bayesian posterior probabilities, of 0.23 and 0.12, even if it does not account for the overweighting of 0.4.

Unfortunately this explanation fails. The value of γ estimated from the data is much larger than the one depicted in Figure 5. Moreover, the qualitative pattern in Figure 5 is effectively rejected by the more flexible Rieger and Wang (2006) form, which generates the estimated probability weighting function in Figure 6 for our data. As can be seen, this estimate shows overweighting of probabilities over almost the entire range. This is the exact *reverse* of the pattern of probability weighting than would

be needed to reconcile our estimates with the Bayesian posterior probability.

4. Alternative Hypotheses

We have just shown that relaxing the Independence Axiom in SEU does not allow us to explain observed behavior. In this section we consider alternative approaches.

Both SEU and conventional RDU assume the reduction of compound lotteries (ROCL) axiom. In SEU it is front and center, and explicit. In RDU it is commonly implied: for example, Quiggin (1982; p.331) writes that “...acceptance of the NM complexity axiom (that the value of a compound lottery depends only on the probability of each outcome) is implicit” in the notation he uses. However, Segal (1987, 1988, 1990, 1992) proposes an alternative version of RDU that does not assume ROCL. Precisely *how* one relaxes ROCL is a matter for important, foundational research. Although it has taken half a century for the implications of Ellsberg (1961) to be formalized in tractable ways, we are much closer to doing so, and virtually all point to the role of ROCL in understanding uncertainty and ambiguity. One popular approach is the “smooth ambiguity model” of Klibanoff, Marinacci and Mukerji (2005), with important parallels in Davis and Paté-Cornell (1994), Ergin and Gul (2009), Nau (2006) and Neilson (2010). Another popular approach is due to Ghirardoto, Maccheroni and Marinacci (2004), generalizing Gilboa and Schmeidler (1989). These models could be used to formally explain how the decision-maker reacts to a non-degenerate subjective posterior distribution *if* that distribution is not collapsed down to its mean by applying ROCL. In effect, the suggestion is that decision makers have some belief that the posterior is “more or less” equal to some value, which may or may not be centered on the true posterior probability. Even if it is centered on the true posterior probability, they will exhibit uncertainty aversion or ambiguity aversion towards it when placing bets that depend on it.

Another source of hypotheses about the behavior we document and characterize is the vast

literature on “background risk.” The idea is to consider the effect on foreground risk taking of the existence of some zero-mean background risk that is uncorrelated with the foreground risk. When posterior probabilities arise from an updating process of the kind considered here, we might reasonably hypothesize that there could be some background risk that the probability is calculated correctly. Even in the case of zero-mean background risks, positive and negative effects on foreground risk aversion can be predicted depending on the use of EUT or RDU models, as illustrated by Eeckhoudt, Gollier and Schlesinger (1996), Gollier and Pratt (1996) and Quiggin (2003). Things become more complicated, and interesting, if this cognitive-load background risk has non-zero mean *and* is asymmetric around the true posterior probability, as suggested by Figures 2, 3 and 4.

Yet another source of alternative hypotheses comes from other specifications of decision-making under risk that relax SEU. For instance, a simple reference-dependent model of “disappointment aversion,” in the spirit of Gul (1991), might generate different utility specifications that can be used to infer latent subjective probabilities. In effect, this is a behavioral, non-SEU version of the basic methodological point of Savage (1972), about the theoretical need to identify the model of decision making under risk at the same time as one identifies subjective probabilities. For Savage (1972), of course, the model was SEU, but one expects that if the evidence demanded that EUT be revised, he would still insist on the need for joint identification.

5. Previous Literature

The closest studies to ours that also use real rewards and incentive-compatible designs are those of Grether (1992) and Holt and Smith (2009), both of whom used a procedure to elicit or infer subjective probabilities in a Bayesian environment of priors and sample realizations.

Grether (1992) employed a variant of the Becker, DeGroot and Marshack (1964) procedure to elicit subjective probabilities. It is a variant because the usual application of that procedure involves

prices rather than probabilities, but the theoretical incentive compatibility of the procedure remains. Unfortunately, so do the problematic behavioral features of the procedure: the incentives for truth-telling are weak (Harrison, 1992), and there is a risk of subject confusion (Rutström, 1998; Plott and Zeiler, 2005). His statistical analysis is essentially descriptive, and there is no structural model.

Holt and Smith (2009) use the same procedures as Grether (1992), and found evidence for the use of Bayes Rule when priors were close to 0.5. However, as priors deviated from 0.5, they observed a tendency to overweight low posterior probabilities and overweight high posterior probabilities. They argued that probability weighting of the posterior cannot account for their findings, since the proportional overweighting (of the same posteriors) is much greater for low priors than for priors closer to 0.5. Their structural estimation is focused on the extent to which probability weighting can account for the observed data. In part this is because they do not need to correct the elicited subjective probability reports for the non-linearity of utility.

