

COLEMAN FUNG
RISK MANAGEMENT
RESEARCH CENTER

**University of California
Berkeley**

Estimating Ambiguity Aversion in a Portfolio Choice Experiment*

David Ahn[†] Syngjoo Choi[‡] Douglas Gale[§] Shachar Kariv[¶]
UC Berkeley UCL NYU UC Berkeley

February 27, 2009

Abstract

We report a laboratory experiment that enables us to estimate parametric models of ambiguity aversion at the level of the individual subject. We use two main specifications, a “kinked” specification that nests Maxmin Expected Utility, Choquet Expected Utility, α -Maxmin Expected Utility, and Contraction Expected Utility and a “smooth” specification that nests the various theories referred to collectively as Recursive Expected Utility. Our subjects solved a series of portfolio-choice problems. The assets are Arrow securities corresponding to three states of nature, where the probability of one state is known and the remaining two are ambiguous. The sample exhibits

*The experiment was conducted at the Experimental Social Science Laboratory (Xlab) at the University of California, Berkeley. We thank Raymond Fisman for detailed comments and suggestions. We are also grateful to Yoram Halevy, Tom Palfrey, Chris Shannon, and Bill Zame for helpful discussions. This paper has also benefited from suggestions by the participants of seminars at Brown, UC Berkeley, UCL, UCLA, and UCSD. We acknowledge the National Science Foundation Grant No. SES-0550224 (Ahn) and SES-0617955 (Gale and Kariv), and the Coleman Fung Risk Management Research Center (OpenLink Fund) at UC Berkeley (Kariv) for financial support.

[†]Department of Economics, University of California, Berkeley, 508-1 Evans Hall # 3880, Berkeley, CA 94720 (E-mail: dahn@econ.berkeley.edu, URL: <http://www.econ.berkeley.edu/~dahn/>).

[‡]Department of Economics, University College London, Gower Street, London WC1E 6BT, UK (Email: syngjoo.choi@ucl.ac.uk, URL: <http://www.homepages.ucl.ac.uk/~uctpsc0>).

[§]Department of Economics, New York University, 19 W. 4th Street, New York, NY, 10012 (E-mail: douglas.gale@nyu.edu, URL: <http://www.econ.nyu.edu/user/galed>).

[¶]Department of Economics, University of California, Berkeley, 508-1 Evans Hall # 3880, Berkeley, CA 94720 (E-mail: kariv@berkeley.edu, URL: <http://econ.berkeley.edu/~kariv/>).

considerable heterogeneity in preferences, as captured by parameter estimates. Nonetheless, there exists a strong tendency to equalize the demands for the securities that pay off in the ambiguous states, a feature more easily accommodated by the kinked specification than by the smooth specification. We also find that a large number of subjects are well described by the ambiguity-neutral Subjective Expected Utility model.

JEL Classification Numbers: D81, C91.

Key Words: uncertainty, ambiguity aversion, Subjective Expected Utility, Maxmin Expected Utility, α -Maxmin Expected Utility, Choquet Expected Utility, Contraction Expected Utility, Recursive Expected Utility, experiment.

1 Introduction

In Savage's (1954) celebrated theory of Subjective Expected Utility (SEU), an individual acts as if a single probability measure governs uncertainty over states of the world. Ellsberg (1961) proposed a thought experiment in which aversion to ambiguity would lead to a violation of the Savage axioms. Subsequent experimental work has repeatedly and robustly confirmed Ellsberg's conjecture. Meanwhile, a large theoretical literature has developed models consistent with this behavior.

In this paper, we present a new experimental data set and use it to estimate parametric models of ambiguity averse behavior. The experimental data is generated by subjects solving a series of randomly generated portfolio-choice problems. In our preferred interpretation, there are three states of nature, denoted by $s = 1, 2, 3$. For each state s , there is an Arrow security that pays one dollar in state s and nothing in the other states. To distinguish the effects of risk (known probabilities) and ambiguity (unknown probabilities), state 2 is assigned an objectively known probability, whereas states 1 and 3 have ambiguous probabilities.

More precisely, subjects are informed that state 2 occurs with probability $\pi_2 = \frac{1}{3}$, whereas states 1 and 3 occur with unknown probabilities $\pi_1 \geq 0$ and $\pi_3 \geq 0$, satisfying $\pi_1 + \pi_3 = \frac{2}{3}$. By letting x_s denote the demand for the security that pays off in state s and p_s denote its price, the budget constraint can be written as $\mathbf{p} \cdot \mathbf{x} = 1$, where $\mathbf{x} = (x_1, x_2, x_3)$ and $\mathbf{p} = (p_1, p_2, p_3)$. Then the subject can choose any non-negative portfolio $\mathbf{x} \geq \mathbf{0}$ satisfying the budget constraint.

Each budget set defines a corresponding portfolio-choice problem. The budget sets are displayed on a computer screen using the graphical interface introduced by Choi, et al. (2007a) and exploited by Choi, et al. (2007b)

for the study of risky decisions. The data generated in this way allows us to estimate models of ambiguity aversion for each individual subject. Estimation at the individual level is crucial because of the possibility of individual heterogeneity.

There is a variety of theoretical models of attitudes toward risk and ambiguity, but they all give rise to one of two main specifications. The first is a “kinked” specification, which can be rationalized by different utility models in the literature, including Maxmin Expected Utility (MEU), Choquet Expected Utility (CEU), α -Maxmin Expected Utility (α -MEU), or Contraction Expected Utility.¹ The second is a “smooth” specification, based on the class of Recursive Expected Utility (REU) models.² The standard SEU model and the MEU model (with a maximal set of prior beliefs) are special cases of both the kinked and smooth specifications.

Each of the two specifications, kinked and smooth, is characterized by two parameters, one of which is associated with risk aversion and the other with ambiguity aversion. The estimated parameters provide summary statistics of attitudes to risk and ambiguity. However, to achieve this we have to adopt a parsimonious parameterization. Once this information is succinctly summarized, we can compare risk and ambiguity attitudes across subjects through their estimated individual parameters and get a broad picture of the heterogeneity of preferences. For any given model, these parameter estimates range from risk neutrality with ambiguity aversion, to ambiguity neutrality with risk aversion, to infinite risk aversion.

Although individual preferences are heterogeneous, about two-thirds of our subjects have a positive degree of ambiguity aversion. In other words, one-third of subjects are well described by the standard SEU model, a much higher proportion than has been found in previous studies.³ Naive tests of significance suggest that nearly half of subjects may be consistent with SEU. Our estimates of risk aversion are similar to other recent estimates.

Unlike urn-based studies, in which the exposure to ambiguity is fixed by the experimenter, in our design subjects can reduce their exposure to ambiguity by choosing portfolios whose payoffs are less dependent on the

¹See, MEU: Gilboa and Schmeidler (1989); CEU: Schmeidler (1989); α -MEU: Ghirardato et al. (2004) and Olszewski (2006); Contraction Expected Utility: Gajdos et al. (2008).