6. Conclusions

Our objective has been to establish a “base camp” for attacking what we believe to be the fundamental characteristic of subjective Bayesian beliefs. This is a recognition that the stochastic process that generates posterior probabilities should be viewed as *more uncertain* than the stochastic process that generates risk when the probability of final outcomes is directly given. In other words, the posterior probability should be viewed as a subjective probability which may be seen by the decision-maker as subject to “uncertainty aversion” that exacerbates the effect of traditional “risk aversion.” If this hypothesis is correct then the decision-maker will make choices that differ from those that would be made if she was neutral towards uncertainty. Consequently, the subjective posterior probability inferred from observed choices will differ depending on whether one allows for the possibility of uncertainty aversion. Previous analyses of subjective Bayesian decision-making,

including our own here, have assumed that the subject is neutral towards the uncertainty that is involved in the use of an inferred posterior probability.

To address this hypothesis one would need theoretical, experimental and econometric extensions of our approach. The theoretical extensions are to state structures in which uncertainty (or ambiguity, as it is often called) plays a non-degenerate role: for example, Gilboa and Schmeidler (1989), Ghirardoto, Maccheroni and Marinacci (2004), Klibanoff, Marinacci and Mukerji (2005) and Gilboa, Postlewaite and Schmeidler (2008). The extensions of experimental and econometric design require that one construct tasks that allow those structures to be identified.

Our approach is motivated by the same puzzle that has spurred the development of models towards uncertainty. Something tells us that behavior towards an event that has a 50% chance of occurring with probability 0 and a 50% probability of occurring with probability 1 *could* reasonably differ from behavior towards an event that has a 100% chance of occurring with probability 0.5. Yet, by some readings and axioms, these are not even two different states of the world, even though one can easily envisage distinct physical processes for each.¹⁰ Our version of this puzzle is that one could reasonably expect behavior to differ when a decision-maker is credibly told that the probability of some event is 0.87 compared to when he is credibly given priors and sample realizations that imply, under Bayes Rule, a posterior probability of 0.87. We hypothesize that the reasons for differences in behavior in these two puzzles are fundamentally the same.

¹⁰ In many physical processes, for example, threshold effects can lead to extreme outcomes rather than intermediate ones. So one person might believe the threshold has been exceeded, and another person might believe it has not.

References

- Andersen, S.; Fountain, J.; Harrison, G.W., and Rutström, E.E., “Estimating Subjective Probabilities,” *Journal of Risk & Uncertainty*, 2013, forthcoming.
- Andersen, S.; Harrison, G.W.; Lau, M.I., and Rutström, E.E., “Valuation Using Multiple Price List Formats,” *Applied Economics*, 39, 2007, 675-682.
- Becker, G.M.; DeGroot, M.H., and Marschak, J., “Measuring Utility By A Single-Response Sequential Method,” *Behavioral Science*, 9, 1964, 226-232.
- Davis, D. B., and Paté-Cornell, M. E., “A Challenge to the Compound Lottery Axiom: A Two-stage Normative Structure and Comparison to Other Theories,” *Theory and Decision*, 37, 1994, 267-309.
- Eeckhoudt, L.; Gollier, C., and Schlesinger, H., “Changes in Background Risk and Risk Taking Behavior,” *Econometrica*, 64, 1996, 683-689.
- Ellsberg, D., “Risk, Ambiguity, and the Savage Axioms,” *Quarterly Journal of Economics*, 75, 1961, 643-669.
- Ergin, H. and Gul, F., “A Theory of Subjective Compound Lotteries,” *Journal of Economic Theory* 144(3), May 2009, 899-929.
- Fiore S.M.; Harrison G.W.; Hughes, C.E., and Rutström E. E., “Virtual Experiments and Environmental Policy,” *Journal of Environmental Economics and Management*, 57, 2009, 65-86.
- Ghirardoto, P.; Maccheroni, F., and Marinacci, M., “Differentiating Ambiguity and Ambiguity Attitude,” *Journal of Economic Theory*, 118, 2004, 133-173.
- Gilboa, I.; Postlewaite, A.W., and Schmeidler, D., “Probability and Uncertainty in Economic Modeling,” *Journal of Economic Perspectives*, 22, 2008, 173-188.
- Gilboa, I., and Schmeidler, D., “Maxmin Expected Utility with a Non-Unique Prior,” *Journal of Mathematical Economics*, 18, 1989, 141-153.
- Gollier, C., and Pratt, J. W., “Risk Vulnerability and the Tempering Effect of Background Risk,” *Econometrica*, 64, 1996, 1109-1123.
- Gonzalez, R., and Wu, G., “On the Shape of the Probability Weighting Function,” *Cognitive Psychology*, 38, 1999, 129-166.
- Grant, S., “Subjective Probability Without Monotonicity: or How Machina's Mom May Also be Probabilistically Sophisticated,” *Econometrica*, 63, January 1995, 159-189.
- Grether, D.M., “Testing Bayes’ Rule and the Representativeness Heuristic: Some Experimental Evidence,” *Journal of Economic Behavior and Organization*, 17, 1992, 31-57.

- Griffin, D., and Tversky A., "The Weighing of Evidence and the Determinants of Confidence," *Cognitive Psychology*, 24, 1992, 411-435.
- Gul, F., "A Theory of Disappointment Aversion," *Econometrica*, 59, 1991, 667-686.
- Harrison, G.W., "Theory and Misbehavior of First-Price Auctions: Reply," *American Economic Review*, 82, 1992, 1426-1443.
- Harrison, G. W.; Johnson, E.; McInnes, M. M., and Rutström, E. E., "Risk Aversion and Incentive Effects: Commensurate," *American Economic Review*, 95, 2005, 897-901.
- Harrison, G.W.; Lau, M.I., and Williams, M.B., "Estimating Individual Discount Rates for Denmark: A Field Experiment," *American Economic Review*, 92, 2002, 1606-1617.
- Harrison, G.W., and Rutström E. E., "Risk Aversion in the Laboratory," in J.C. Cox and G.W. Harrison (eds.), *Risk Aversion in Experiments* (Bingley, U.K: Emerald, Research in Experimental Economics, Vol. 12, 41-197, 2008a).
- Hey, J.D., and Orme C., "Investigating Generalizations of Expected Utility Theory Using Experimental Data," *Econometrica*, 62, 1994, 1291-1326.
- Holt, C.A., and Laury, S.K., "Risk Aversion and Incentive Effects," *American Economic Review*, 92, 2002, 1644-1655.
- Holt, C.A., and Smith, A.M., "An Update on Bayesian Updating," *Journal of Economic Behavior & Organization*, 69, 2009, 125-134.
- Klibanoff, P.; Marinacci, M., and Mukerji, S., "A Smooth Model of Decision Making Under Ambiguity," *Econometrica*, 73, 2005, 1849-1892.
- Machina, M.J., and Schmeidler, D., "A More Robust Definition of Subjective Probability," *Econometrica*, 60, 1992, 745-780.
- Machina, M.J., and Schmeidler, D., "Bayes without Bernoulli: Simple Conditions for Probabilistically Sophisticated Choice," *Journal of Economic Theory*, 67, 1995, 106-128.
- Nau, R. F., "Uncertainty Aversion with Second-order Utilities and Probabilities," *Management Science*, 52, 2006, 136-156.
- Neilson, W. S., "A Simplified Axiomatic Approach to Ambiguity Aversion," *Journal of Risk and Uncertainty* 41, 2010, 113-124.
- Plott, C.R., and Zeiler, K., "The Willingness to Pay-Willingness to Accept Gap, the 'Endowment Effect,' Subject Misconceptions, and Experimental Procedures for Eliciting Valuations," *American Economic Review*, 95, 2005, 530-545.
- Prelec, D., "The Probability Weighting Function," *Econometrica*, 66, 1998, 497-527.

- Quiggin, J., "A Theory of Anticipated Utility," *Journal of Economic Behavior & Organization*, 3, 1982, 323-343.
- Quiggin, J., "Background Risk in Generalized Expected Utility Theory," *Economic Theory*, 22, 2003, 607-611.
- Rieger, M.O., and Wang, M., "Cumulative Prospect Theory and the St. Petersburg Paradox," *Economic Theory*, 28, 2006, 665-679.
- Rutström, E.E., "Home-Grown Values and the Design of Incentive Compatible Auctions," *International Journal of Game Theory*, 27, 1998, 427-441.
- Segal, U., "The Ellsberg Paradox and Risk Aversion: An Anticipated Utility Approach," *International Economic Review*, 28, 1987, 175-202.
- Segal, U., "Does the Preference Reversal Phenomenon Necessarily Contradict the Independence Axiom?" *American Economic Review*, 78, 1988, 233-236.
- Segal, U., "Two-Stage Lotteries Without the Reduction Axiom," *Econometrica*, 58, 1990, 349-377.
- Segal, U., "The Independence Axiom Versus the Reduction Axiom: Must We Have Both?" in W. Edwards (ed.), *Utility Theories: Measurements and Applications* (Boston: Kluwer Academic Publishers, 1992).
- Savage, L.J., "Elicitation of Personal Probabilities and Expectations," *Journal of American Statistical Association*, 66, 1971, 783-801.
- Savage, L.J., *The Foundations of Statistics* (Dover: New York, Second Edition, 1972).
- Tversky, A., and Kahneman, D., "Advances in Prospect Theory: Cumulative Representations of Uncertainty," *Journal of Risk & Uncertainty*, 5, 1992, 297-323.
- Wilcox, N.T., "Predicting Individual Risky Choices Out-of-Context: A Critical Stochastic Modeling Primer and Monte Carlo Study," in J. Cox and G.W. Harrison (eds.), *Risk Aversion in Experiments* (Bingley, UK: Emerald, Research in Experimental Economics, Volume 12, 2008).
- Wilcox, N.T., "'Stochastically More Risk Averse': A Contextual Theory of Stochastic Discrete Choice Under Risk," *Journal of Econometrics*, 162, 2011, 89-104.

Figure 1: Typical Lottery Choice Task




Table 1: Decision Sheet in Betting Task

Make a bet with all bookies.