²See, Ergin and Gul (2004), Klibanoff et al. (2005), Nau (2005), Seo (2007), and related work by Halevy and Feltkamp (2005), Giraud (2006) and Ahn (2008).

³In variants of Ellsberg reported in Camerer (1995), subjects regularly pick the less ambiguous urn and pay high premium to avoid ambiguity – around 10-20 percent of expected value.

ambiguous states. In the limit, when $x_1 = x_3$, there is no effective exposure to ambiguity. We find there is a strong tendency for subjects to equalize their demands for the securities that pay off in the ambiguous states. This feature of the data can be rationalized by the kinked specification, but is harder to reconcile with the smooth specification.

To our knowledge, this is the first study to simultaneously estimate different models of ambiguity preferences using experimental data. No model appears to be a “winner” across all subjects, but some do better than others. The absolute and relative abilities of the models to fit observed individual data vary across subjects, suggesting that a variety of models may be needed to explain the different choices patterns in the population. However, some patterns appear consistently across many subjects. In particular, the tendency of subjects to hedge ambiguity, i.e., to select portfolios that equalize the demands for the securities that pay off in the ambiguous states, is more easily explained by the kinked specification.

The rest of the paper is organized as follows. Section 2 provides a discussion of some related literature. The experimental design and procedures are described in Section 3. Section 4 provides descriptive statistics, some case studies of individual subjects, and preliminary consistency tests using revealed preference analysis. Section 5 describes the different theoretical models that lie behind the two main specifications that we estimate. Section 6 contains the econometric analysis and Section 7 contains some concluding remarks. The experimental instructions and individual-level data are contained in online appendices.

2 Related Literature

We will not attempt to review the large and growing experimental literature on ambiguity aversion. Camerer and Weber (1992) and Camerer (1995) provide excellent, though now somewhat dated, surveys that the reader may wish to consult. Instead, we focus attention on some recent papers that are particularly relevant to our study.

Halevy (2007) presents a cleverly designed experiment that allows him to distinguish between four models of ambiguity aversion – SEU, MEU or CEU, Recursive Nonexpected Utility (Segal, 1987, 1990), and REU. Subjects are asked their reservation values for four different urns, representing different types of ambiguity (pure risk, pure ambiguity and two types of compound lotteries). The different models of ambiguity aversion generate different predictions about how the urns will be ordered. For each subject,

there will be a unique model that predicts (is consistent with) the subject's reservation values. Halevy (2007) concludes that no single model predicts all the observed behaviors. In fact, all models are represented in the pool of subjects.

Hayashi and Wada (forthcoming) examine attitudes toward imprecise information about uncertainty. Subjects are provided with information in the form of objective restrictions on the probability distribution of states of nature and observe how the subjects' reservation values vary with the objective restrictions on probabilities. Attitudes to imprecise information are obviously closely related to ambiguity aversion. Hayashi and Wada (forthcoming) observe that both the α -MEU model and the contraction model of Gajdos et al. (2008) are systematically violated.

A few recent studies use variants of the Ellsberg urn problems to estimate parameter values or functional forms for individual subjects. Abdellaoui et al. (2008) capture attitudes towards uncertainty and ambiguity by fitting different source functions converting subjective (choice-based) probabilities into willingness to bet. They find considerable heterogeneity in subjects' preferences both in an Ellsberg urn experiment and in experiments using neutrally occurring uncertainties. Hey et al. (2007, 2008) create ambiguity in the laboratory using a Bingo Blower. Unless the number of balls in the Bingo Blower is small, the composition of balls of different colors is missing information. Their results are rather discouraging for the new REU theories.

Finally, Bossaerts, et al. (2008) study the impact of ambiguity and ambiguity aversion on portfolio holdings and asset prices in a financial market experiment. The experimental procedures were adapted from those used by Bossaerts, et al. (2007) to study markets with pure risk. Bossaerts, et al. (2008) point out that there is substantial heterogeneity in ambiguity preferences and that there is a positive correlation between risk aversion and ambiguity aversion.

We share Halevy's (2007) point of view that different models might be needed to describe the behaviors of different subjects, but we go further. In addition to allowing for different models of ambiguity aversion we want to measure the degrees of ambiguity aversion exhibited by subjects who conform to the same model. This last point is particularly important. As the recent evidence shows, individual heterogeneity requires us to study behavior at the individual level in order to properly understand attitudes to risk and ambiguity. For each individual subject, our experimental design allows us to observe a larger number of choices, in a wider variety of settings, than the typical Ellsberg urn-based experiment. As Choi et al. (2007a) emphasize, a choice from a convex budget set provides more information

about preferences than a choice from a discrete set and a larger number of independent observations gives more precise estimates of the parameters of interest.

3 Experimental Design

The experiment was conducted at the Experimental Social Science Laboratory (Xlab) at the University of California, Berkeley under the Xlab Master Human Subjects Protocol. The 154 subjects in the experiment were recruited from all undergraduate classes and staff at UC Berkeley. After subjects read the instructions, the instructions were read aloud by an experimenter. At the end of the instructional period subjects were asked if they had any questions or difficulties understanding the experiment. No subject reported difficulty understanding the procedures or using the computer interface. Each experimental session lasted about one and a half hours. Payoffs were calculated in terms of tokens and then converted into dollars, where each token was worth \$0.50. Earnings were paid in private at the end of the experimental session.

The experimental procedures described below are identical to those described by Choi et al. (2007a) and used by Choi et al. (2007b) to study a portfolio choice problem with two risky assets. Each experimental session consisted of 50 independent decision problems. These decision problems were presented using a graphical interface. On a computer screen, subjects saw a graphical representation of a three-dimensional budget set. An example of one such budget set is illustrated in the experimental instructions reproduced in Appendix I.⁴

There are three axes in the diagram, labeled x , y and z . The axes are scaled from 0 to 100 tokens and are held constant throughout a given experimental session. Each of the axes corresponds to one of three accounts, x , y and z . The subject's decision problem is to select an allocation from the budget set, that is, to allocate his wealth among the three accounts while satisfying the budget constraint. For each round, the computer selected a budget set randomly subject to the constraints that each intercept lies between 0 and 100 tokens and at least one intercept must be greater than 50 tokens. The budget sets selected for each subject in different decision problems were independent of each other and of the sets selected for any of the other subjects in their decision problems.

⁴Online Appendix I: http://emlab.berkeley.edu/~kariv/ACGK_I_A1.pdf.

The resolution compatibility of the budget sets was 0.2 tokens. At the beginning of each decision round, the experimental program dialog window went blank and the entire setup reappeared. The appearance and behavior of the pointer were set to the Windows mouse default and the pointer was automatically repositioned randomly on the budget constraint at the beginning of each decision round. Subjects could use the mouse or the keyboard arrows to move the pointer on the computer screen to the desired allocation. Choices were restricted to allocations on the budget constraint, so that subjects could not violate budget balancedness. Subjects could either left-click or press the Enter key to record their allocation. The process was repeated until all 50 rounds were completed.

Subjects were told that the payoff in each decision round was determined by the number of tokens in each account and that, at the end of each round, the computer would randomly select one of the accounts, x , y or z . Subjects were only informed that account y was with probability $\pi_y = \frac{1}{3}$ and accounts x and z were selected with unknown probabilities π_x and π_z such that $\pi_x + \pi_z = \frac{2}{3}$. In practice, π_x was drawn from the uniform distribution over $[0, \frac{2}{3}]$. This distribution was not announced to the subjects. If the distribution had been revealed to the subjects, the decision problem would have involved compound risk rather than ambiguity.

During the course of the experiment, subjects were not provided with any information about the account that had been selected in each round. Instead, at the end of the experiment, the experimental program randomly selected one decision round from each participant and used that round to determine the subject's payoff. Each round had an equal probability of being chosen, and the subject was paid the amount he had earned in that round. Note that by selecting a single decision round for the payoff we prevent subjects from diversifying their risk across the 50 rounds.

4 Nonparametric Analysis

In this section, we take an initial look at some broad features of the experimental data as a prelude to our estimation of parametric models of ambiguity aversion. We begin with an overview of the basic features of the aggregate data.

4.1 Aggregate behavior

A subject can avoid ambiguity completely by demanding *equal* amounts of the securities that pay off in the ambiguous states $x_1 = x_3$. The resulting

portfolio pays an amount x_2 with probability $\frac{1}{3}$ and an amount $x_1 = x_3$ with probability $\frac{2}{3}$, thus eliminating any ambiguity regarding the probability distribution of payoffs. Similarly, choosing x_1 close to x_3 reduces exposure to ambiguity, without eliminating it altogether. For any portfolio $\mathbf{x} = (x_1, x_2, x_3)$ and any pair of securities s and $s' \neq s$, we define the *relative demand* to be the demand for the security that pays off in state s as a fraction of the sum of demands for securities that pay off in states s and s'

$$\frac{x_s}{x_s + x_{s'}}.$$

The proximity of this ratio to $1/2$ measures the extent to which the demands for securities s and s' are equalized.

Figure 1 below depicts a kernel density estimate of $x_1/(x_1 + x_3)$ and compares it with kernel density estimates of $x_1/(x_1 + x_2)$ and $x_3/(x_2 + x_3)$, which measure the extent to which subjects equalize payoffs in two states, exactly one of which is ambiguous. Before calculating these densities, we screen the data for safe ($x_1 = x_2 = x_3$) and boundary ($x_s = x_{s'} = 0$ for some $s \neq s'$) portfolios using a narrow confidence interval of two tokens.⁵ The safe and boundary portfolios account for 20.0 and 6.5 percent of all portfolios. Perhaps as expected, the three distributions are nearly symmetric and concentrated near the midpoint $1/2$. More interestingly, the mode is more pronounced in the distributions of relative demands for securities that pay off in ambiguous states, x_1 and x_3 . This provides clear evidence of ambiguity aversion.⁶

[Figure 1 here]

The percentage of portfolios for which $x_1/(x_1 + x_3)$ lies between 0.45 and 0.55 is 32.4, and this increases to 41.6 percent if we consider relative demands lying between the bounds 0.4 and 0.6. The corresponding percentages for $x_1/(x_1 + x_2)$ are 26.4 and 36.8 and for $x_3/(x_2 + x_3)$ they are 28.2 and 38.4, respectively.⁷ The tendency to equate x_1 and x_3 could, of course, result from simple risk aversion, but this is where the unambiguous and risky state are useful. The greater tendency to equate the demands for x_1 and x_3 suggests an aversion to ambiguity rather than just risk.

⁵This accounts for small mistakes resulting from the slight imprecision of subjects' handling of the mouse.

⁶Two-sample Kolmogorov-Smirnov tests for equality of distribution functions show that no two relative demands have the same distribution. This is as expected, as our subjects were given a large and rich menu of randomly determined budget sets.

⁷We also performed the analysis for each half of the data, and found a very high concordance in the two sets of distributions.

4.2 Individual behavior

The aggregate data above tell us little about the choice behavior of individual subjects. For that purpose we make within-subject comparisons of the number of unambiguous portfolios for which $x_1/(x_1 + x_3)$ lies between 0.45 and 0.55 and the *average* number of portfolios for which $x_1/(x_1 + x_2)$ or $x_2/(x_2 + x_3)$ lies between these bounds. Before calculating these numbers, we again screen the data for safe and boundary portfolios. This results in many fewer observations for a small number of subjects. Figure 2 presents the data as points in a scatterplot. The most notable feature of the distribution in Figure 2 is that the data are concentrated above the diagonal and skewed to the upper left. This provides evidence on both the prominence and the heterogeneity of subjects' attitudes toward ambiguity.

[Figure 2 here]

Next, Figure 3 depicts particular portfolios chosen by the individual subjects who serve to illustrate ideal types in terms of token shares (left panel) and budget shares (right panel) for the three securities as points in the unit simplex. The vertices of the unit simplex correspond to portfolios consisting of one of the three securities. Each point in the simplex represents a portfolio as a convex combination of the extreme points. For any portfolio $\mathbf{x} = (x_1, x_2, x_3)$, we define the *token share* of the security that pays off in state s to be the number of tokens payable in state s as a fraction of the sum of tokens payable in all three states

$$\frac{x_s}{x_1 + x_2 + x_3}.$$

We also define the *budget share* (or expenditure share) of the security that pays off in state s to be the expenditure on tokens invested in this security as a fraction of total expenditure. Since prices are normalized so that total expenditure equals unity, the budget share is simply $p_s x_s$. We note that we have chosen subjects whose behavior corresponds to one of several prototypical notions of risk or ambiguity aversion and illustrate the striking regularity within subjects and heterogeneity across subjects that is characteristic of all our data. We use the same subjects for illustrative purposes later as well.

[Figure 3 here]

Figure 3A depicts the choices of a subject (ID 11) who always chose nearly equal portfolios $x_1 = x_2 = x_3$, suggesting infinite risk aversion. Figure

3B shows a very different case, the choices of a subject (ID 20) who, with a few exceptions, invested all his tokens in the cheapest security, behavior which is consistent with pure risk neutrality. Figure 3C depicts the portfolio choices of a subject (ID 31) who equalizes expenditures $p_1x_1 = p_2x_2 = p_3x_3$, rather than tokens, across the three securities. This behavior is consistent with a logarithmic von Neumann-Morgenstern utility function (with respect to money).

A more interesting regularity is illustrated in Figure 3D, which shows the portfolio choices of a subject (ID 23) who, with very few exceptions, invested nearly equal amounts in the securities that pays off in the ambiguous states $x_1 = x_3 \neq x_2$. Figure 3E depicts the choices of a subject (ID 12) with a similar regularity, albeit implemented less precisely. Finally, Figure 3F shows the choices of a subject (ID 37) who did not demand nearly equal amounts of the securities that pay off in the ambiguous states $x_1 \neq x_3$, but these demands were much closer to each other than to the demand for the security that pays off in the unambiguous state x_2 . The behaviors of these subjects suggest ambiguity aversion, in the sense that they are trying to reduce the sensitivity of their payoffs to states with ambiguous probabilities.