Bookie	Stake	Odds offered		Earnings including the stake of £3		I will bet on (circle)	
		White	Blue	White	Blue		
1	£3	20.00	1.05	£60.00	£3.15	W	B
2	£3	10.00	1.11	£30.00	£3.33	W	B
3	£3	6.67	1.18	£20.00	£3.54	W	B
4	£3	5.00	1.25	£15.00	£3.75	W	B
5	£3	4.00	1.33	£12.00	£4.00	W	B
6	£3	3.33	1.43	£10.00	£4.29	W	B
7	£3	2.86	1.54	£8.58	£4.62	W	B
8	£3	2.50	1.67	£7.50	£5.00	W	B
9	£3	2.22	1.82	£6.66	£5.46	W	B
10	£3	2.00	2.00	£6.00	£6.00	W	B
11	£3	1.82	2.22	£5.46	£6.66	W	B
12	£3	1.67	2.50	£5.00	£7.50	W	B
13	£3	1.54	2.86	£4.62	£8.58	W	B
14	£3	1.43	3.33	£4.29	£10.00	W	B
15	£3	1.33	4.00	£4.00	£12.00	W	B
16	£3	1.25	5.00	£3.75	£15.00	W	B
17	£3	1.18	6.67	£3.54	£20.00	W	B
18	£3	1.11	10.00	£3.33	£30.00	W	B
19	£3	1.05	20.00	£3.15	£60.00	W	B

Table 2: Experimental Design Parameters

Posterior Probability of White	Total Dice (N)	Actual White Faces (w)	Actual Blue Faces (b)	Number of Observations
0.12	5	0	5	72
	9	2	7	110
	17	6	11	47
0.23	3	0	3	77
	5	1	4	218
	9	3	6	112
	17	7	10	106
0.4	3	1	2	110
	5	2	3	412
	9	4	5	117
	17	8	9	82
0.6	3	2	1	154
	5	3	2	394
	9	5	4	87
	17	9	8	44
0.77	3	3	0	55
	5	4	1	227
	9	6	3	95
	17	10	7	73
0.88	5	5	0	45
	7	7	2	53
	17	11	6	81

Figure 2: Estimates of Subjective Probability under SEU

Black is 0.60 posterior, Blue is 0.77 posterior and Red is 0.88 posterior




Figure 3: Estimates of Subjective Probability under SEU Assuming Risk Neutrality

Black is 0.60 posterior, Blue is 0.77 posterior and Red is 0.88 posterior




Figure 4: Estimates of Subjective Probability Under SEU: The Effect of Risk Aversion
Solid line is general SEU model, and dashed line is SEU assuming risk neutrality
Black is 0.60 posterior, blue is 0.77 posterior and red is 0.88 posterior




Figure 5: Possible Probability Weighting and Decision Weights
 Tversky & Kahneman (1992) probability weighting function





Figure 6: Actual Probability Weighting and Decision Weights
 Reiger & Wang (2006) probability weighting function



Appendix 1: Experimental Instructions (NOT FOR PUBLICATION)

The exact layout of the instructions can be viewed by accessing machine-readable versions available from the authors. A horizontal line indicates a page break in the original instructions.

A. Risk Attitudes Task

Stage 2: INSTRUCTIONS

We will now continue with Stage 2 of the experiment.

This stage is about choosing between lotteries with varying prizes and chances of winning. You will be shown 20 lottery pairs, and from each pair you will choose the lottery you prefer. You will actually get the chance to play one of the lotteries you choose, and will be paid according to the outcome of that lottery, so you should think carefully about your preferences.

On the accompanying sheet there is an example lottery pair.

The outcome of the lotteries will be determined by the roll of a 100-sided die that is numbered from 1 to 100. The numbers that will determine each outcome are shown below each lottery.

In the example the left lottery pays five pounds (£5) if the number is between 1 and 40 (a 40% chance), and it pays fifteen pounds (£15) if it is between 41 and 100 (a 60% chance).

The lottery on the right pays five pounds (£5) if the number drawn is between 1 and 50 (a 50% chance), ten pounds (£10) if the number is between 51 and 90 (a 40% chance), and fifteen pounds (£15) if the number is between 91 and 100 (a 10% chance).

The size of the pie slices represent the chances of earning each payoff.

Each lottery pair will be shown on a separate sheet of paper. On each sheet you should indicate your preferred lottery by ticking the appropriate box. After you have worked through all the lottery pairs, please raise your hand.

You will then roll a 20-sided die to determine which pair of lotteries will be played out, and the 100-sided die to determine the outcome of the chosen lottery.

For instance, suppose the lottery on the accompanying page was chosen to pay off and you rolled a 42 on the 100-sided die. If you had picked the lottery on the left you would win £15, while if you had picked the lottery on the right you would have won £5.

Therefore, your payoff is determined by three things:

- which lottery pair is chosen to be played out using the 20-sided die;
- which lottery you selected, the left or the right, for the chosen lottery pair;
- the outcome of that lottery when you roll the 100-sided die.