The data for the full set of subjects are available in Appendix II, where we also show, for each subject, the relationships between the log-price ratio $\ln(p_1/p_3)$ and the relative demand $x_1/(x_1 + x_3)$ and between $\ln(p_1/p_2)$ and $x_1/(x_1 + x_2)$.⁸ These scatterplots illustrate the sensitivity of portfolio decisions to changes in relative prices. We emphasize again that for most subjects the data are much less regular and, for those subjects, it is more difficult to see these relationships in a scatterplot. Nevertheless, the portfolio choices for the full set of subjects reveal striking regularities within and marked heterogeneity across subjects.

4.3 Testing rationality

The most basic question to ask about choice data is whether it is consistent with individual utility maximization. In principle, the presence of ambiguity could cause not just a departure from expected utility, but a more fundamental departure from rationality. Thus, before calibrating particular utility functions, we first test whether choices can be utility-generated. Afriat (1967) and Varian (1982, 1983) show that choices from a finite number of budget sets are consistent with maximization of a well-behaved (piecewise linear, continuous, increasing, and concave) utility function if and only if

⁸Online Appendix II: http://emlab.berkeley.edu/~kariv/ACGK_I_A2.pdf.

they satisfy the Generalized Axiom of Revealed Preference (GARP). Since GARP offers an exact test (either the data satisfy GARP or they do not), we assess how nearly individual choice behavior complies with GARP by using Afriat's (1972) Critical Cost Efficiency Index (CCEI), which measures the fraction by which each budget constraint must be shifted in order to remove all violations of GARP. By definition, the CCEI is between 0 and 1: indices closer to 1 mean the data are closer to perfect consistency with GARP and hence to perfect consistency with utility maximization.

Over all subjects, the CCEI scores averaged 0.945. We interpret this number as a confirmation that subject choices are generally consistent with utility maximization. To calibrate the CCEI, we use a test designed by Bronars (1987). We simulate the choices of random agents whose choices are uniformly distributed on the feasible region of the budget hyperplane. Figure 4 shows the distribution of CCEI scores generated by 25,000 random agents and compares this with the observed distribution. The histograms in Figure 4 make plain that the significant majority of our subjects came much nearer to consistency with utility maximization than random agents would have done and that their CCEI scores were only slightly worse than the score of a perfect utility maximizer. Note that there is no significance threshold for the CCEI. If we follow Varian's (1991) suggestion and choose the 0.95 efficiency level as our critical value, we find that 93 subjects (60.4 percent) have CCEI scores above this threshold, while none of the random agents has a CCEI score that high.

[Figure 4 here]

We refer the interested reader to Choi et al. (2007a, 2007b) for more details on testing for consistency with GARP and other measures that have been proposed for this purpose by Varian (1991) and Houtman and Maks (1985). In practice, all these measures yield similar conclusions. Appendix III lists, by subject, the number of violations GARP, and also reports the values of the three goodness-of-fit indices.⁹ Choi et al. (2007a) demonstrate that if utility maximization is not in fact the correct model, then the experiment is sufficiently powerful to detect it. Finally, we note that subjects would be unlikely to produce behavior that is consistent with utility maximization if they had any difficulties understanding the decision problem or using the computer program.

⁹Online Appendix III: http://emlab.berkeley.edu/~kariv/ACGK_I_A3.pdf.

5 Models of Ambiguity

In this section we introduce the two parametric utility specifications to be estimated. The first is a “kinked” specification. It can be derived as a special case of a variety of utility models: MEU, CEU, Contraction Expected Utility, and α -MEU. The second is a “smooth” specification that can be derived from REU. Our discussion focuses on the restrictions – in the form of specific functional form assumptions – that we place on the general models in order to obtain parametric versions that can be estimated.

Each of our specifications is characterized by two parameters, one of which can be identified with risk aversion and one of which can be identified with ambiguity aversion. The advantage of a parsimonious specification is that we can summarize subjects’ preferences in terms of a few parameters; the disadvantage is that it may restrict behavior compared to the most general version. We report simple empirical distributions of the estimated parameters.

In addition to the two main specifications, we consider two important special cases. The first corresponds to SEU in the sense of Savage, while the second corresponds to an extreme form of MEU. Each is derived by setting the ambiguity parameter equal to some extreme value.

Our first parametric assumption relates to attitudes toward risk. We assume that risk preferences are represented by a von Neumann-Morgenstern utility function $u(x)$ with constant absolute risk aversion (CARA),

$$u(x) = -e^{-\rho x},$$

where x is the number of tokens and ρ is the coefficient of absolute risk aversion. This specification has two advantages. First, it is independent of the (unobservable) initial wealth level of the subjects. Second, it accommodates portfolios where $x_s = 0$ for some state s even when initial income is zero.

5.1 A kinked specification

The kinked utility function is so-called because the indifference curves have a “kink” at all portfolios where $x_1 = x_3$. The parametric specification we use has the form

$$U(\mathbf{x}; \alpha, \rho) = \alpha \left[-\frac{2}{3} \exp\{-\rho \min\{x_1, x_3\}\} - \frac{1}{3} \exp\{-\rho x_2\} \right] + (1 - \alpha) \left[-\frac{2}{3} \exp\{-\rho \max\{x_1, x_3\}\} - \frac{1}{3} \exp\{-\rho x_2\} \right], \quad (1)$$

where α is the ambiguity parameter and ρ is the coefficient of risk aversion. The distinguishing feature of this specification is its dependence on the minimum and maximum payoffs, $\min\{x_1, x_3\}$ and $\max\{x_1, x_3\}$, between the two

ambiguous states, 1 and 3. The agent knows that the probabilities of states 1 and 3 lie between 0 and $\frac{2}{3}$. In the best case scenario, the probability of the state in which he receives $\max\{x_1, x_3\}$ is $\frac{2}{3}$; in the worst case scenario, it is zero. What Equation (1) says is that the agents's utility is a weighted average, with weights α and $1 - \alpha$, of the expected utility in the worst-case and best-case scenarios.

We now demonstrate how this kinked functional form can be generated by different classes of preferences.

5.1.1 Maxmin Expected Utility with flexible priors

The Maxmin Expected Utility (MEU) model of Gilboa and Schmeidler (1989) evaluates a portfolio by its minimal expected utility over a set of subjective prior beliefs. This minimization over a non-singleton set can be interpreted as aversion to ambiguity. The general form of the MEU model is

$$U(\mathbf{x}) = \min_{\pi \in \Pi} \int_S u(x_s) d\pi(s),$$

where $\Pi \subseteq \Delta S$ is a closed convex set of prior beliefs over states.