This is not a test of whether you can pick the best lottery in each pair, because none of the

lotteries are necessarily better than the others. Which lotteries you prefer is a matter of personal taste.

Please work silently, and think carefully about each choice.

All payoffs are in cash, and are in addition to the £5 show-up fee that you receive just for being here.

B. Betting Task

Stage 3: INSTRUCTIONS

What you will do

In this stage of the experiment you will be betting on the outcomes of uncertain events. Usually we bet on events like football matches or elections, but in this task the events will be random choices made by the experimenter between two boxes, one blue and the other white. The experimenter will not tell you which box was chosen. At the start each box will have the same chance of being chosen, but once it has been chosen the experimenter will give you some information to help you work out the chances that it was blue or white. Armed with this information, you will make bets on which box was chosen.


The procedure, which is summarized on the accompanying picture, is as follows. The experimenter will first choose the box by rolling a 6-sided die with three blue and three white sides. If blue comes up he will choose the blue box, if white comes up he will choose the white one.

Both the white and blue boxes contain several dice, each having 10 sides. Both boxes have the same number of dice, which will vary over the course of the experiment. The dice in the blue box always have 6 blue sides and 4 white ones, while those in the white box have 4 blue sides and 6 white ones.

The experimenter will roll all the dice in the chosen box and tell you how many blue and white sides came up. He will not tell you which box was chosen.

Because the dice in the blue box have more blue sides than those in the white box, knowing the number of blue and white sides that come up can help you work out the chances that each box was chosen. For example, if more blue sides come up this means it is more likely to be the blue box, and if more white sides come up it is more likely to be the white box.

Once you have the information about the dice rolls, you will then make bets on which box was chosen.



About betting

You will be making bets with several betting houses or “bookies,” just as you might bet on a football game or a horse race.

To familiarize you with betting, we will illustrate how it works with the example of a horse race.

Imagine a two horse race between Blue Bird and White Heat. Several bookies offer different odds for both horses. The table below shows the odds offered by three bookies along with the amounts they would pay if you staked £10 on the winning horse. The earnings are calculated by multiplying the odds by the stake. In this experiment you will be making bets on which box was chosen using a table like this. **At this point you should take some time to study the table.**

Bookie	Stake	Odds offered		Earnings including the stake of £10	
		Blue Bird	White Heat	Blue Bird	White Heat
A	£10	5.00	1.25	£50.00	£12.50
B	£10	3.33	1.43	£33.33	£14.30
C	£10	2.00	2.00	£20.00	£20.00

Below are three important points about betting.

- 1. Your belief about the chances of each outcome is a personal judgment that depends on information you have about the different events.** For the horse race, you may have seen previous races or read articles about them. In the experiment the information you have about whether the blue or white box was chosen will be how many blue and white faces came up.
- 2. Even if you believe Event X is more likely to occur than Event Y, you may want to bet on Y because you find the odds attractive.** For example, even if you believe White Heat is most likely to win you may want to bet on Blue Bird because you find the odds attractive. To illustrate, suppose you personally believe that Blue Bird has a 40% chance of winning and White Heat has a 60% chance of winning. This means that if you bet £10 on Blue Bird with Bookie A you believe there is a 40% chance of receiving £50.00 and a 60% chance of receiving nothing. You may find this more attractive than betting on White Heat, which you believe offers a 60% chance of 12.50 and a 40% chance of nothing.
- 3. Your choices might also depend on your willingness to take risks or to gamble.** There is no right choice for everyone. In a horse race you might want to bet on the longshot since it will bring you more money if it wins, but you also might want to bet on the favorite since it is more likely to win something.

For each bookie, whether you would choose to bet on Blue Bird or White Heat will depend on three things: your judgment about how likely it is each horse will win, the odds

offered by the bookie, and how much you like to gamble or take risks.

Your choices

Now you are familiarized with odds, we can go back to the experimental betting task. Recall that the experimenter will first make a random choice of a blue or white box. Then he will roll the dice in the chosen box and tell you how many white and blue sides came up. Then you will consider the chances that the box chosen was blue or white, and make a series of bets.

You have a booklet of record sheets. Each record sheet shows the bookies you will be dealing with, and the odds they offer. There are 19 bookies on each sheet, and each offer different odds for the two outcomes. **Take a minute to look at one such record sheet, shown on the next page.**

There will be 30 separate events, and 19 bookies offer odds for each event. **You will make bets at all 19 bookies for all 30 events.**

For each bet, you have a £3 stake, and the record sheet shows the payoffs you will receive if you bet on the box that was actually chosen.

There is a separate record sheet for each of the 30 events. On each sheet you should circle W or B to indicate the bet you want to make with **all 19 bookies.**

One and only one of the bets in the entire experiment will pay off for real. Therefore, please consider each bet as if it is the only one that will be paid out. After you have placed all your bets, you will roll a 30-sided die to determine which event will be played out, and a 20-sided die to determine which bookie will determine your earnings.

All payoffs are in cash, and are in addition to the £5 show-up fee that you receive just for being here.

C. Demographic Questionnaire

In this survey most of the questions asked are descriptive. We will not be grading your answers and your responses are completely confidential. Please think carefully about each question and give your best answers.

1. What is your age? _____ years
2. What is your sex? (Circle one number.)
01 Male 02 Female

3. Which of the following categories best describes you? (Circle one number.)

- 01 British
- 02 Irish
- 03 Any other white background
- 04 White and Black Caribbean
- 05 White and Black African
- 06 White and Asian
- 07 Any other mixed background
- 08 Indian
- 09 Pakistani
- 10 Bangladeshi
- 11 Any other Asian background
- 12 Caribbean
- 13 African
- 14 Any other Black background
- 15 Chinese
- 16 Any other ethnic group
- 17 Not stated
- 18 Prefer not to say

4. What is your main field of study? (Circle one number.)

- 01 Accounting
- 02 Economics
- 03 Finance
- 04 Business Administration, other than Accounting, Economics, or Finance
- 05 Education
- 06 Engineering
- 07 Health and Medicine
- 08 Biological and Biomedical Sciences
- 09 Math, Computer Sciences, or Physical Sciences
- 10 Social Sciences or History
- 11 Law
- 12 Psychology
- 13 Modern Languages and Cultures
- 14 Other Fields

5. What is your year of studies? (Circle one number.)

- | | |
|----------------|-------------|
| 01 First year | 04 Masters |
| 02 Second year | 05 Doctoral |
| 03 Third year | |

6. What is the highest level of education you expect to complete? (Circle one number)

- 01 Bachelor's degree
- 02 Master's degree

- 03 Doctoral degree
- 04 Professional qualification

7. As a percentage, what is your current average mark if you are doing a Bachelor's degree, or what was it when you did a Bachelor's degree? This mark should refer to all your years of study for this degree, not just the current year. Please pick one by rounding up or down to the nearest number:

- 01 Above 70%
- 02 Between 60 - 69%
- 03 Between 50 - 59%
- 04 Between 40 - 49%
- 05 Less than 40%
- 06 Have not taken courses for which grades are given.

8. What is your citizenship status?

- 01 British Citizen
- 02 EU Citizen (non-British Citizen)
- 03 Non-EU Citizen

9. Are you currently:

- 01 Single and never married?
- 02 Married?
- 03 Separated, divorced or widowed?

10. How many people live in your household? Include yourself, your spouse and any dependents. Do not include your parents or roommates unless you claim them as dependents.

11. Please circle the category below that describes the total amount of income before tax earned in the calendar year 2007 by the people in your household (as "household" is defined in question 10).

[Consider all forms of income, including salaries, tips, interest and dividend payments, scholarship support, student loans, parental support, social security, alimony, and child support, and others.]

- 01 Less than £10,000
- 02 £10,000 – £19,999
- 03 £20,000 – £29,999
- 04 £30,000 - £49,999
- 05 Over £50,000

12. Please circle the category below that describes the total amount of income before tax earned in the calendar year 2007 by your parents.

[Consider all forms of income, including salaries, tips, interest and dividend payments, social security, alimony, and child support, and others.]

- 01 Less than £10,000

- 02 £10,000 – £19,999
- 03 £20,000 – £29,999
- 04 £30,000 - £49,999
- 05 Over £50,000
- 06 Don't Know

13. Do you currently smoke cigarettes? (Circle one number.)

- 00 No
- 01 Yes

If yes, approximately how much do you smoke in one day? _____ cigarettes.

Appendix 2: Derivations (NOT FOR PUBLICATION)

The posterior probability of the W outcome given the sample information is given by

$$\pi = p(W | \text{sample}) = p(\text{sample} | W) \times p(W) / p(\text{sample}),$$

and the posterior probability of the B outcome given the sample information is given by

$$1 - \pi = 1 - p(W | \text{sample}) = p(B | \text{sample}) = p(\text{sample} | B) \times p(B) / p(\text{sample}).$$

The prior probabilities of getting the W and B outcomes are $p(W) = p(B) = 0.5$. The conditional probability of getting a specific sample given the W outcome follows a binomial probability distribution and is

$$p(\text{sample} | W) = [N! / (w! (N-w)!)] p(w)^w (1-p(w))^{(N-w)},$$

where $p(w)$ is the probability of the w outcome given the W outcome. The conditional probability of getting a specific sample given the B outcome is

$$p(\text{sample} | B) = [N! / (b! (N-b)!)] p(b)^b (1-p(b))^{(N-b)},$$

where $p(b)$ is the probability of the b outcome given the B outcome. The probability of getting the sample information is:

$$p(\text{sample}) = 0.5 \times p(\text{sample} | W) + 0.5 \times p(\text{sample} | B).$$