Connecting the general MEU model to our kinked specification assumes that the utility over tokens takes the CARA form and that the set of priors is symmetric about $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. In particular, the set of priors is

$$\Pi_\delta = \left\{ \pi : \pi_2 = \frac{1}{3}, \frac{1}{3} - \delta \leq \pi_1 \leq \frac{1}{3} + \delta, \pi_3 = \frac{2}{3} - \pi_1 \right\}$$

for some $0 \leq \delta \leq \frac{1}{3}$. Larger values of δ indicate a larger set of priors, hence more ambiguity. This reduces the general MEU model to the following two-parameter formula:

$$U(\mathbf{x}; \delta, \rho) = - \left(\frac{1}{3} + \delta \right) \exp\{-\rho \min\{x_1, x_3\}\} - \frac{1}{3} \exp\{-\rho x_2\} \\ - \left(\frac{1}{3} - \delta \right) \exp\{-\rho \max\{x_1, x_3\}\}.$$

This equation is exactly Equation (1) with a change of variables, letting $\alpha = \frac{1}{2} + \frac{3}{2}\delta$.

5.1.2 Choquet Expected Utility with flexible capacity

The Choquet Expected Utility (CEU) model of Schmeidler (1989) is related to MEU and takes the following general form:

$$U(\mathbf{x}) = \int_S u(x_s) d\nu(s),^{10}$$

¹⁰The exact formula for integration with respect to a capacity can be found in Schmeidler (1989).

where ν is a nonadditive capacity over the state space. Ambiguity in the CEU model is captured by the convexity of the capacity ν .¹¹

Any CEU representation with a convex capacity can be rewritten as an MEU representation where the set of priors is the core of the capacity. Correspondingly, if we assume CARA utility over tokens and that the capacity is symmetric over the two ambiguous states, then the CEU model reduces to the parameterized MEU model with symmetric priors presented in the previous section. In particular, if the capacity obeys:

$$\begin{aligned} \nu(\{1\}) = \nu(\{3\}) &= \frac{1}{3} - \delta, & \nu(\{2\}) &= \frac{1}{3}, \\ \nu(\{1, 2\}) = \nu(\{2, 3\}) &= \frac{2}{3} - \delta, & \nu(\{1, 3\}) &= \frac{2}{3}, \end{aligned}$$

for some $0 \leq \delta \leq \frac{1}{3}$, then the implied Choquet integral reduces to Equation (1), via the same change of variables $\alpha = \frac{1}{2} + \frac{3}{2}\delta$.

5.1.3 Contraction Expected Utility with fixed information

The contraction model of Gajdos et al. (2008) incorporates objective information about the set of possible prior distributions over states. It enriches the standard subjective setup by considering acts or portfolios paired with some set of objectively known possible priors. The agent partially contracts this set towards its center and then applies the MEU criterion to this smaller set of priors. The general representation is

$$U(\mathbf{x}) = \min \left\{ \int_S u(x_s) d\pi(s) : \pi \in (1 - \epsilon)\{s(\Pi)\} + \epsilon\Pi \right\},$$

where $s(\Pi) \in \Delta S$ is the Steiner point (a geometric notion of the center) of the set Π of objectively specified priors.¹² Larger values of $\epsilon \in [0, 1]$ place more weight on the entire set of possible priors Π and, hence, suggest more ambiguity.

The experimental choice problem can be represented in this form, where every portfolio is paired with the same set of objective priors, namely $\Pi = \{\pi : \pi_2 = \frac{1}{3} \text{ and } \pi_1 + \pi_3 = \frac{2}{3}\}$. Its Steiner point is $s(\Pi) = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. As Hayashi et al. (forthcoming) mention, the contraction model with a fixed set of possible priors is identical to a special form of the MEU model. To be specific, maintaining the CARA form for utility over tokens, the contraction

¹¹A capacity is convex if $\nu(A \cup B) + \nu(A \cap B) \geq \nu(A) + \nu(B)$ for any sets A and B .

¹²The convex combination of two sets A and B is defined as the union of their pointwise convex combinations: $\lambda A + (1 - \lambda)B = \{\lambda a + (1 - \lambda)b : a \in A, b \in B\}$.

model reduces to:

$$U(\mathbf{x}; \epsilon, \rho) = -\left(\frac{1+\epsilon}{3}\right) \exp\{-\rho \min\{x_1, x_3\}\} - \frac{1}{3} \exp\{-\rho x_2\} \\ - \left(\frac{1-\epsilon}{3}\right) \exp\{-\rho \max\{x_1, x_3\}\}.$$

This is exactly the MEU model above with $\delta = \frac{\epsilon}{3}$ and is the kinked specification in Equation (1) with $\alpha = \frac{1-\epsilon}{2}$.

5.1.4 α -Maxmin Expected Utility with fixed priors

A proposed generalization of MEU is α -Maxmin Expected Utility (α -MEU) characterized by Ghirardato et al. (2004) and Olszewski (2006), which evaluates each portfolio by a convex combination of its minimal and maximal expected utilities over some set of subjective prior beliefs over states.¹³ The general form of the α -MEU model is

$$U(\mathbf{x}) = \alpha \cdot \min_{\pi \in \Pi} \int u(x_s) d\pi(s) + (1 - \alpha) \cdot \max_{\pi \in \Pi} \int u(x_s) d\pi(s),$$

where $\Pi \subseteq \Delta S$ is a closed convex set of distributions over states and $\alpha \in [0, 1]$ reflects the relative weight of the worst versus the best possible expected utility of \mathbf{x} given Π . Hence, α serves as a parameter reflecting ambiguity aversion. (In the most general case, the α -MEU parameter could depend on the portfolio under consideration $\alpha(\mathbf{x})$.)

If we assume that u has the CARA form and that the set of priors Π is the entire set of distributions consistent with the objective information in the experiment, $\Pi = \{\pi : \pi_2 = \frac{1}{3}\}$, this reduces to the two-parameter formula in Equation (1). The weight α and the set of priors Π in the α -MEU model cannot be separately identified. In fact, Siniscalchi (2006) proves that the α -MEU and MEU models are generally confounded in the symmetric case: any MEU representation with some fixed symmetric set of priors can be rewritten as one of a continuum of α -MEU representations with arbitrarily small alternative sets of priors.

In all of the models described above, the parameter α that appears in Equation (1) depends on the set Π (or the capacity ν in the case of CEU). Unless the set Π is objectively known, knowledge of the estimated parameter α does not allow us to characterize the degree of ambiguity aversion independently from the degree of ambiguity in the decision problem. In any case, the lack of identification is endemic to these theoretical models, rather than a feature of our data. When we adopt the MEU interpretation, we are fixing $\alpha = 1$

¹³In the most general version, $\alpha(\mathbf{x})$ could depend on the portfolio \mathbf{x} under consideration.

and allowing the set of priors to vary; when we adopt the α -MEU interpretation, we are fixing the set of priors and allowing α to vary. To simplify the exposition and facilitate comparisons, we adopt the second convention as our main interpretation in the sequel.

5.2 A smooth specification

Our second utility specification is differentiable everywhere. The utility of a portfolio $\mathbf{x} = (x_1, x_2, x_3)$ takes the form

$$U(\mathbf{x}; \alpha, \rho) = \int_0^{\frac{2}{3}} -\exp \left\{ -\alpha \begin{pmatrix} -\pi_1 \exp\{-\rho x_1\} - \frac{1}{3} \exp\{-\rho x_2\} \\ -(\frac{2}{3} - \pi_1) \exp\{-\rho x_3\} \end{pmatrix} \right\} d\pi_1, \quad (2)$$

This specification involves two iterated integrals. First, the formula inside the parentheses is the expected value of the CARA utility of the portfolio \mathbf{x} when the probability of the first state is known to be π_1 . Next, the integral ranging from 0 to $\frac{2}{3}$ takes the expectation of these expected utilities with respect to the uniform distribution for π_1 , with each expected utility transformed using a CARA aggregator.

While the kinked specification can be interpreted using a variety of different models, the smooth specification is really motivated by a single model. A recent view of ambiguity aversion (Ergin and Gul, 2004; Klibanoff et al., 2005; Nau, 2005; and Seo, 2007; as well as related work by Halevy and Feltkamp, 2005; Giraud, 2006; and Ahn, 2008) assumes the agent has a subjective (second-order) distribution μ over the possible (first-order) prior beliefs π over states. Unsure which of the possible first-order prior beliefs actually governs the states, the agent transforms the expected utilities for all prior beliefs π by a concave function φ before integrating these utilities with respect to his second-order distribution μ . This procedure is entirely analogous to the transformation of wealth into cardinal utility before computing expected utility under risk. The concavity of this transformation captures ambiguity aversion. We follow Halevy (2007) in referring to this model as Recursive Expected Utility (REU), owing to its recursive double expectation.

The general form of the REU model is

$$U(\mathbf{x}) = \int_{\Delta S} \varphi \left(\int_S u(x_s) d\pi(s) \right) d\mu(\pi),$$

where $\mu \in \Delta(\Delta(S))$ is a (second-order) distribution over possible priors π on

S and $\varphi : u(\mathbf{R}_+) \rightarrow \mathbf{R}$ is a possibly nonlinear transformation over expected utility levels.¹⁴

To facilitate comparison with the kinked specification, we reduce the REU model to two parameters. Assuming that

$$\varphi(z) = -e^{-\alpha z},$$

which replicates the constant curvature of u , and that μ is uniformly distributed over the set of priors consistent with the objective information $\Pi = \{\pi : \pi_2 = \frac{1}{3}\}$, this specializes to the two-parameter formula in Equation (2). Here, α reflects the curvature of the aggregator φ and hence measures the degree of ambiguity aversion.

One of the crucial features of the REU specification is its reliance on a cardinal utility indicator. Unlike the preferences generated by SEU, MEU and α -MEU, which are invariant to affine transformations of the utility function $u(\cdot)$, the preferences generated by REU are not independent of a change in the scale of utility. For example, if we introduce a scale parameter and set $u(x) = -Ae^{-\rho x}$, the concavity of the transformation φ implies that the ranking of uncertain prospects will not be invariant to changes in A . Since the parameters α and A enter Equation (2) only in the form of the product αA , we can estimate αA but cannot identify the values of α and A separately. If we assume a common scale factor for all subjects, say $A = 1$, interpersonal comparisons of ambiguity aversion will still be affected by risk aversion. A higher coefficient of absolute risk aversion, ρ , will reduce the range of the function $u(x) = -e^{-\rho x}$ and, hence, will reduce the ambiguity to which the agent is exposed. We can normalize the ambiguity parameters to take into account the different ranges of expected utility for different subjects, but the meaning of such comparisons is not clear. We will return to this subject in the next section when we discuss the parameter estimates of the REU model.

5.3 Restricted specifications

5.3.1 Ambiguity neutrality: Subjective Expected Utility with a fixed prior

Subjective expected utility (SEU) is a special case of both the kinked and smooth formulations:

$$U(\mathbf{x}; \rho) = -\frac{1}{3} \exp\{-\rho x_1\} - \frac{1}{3} \exp\{-\rho x_2\} - \frac{1}{3} \exp\{-\rho x_3\}.$$

¹⁴Here, $\Delta(\Delta(S))$ denotes the space of all probability measures over $\Delta(S)$, the set of all probability distributions on S .

This corresponds to the kinked specification in Equation (1) with $\alpha = \frac{1}{2}$ and to the smooth specification in Equation (2) with $\alpha = 0$ and provides a benchmark for probabilistic sophistication within these specifications.

To derive this formula directly, recall that the general SEU model of Savage (1954) consists of a utility function u which is integrated with respect to a single subjective probability distribution π . The general form for the utility of a portfolio $\mathbf{x} = (x_1, x_2, x_3)$ is:

$$U(\mathbf{x}) = \int_S u(x_s) d\pi(s)$$

where π is a subjective probability over states of the world and u is a cardinal utility index over tokens. If we assume that the agent believes the ambiguous states in our experimental choice problem are equally probable, that is, her prior belief over states is $\pi = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ and has CARA utility over tokens, this specializes to the above formula.

5.3.2 Extreme ambiguity aversion: Maxmin Expected Utility with maximal priors

The opposite special case for both the kinked and smooth specifications is the following restricted formulation:

$$U(\mathbf{x}; \rho) = -\frac{2}{3} \exp\{-\rho \min\{x_1, x_3\}\} - \frac{1}{3} \exp\{-\rho x_2\}.$$

This corresponds to the kinked specification in Equation (1) with $\alpha = 1$ and to the smooth specification in Equation (2) as $\alpha \rightarrow \infty$ and provides the opposite benchmark of the most ambiguity averse subspecification within these models.

5.4 Properties of demand

Before proceeding to the estimation of the parametric models, it is important to understand the implications of the different models of ambiguity aversion for individual behavior. For this reason, we first illustrate the relative demand functions for the different models and different parameter values. This exercise will allow us to see how changes in risk and ambiguity aversion affect the elasticity of demands and how the kinked and smooth specifications affect the shape of the demand curves. These differences will be important in understanding how the models fit the data.

The simulated demand functions are illustrated in Figure 5 below. The figure shows the relationships between the log-price ratio $\ln(p_1/p_3)$ and the

optimal relative demand $x_1^*/(x_1^* + x_3^*)$ (left panels) and between $\ln(p_1/p_2)$ and $x_1^*/(x_1^* + x_2^*)$ (right panels), for each model, using a range of parameter values (each panel assumes a different value for α).¹⁵ Comparing the two relative demands illustrate the differences in the tradeoffs agents make in their ambiguity preferences and risk preferences under the two specifications.

[Figure 5 here]

Figure 5A illustrates the relative demands for the kinked specification in Equation (1). If the prices of the securities that pay off in the ambiguous states, p_1 and p_3 , are similar ($\ln(p_1/p_3)$ is close to zero), then the optimum portfolio choice satisfies $x_1^* = x_3^*$ and is insensitive to ambiguity. The only effect of increasing the level of ambiguity aversion α is to make this intermediate range of price ratios larger. The key feature of the kinked model is this flat range in which the portfolio satisfies $x_1^* = x_3^*$ and is insensitive to ambiguity. By contrast, the choice of a portfolio without ambiguity ($x_1 = x_2$) is a knife-edge case in the smooth model.