The posterior probabilities of the W and B outcomes given the sample information are then:

$$\pi = [N! / (w! (N-w)!)] p(w)^w (1-p(w))^{(N-w)} /$$

$$[N! / (w! (N-w)!)] p(w)^w (1-p(w))^{(N-w)} + [N! / (b! (N-b)!)] p(b)^b (1-p(b))^{(N-b)}]$$

and

$$1 - \pi = [N! / (b! (N-b)!)] p(b)^b (1-p(b))^{(N-b)} /$$

$$[N! / (w! (N-w)!)] p(w)^w (1-p(w))^{(N-w)} + [N! / (b! (N-b)!)] p(b)^b (1-p(b))^{(N-b)}].$$

Appendix 3: Detailed Estimation Results (NOT FOR PUBLICATION)

Apologies for these tables being “in *Stata*,” but the format is reasonably standard. Estimates of the behavioral error terms use a log-transform to ensure that the parameter is non-negative; the delta method is used under Table A1 to illustrate how one should infer the correct underlying parameter. Similarly, the subjective probabilities in Tables A2 and A3 are constrained to lie in the unit interval, using the transform $\pi = 1/(1+\exp(\kappa))$, where κ is the parameter estimated and π is the inferred probability. Thus the numerical algorithm finding the maximum likelihood estimates can vary κ between $\pm\infty$ to evaluate numerical derivatives, while π is constrained to lie in the unit interval. Again, the delta method can be used to infer point estimates and standard errors for π from estimates of κ . Finally, the number of observations in Tables A2, A3 and A4 are inflated because we employ “frequency weights” of 50 for every observed choice from the risk task: there are actually 2,220 choices, as shown in Table A1, but these appear then to be $50 \times 2,220 = 111,000$ observations. These weights ensure that the estimated risk parameters are based primarily on the choices from the risk tasks. The number of observations in Tables A2 and A3 is greater because their analysis contains the data from both the risk and the belief task.

Table A1: Risk Attitudes, Estimated only on Choices from Risk Task

					Number of obs	=	2220	
					Wald chi2(8)	=	24.10	
Log pseudolikelihood = -1335.7704					Prob > chi2	=	0.0022	
(Std. Err. adjusted for 111 clusters in id)								

			Robust					
		Coef.	Std. Err.	z	P> z	[95% Conf. Interval]		

r								
	RA_then_B		-.1825245	.0680085	-2.68	0.007	-.3158187	-.0492303
	female		.18269	.0745464	2.45	0.014	.0365817	.3287983
	teenager		-.0916832	.07621	-1.20	0.229	-.241052	.0576857
	white		.2132524	.1225167	1.74	0.082	-.0268759	.4533807
	brit		-.2231895	.1430793	-1.56	0.119	-.5036198	.0572407
	income_high		.1324355	.0660586	2.00	0.045	.0029631	.261908
	graduate		-.1003567	.097986	-1.02	0.306	-.2924058	.0916923
	math		-.1108413	.081391	-1.36	0.173	-.2703647	.0486821
	_cons		.5896646	.1104044	5.34	0.000	.373276	.8060531

LNmuRA								
	_cons		-2.106163	.0872657	-24.14	0.000	-2.277201	-1.935126

Table A2: Subjective Expected Utility Model

Log pseudolikelihood = -83967.476

Number of obs = 163649
Wald chi2(8) = 40.75
Prob > chi2 = 0.0000

(Std. Err. adjusted for 111 clusters in id)

		Coef.	Robust Std. Err.	z	P> z	[95% Conf. Interval]	
r							
	RA_then_B	-.1228562	.0562522	-2.18	0.029	-.2331085	-.012604
	female	.1356762	.0603764	2.25	0.025	.0173406	.2540117
	teenager	-.0501637	.0587284	-0.85	0.393	-.1652693	.064942
	white	.1370246	.0795495	1.72	0.085	-.0188895	.2929388
	brit	-.2015224	.1008559	-2.00	0.046	-.3991962	-.0038486
	income_high	.0605295	.0564305	1.07	0.283	-.0500723	.1711313
	graduate	-.0340874	.0781271	-0.44	0.663	-.1872136	.1190388
	math	-.1036242	.0602806	-1.72	0.086	-.2217719	.0145235
	_cons	.6124987	.0774014	7.91	0.000	.4607946	.7642027
sprob1_							
	RA_then_B	.0219064	.0409056	0.54	0.592	-.0582672	.10208
	N_ascending	.0569943	.029496	1.93	0.053	-.0008168	.1148053
	female	.067758	.0342876	1.98	0.048	.0005555	.1349604
	teenager	.0349921	.0376387	0.93	0.353	-.0387784	.1087625
	white	.0594294	.0469789	1.27	0.206	-.0326476	.1515063
	brit	-.0515731	.0592538	-0.87	0.384	-.1677085	.0645622
	income_high	.0274503	.0359471	0.76	0.445	-.0430046	.0979053
	graduate	.0621187	.0587186	1.06	0.290	-.0529677	.1772051
	math	.0071615	.0419263	0.17	0.864	-.0750125	.0893355
	_cons	-.3617571	.0665444	-5.44	0.000	-.4921816	-.2313325
sprob2_							
	RA_then_B	-.029905	.0790267	-0.38	0.705	-.1847945	.1249845
	N_ascending	.0905342	.0575997	1.57	0.116	-.0223592	.2034277
	female	.0608162	.070985	0.86	0.392	-.0783119	.1999442
	teenager	-.0127714	.0842156	-0.15	0.879	-.177831	.1522882
	white	.1131135	.1018507	1.11	0.267	-.0865101	.3127372
	brit	-.2838021	.121042	-2.34	0.019	-.5210401	-.046564
	income_high	.0126633	.0794306	0.16	0.873	-.1430178	.1683445
	graduate	-.0334664	.100381	-0.33	0.739	-.2302096	.1632767
	math	-.0139351	.08934	-0.16	0.876	-.1890383	.161168
	_cons	-.4708455	.1322005	-3.56	0.000	-.7299537	-.2117373
sprob3_							
	RA_then_B	.0463894	.1224547	0.38	0.705	-.1936173	.2863962
	N_ascending	.001102	.096178	0.01	0.991	-.1874033	.1896074
	female	.1389314	.1119441	1.24	0.215	-.080475	.3583377
	teenager	-.0332965	.1085275	-0.31	0.759	-.2460065	.1794134
	white	.1880194	.1815643	1.04	0.300	-.16784	.5438789
	brit	-.2960438	.2121002	-1.40	0.163	-.7117525	.1196649
	income_high	-.0157023	.1366839	-0.11	0.909	-.2835979	.2521932
	graduate	-.0864282	.1893455	-0.46	0.648	-.4575386	.2846822
	math	-.003462	.1379542	-0.03	0.980	-.2738473	.2669233
	_cons	-.7170649	.2266169	-3.16	0.002	-1.161226	-.2729039
LNmuRA							
	_cons	-2.160693	.0832332	-25.96	0.000	-2.323827	-1.997559
LNmuB							
	_cons	-3.38979	.1201996	-28.20	0.000	-3.625377	-3.154204