Figure 5B depicts the relative demands generated by the smooth specification in Equation (2), for different values of α and ρ . The relationships between the log-price ratio $\ln(p_1/p_3)$ and the optimal relative demand $x_1^*/(x_1^* + x_3^*)$ and between $\ln(p_1/p_2)$ and $x_1^*/(x_1^* + x_2^*)$, which illustrate the tradeoffs that the agent makes between the payoffs in ambiguous states and between the the payoffs in an ambiguous state and an unambiguous state, are smooth for all price ratios. The ambiguity aversion parameter α flattens the $x_1^*/(x_1^* + x_3^*)$ curves in a manner qualitatively similar to having increased risk aversion.

6 Parametric Analysis

6.1 Econometric specification

The data generated by an individual's choices are denoted by $\{(\mathbf{x}^i, \mathbf{p}^i)\}_{i=1}^{50}$, where $\mathbf{x}^i = (x_1^i, x_2^i, x_3^i)$ is the actual portfolio chosen by the subject and $\mathbf{p}^i = (p_1^i, p_2^i, p_3^i)$ denotes the vector of security prices. For each subject n and for each specification, we generate estimates of the ambiguity and risk aversion parameters, $\hat{\alpha}_n$ and $\hat{\rho}_n$, using nonlinear least squares (NLLS). These

¹⁵The figures are difficult to see in the small black and white format required in the printed version. We refer the reader to color figures which are available in the electronic version of the paper at http://econ.berkeley.edu/~kariv/ACGK_1.pdf.

estimates are chosen to minimize

$$\sum_{i=1}^{50} \|\mathbf{x}^i - \mathbf{x}^*(\mathbf{p}^i; \alpha_n, \rho_n)\|,$$

where $\|\cdot\|$ denotes the three-dimensional Euclidean norm.¹⁶

Before proceeding to estimate the parameters, we make two observations about the econometric specification:

- First, we omit the six subjects with CCEI scores below 0.80 (ID 1, 59, 61, 81, 82 and 93) as their choices are not sufficiently consistent to be considered utility-generated. Additionally, Afriat’s (1967) theorem tells us that when a rationalizing utility function exists, it may be chosen to be well-behaved (piecewise linear, continuous, increasing, and concave). In particular, our analysis over linear budget sets does not allow us to distinguish between risk- or ambiguity-loving behavior, on the one hand, and risk- or ambiguity-neutral behavior, on the other. We therefore restrict the parameters so that preferences are always risk and ambiguity averse. This requires $\rho \geq 0$ in both specifications, $1/2 \leq \alpha \leq 1$ in the kinked specification and $0 \leq \alpha$ in the smooth specification.
- Secondly, when the parameter measuring risk aversion ρ is large, ambiguity aversion cannot be separately identified, since all of the hedging between the securities that pay off in the ambiguous states, x_1 and x_3 , can be attributed to the extreme risk aversion. To avoid this identification problem, we screen out the four subjects (ID 11, 24, 112 and 119) who almost always chose the safe portfolio $x_1 = x_2 = x_3$. The preferences of these subjects are easily identifiable from the scatterplots of their choices. Finally, because of computational difficulties when α is large, we also impose the restriction $\alpha \leq 2$ in the smooth specification. This involves minimal loss in fit, since the predicted choices with such a high levels of ambiguity aversion are virtually identical.

This leaves a set of 144 subjects (93.5 percent) with consistent non-extreme risk preferences for whom we recover the underlying preferences by estimating both specifications. We emphasize again that our estimations will be

¹⁶For simplicity, the estimation technique for both specifications is NLS, rather than a structural model using maximum likelihood (ML). We favor the NLS approach, because it provides a good fit and offers flexibility, tractability and straightforward interpretation. The NLS estimation is still computationally intensive for even moderately large data sets. We use bootstrapping to approximate standard errors.

done for each subject n separately, generating separate estimates $\hat{\alpha}_n$ and $\hat{\rho}_n$.

6.2 Econometric results

To economize on space, the individual-level estimates are relegated to Appendix IV.¹⁷ Table 1 provides a population-level summary of the individual-level estimation results by reporting summary statistics and percentile values. As noted above, the smooth specification in Equation (2) is not invariant to affine utility transformations. To this end, in Table 1, we present the statistics for the estimated raw ambiguity parameter $\hat{\alpha}$, as well as the statistics for different normalized ambiguity parameters $\hat{\alpha}^t$ defined implicitly as a function of $\hat{\alpha}$ as follows:

$$\hat{\alpha}^t = \frac{t\hat{\alpha}}{1 - e^{t\hat{\rho}}}.$$

This normalization readjusts the level of cardinal utility for t tokens to be constant across subjects with varying degrees of risk aversion. The formula can be obviously altered to normalize the comparison for different levels where the parameter $\hat{\alpha}^t$ reflects the curvature of the second-order expected utility index in the smooth specification, thus measuring absolute ambiguity aversion.

[Table 1 here]

Using the kinked specification, of the 144 subjects listed in Appendix IV, 56 subjects (38.9 percent) have non-kinky preferences $\hat{\alpha}_n \approx 1/2$ so their choices are well approximated by SEU.¹⁸ We cannot reject the hypothesis that $\hat{\alpha}_n = 1/2$ for a total of 67 subjects (46.5 percent) at the 95 percent significance level. The remainder appear to have significant degrees of ambiguity aversion $\hat{\alpha}_n > 1/2$. Of those, a single subject (ID 15) displayed infinite ambiguity aversion $\hat{\alpha}_n \approx 1$. His behavior is consistent with MEU. We reject the hypothesis that $\hat{\alpha}_n = 1$ for all other subjects.

Similarly, using the smooth specification, 44 subjects (30.6 percent) have $\hat{\alpha}_n \approx 0$, indicating ambiguity neutrality but we cannot reject the hypothesis that $\hat{\alpha}_n = 0$ for a total of only 38 subjects (26.4 percent) at the 95 percent significance level. The behavior of these subjects is consistent with SEU.

¹⁷Online Appendix IV: http://emlab.berkeley.edu/~kariv/ACGK_I_A4.pdf.

¹⁸In comparison, Halevy (2007) reports that only 28 of his 142 subjects (19.7 percent) behave as if they were ambiguity neutral.

Additionally, seven subjects (4.9 percent) have boundary ambiguity aversion parameter value $\hat{\alpha}_n \approx 2$. We can reject the hypothesis that $\hat{\alpha}_n = 2$ for all other subjects. Finally, in both specifications, a significant fraction of our subjects have moderate levels of risk aversion $\hat{\rho}_n$, which are within the range of estimates reported in Choi et al. (2007b).

Figure 6 presents the data from Appendix IV graphically in the form of scatterplots of the estimates, and illustrates the heterogeneity of preferences that we find in both specifications. Figure 6A shows a scatterplot of $\hat{\alpha}_n$ and $\hat{\rho}_n$, in the kinked specification. Figure 6B shows the scatterplot for the smooth specification. To facilitate presentation of the data, Figure 6C shows the same scatterplot after omitting 16 subjects, whose $\hat{\alpha}_n$ value is higher than 1/2. Certainly, as alluded to earlier in the paper, the estimated ambiguity coefficients $\hat{\alpha}_n$ that come out of the smooth specification are directly comparable only across subjects with similar risk attitudes. This comparison is more delicate when the estimated risk coefficients $\hat{\rho}_n$ are different. Finally, note that in both the kinked and smooth specifications there is considerable heterogeneity in both parameters, $\hat{\alpha}_n$ and $\hat{\rho}_n$, and that their values are positively correlated ($r^2 = 0.157$ and $r^2 = 0.229$, respectively).