Table A3: Subjective Expected Utility Model Assuming Risk Neutrality

				Number of obs	=	163649
				Wald chi2(0)	=	.
Log pseudolikelihood = -92280.438				Prob > chi2	=	.
(1) [r]_cons = .001						
(2) [LNmuRA]_cons = 0						
(Std. Err. adjusted for 111 clusters in id)						

		Coef.	Robust Std. Err.	z	P> z	[95% Conf. Interval]

r						
	_cons	.001	.	.	.	.

sprob1_						
	RA_then_B	.0963775	.0424267	2.27	0.023	.0132226 .1795323
	N_ascending	.0873922	.0426211	2.05	0.040	.0038564 .170928
	female	.0288932	.0394497	0.73	0.464	-.0484268 .1062133
	teenager	.0890286	.044747	1.99	0.047	.0013261 .1767312
	white	.0087767	.0719191	0.12	0.903	-.1321822 .1497355
	brit	.0160792	.0828602	0.19	0.846	-.1463237 .1784821
	income_high	-.0149373	.0436778	-0.34	0.732	-.1005442 .0706697
	graduate	.1406618	.0646784	2.17	0.030	.0138944 .2674291
	math	.0450672	.045673	0.99	0.324	-.0444502 .1345846
	_cons	-.6013873	.073373	-8.20	0.000	-.7451956 -.4575789

sprob2_						
	RA_then_B	.1234993	.1097776	1.12	0.261	-.091661 .3386595
	N_ascending	.1555636	.095968	1.62	0.105	-.0325301 .3436574
	female	-.0807038	.119323	-0.68	0.499	-.3145726 .153165
	teenager	.0722725	.1196011	0.60	0.546	-.1621412 .3066863
	white	-.0715795	.1800304	-0.40	0.691	-.4244326 .2812735
	brit	-.2557715	.2171422	-1.18	0.239	-.6813624 .1698194
	income_high	-.1212199	.0993021	-1.22	0.222	-.3158485 .0734087
	graduate	-.0060656	.1683454	-0.04	0.971	-.3360166 .3238853
	math	.1142384	.1456307	0.78	0.433	-.1711925 .3996693
	_cons	-.8389937	.2400569	-3.49	0.000	-1.309497 -.3684908

sprob3_						
	RA_then_B	.399608	.2611878	1.53	0.126	-.1123107 .9115266
	N_ascending	-.0146083	.2044517	-0.07	0.943	-.4153264 .3861097
	female	-.0412008	.3041888	-0.14	0.892	-.6373999 .5549984
	teenager	.1054203	.1991565	0.53	0.597	-.2849192 .4957599
	white	-.0599099	.3763651	-0.16	0.874	-.797572 .6777522
	brit	-.2222191	.4483063	-0.50	0.620	-1.100883 .6564451
	income_high	-.2307703	.2397852	-0.96	0.336	-.7007408 .2392001
	graduate	-.0780133	.3695222	-0.21	0.833	-.8022634 .6462368
	math	.2055527	.2809535	0.73	0.464	-.345106 .7562114
	_cons	-1.335423	.4636354	-2.88	0.004	-2.244132 -.4267146

LNmuB						
	_cons	-1.46519	.0761123	-19.25	0.000	-1.614367 -1.316012