[Figure 6 here]

Finally, Figure 7 below shows the relationship between log-price ratio $\ln(p_1/p_3)$ and the *actual* relative demand $x_1/(x_1 + x_3)$ (blue) and *estimated* relative demand $\hat{x}_1/(\hat{x}_1 + \hat{x}_3)$ (red) in the kinked (left panels) and smooth (right panels) specifications for the same group of subjects that we followed in the non-parametric analysis.¹⁹ Note that $\hat{x}_1/(\hat{x}_1 + \hat{x}_3)$ is calculated using the individual-level estimates, $\hat{\alpha}_n$ and $\hat{\rho}_n$. We reemphasize that we carefully selected these subjects in order to illustrate salient features of the data. The figures for the full set of subjects and for all models are available in Appendix V, which also depicts the relationship between the log-price ratio $\ln(p_1/p_2)$ and the actual relative demand $x_1/(x_1 + x_2)$ (blue) and estimated relative demand $\hat{x}_1/(\hat{x}_1 + \hat{x}_2)$, as well as the actual and estimated portfolios in terms of token and expenditure shares represented as points in a simplex.²⁰

[Figure 7 here]

The first subject (ID 11) very precisely implemented infinite risk aversion preferences. The ambiguity aversion parameter of this subject is thus

¹⁹ These figures are difficult to see in a black and white format. Color figures are available in the electronic version of the paper at http://econ.berkeley.edu/~kariv/ACGK_I.pdf.

²⁰ Online Appendix V: http://emlab.berkeley.edu/~kariv/ACGK_I_A5.pdf.

unidentified. Figure 7A is therefore omitted. Figure 7B shows the relationship between $\ln(p_1/p_3)$ and the estimated relative demand $x_1/(x_1 + x_3)$ for the subject (ID 20) who most closely approximated risk neutral preferences with $(\hat{\alpha}, \hat{\rho})_{\text{kinked}} = (0.521, 0.000)$ and $(\hat{\alpha}, \hat{\rho})_{\text{smooth}} = (0.009, 0.001)$. Both the kinked and smooth specifications suggest a nontrivial degree of ambiguity aversion, which is driven by a few exceptional choices where this subject chose nearly unambiguous portfolios $x_1 = x_3$. Notice that both models do a very good job of predicting his boundary portfolios, but perform less well in predicting the “outliers.” This subject appears to be very close to risk neutrality, but perhaps reveals some degree of ambiguity aversion through his outlying portfolios.

Figure 7C shows the subject (ID 31) who precisely implemented logarithmic preferences with $(\hat{\alpha}, \hat{\rho})_{\text{kinked}} = (0.500, 0.076)$ and $(\hat{\alpha}, \hat{\rho})_{\text{smooth}} = (0.000, 0.076)$. Nonetheless, the exponential form performs quite well in terms of fit. More interestingly, Figure 7D shows the relationship for a subject (ID 23) with $(\hat{\alpha}, \hat{\rho})_{\text{kinked}} = (0.728, 0.419)$ and $(\hat{\alpha}, \hat{\rho})_{\text{smooth}} = (0.004, 0.487)$, who quite precisely chose unambiguous portfolios $x_1 = x_3$. The estimated ambiguity parameter for the kinked specification is among the highest in the sample. This subject thus closely approximates MEU preferences.

Figure 7E shows the fitted relationships for a subject (ID 12) with $(\hat{\alpha}, \hat{\rho})_{\text{kinked}} = (0.569, 0.054)$ and $(\hat{\alpha}, \hat{\rho})_{\text{smooth}} = (0.000, 0.059)$, who with few exceptions chose unambiguous portfolios ($x_1 = x_3$). This subject’s departures from unambiguous portfolios are precipitated by extreme log-price ratios $\ln(p_1/p_3)$. This is entirely consistent with α -MEU. Furthermore, the kinked specification provides an improved fit over the restricted MEU specification, especially in picking up the few extremely ambiguous portfolios. On the other hand, the smooth specification fails to pick up this subject’s aversion to ambiguity. Clearly, it cannot accommodate both the large interval of intermediate log-price ratios $\ln(p_1/p_3)$ in which unambiguous portfolios $x_1 = x_3$ are chosen and the extremely ambiguous portfolios $x_1 = 0$ or $x_3 = 0$ chosen for low and high price ratios, respectively. The estimated relationship for the smooth specification makes clear that the model gives up the former to improve its fit for the latter. A review of the full set of subjects shows, in many cases, a substantial number of unambiguous portfolios ($x_1 = x_3$), which can only be accommodated by the kinked specification.

Finally, Figure 7F shows the fitted relationship for a subject (ID 37) with $(\hat{\alpha}, \hat{\rho})_{\text{kinked}} = (0.880, 0.075)$ and $(\hat{\alpha}, \hat{\rho})_{\text{smooth}} = (2.000, 0.074)$. This subject chose portfolios with smaller differences between x_1 and x_3 than between any other pair of securities, but did not usually move to the extreme of choosing unambiguous portfolios ($x_1 = x_3$). Although this subject is

averse to ambiguity, the fit of the kinked specification barely improves on the restricted MEU specification. On the other hand, the smooth specification improves the fit considerably over either restricted specifications MEU or SEU. This is especially true for extreme log-price ratios $\ln(p_1/p_3)$, where MEU predicts values for x_1 far too close to x_3 , while SEU predicts values far too distant.

7 Conclusion

The presence of ambiguity aversion in human subjects has been repeatedly demonstrated in the laboratory. Recent theoretical developments have given greater precision to the concept of ambiguity aversion, but there have been few systematic empirical studies of these models. In this paper, we have used a rich data set, containing a relatively large number of choices for each subject, to estimate two classes of parametric utility functions.

In addition to identifying the existence of ambiguity aversion and the type of model that is capable of explaining it, we simultaneously estimate measures of risk and ambiguity aversion. This is the first step in answering an important series of questions: How important is ambiguity aversion? How well do different specifications fit the data? How important are the differences between the general functional forms, i.e., the kinked or smooth specification, and special cases such as SEU and MEU?

Our study confirms the heterogeneity of individual attitudes toward ambiguity that have been found in previous studies. This heterogeneity takes two forms. First, some subjects are better described by one or other of the models we estimate. Second, within a given model, the estimated measures of risk and ambiguity aversion are heterogeneous. We are able to provide a more precise description of this heterogeneity than previous studies because we have estimated the parameters of the different specifications at the individual level.

Our individual-level analyses show that preferences vary widely across subjects and range from risk neutrality with ambiguity aversion, to ambiguity neutrality with risk aversion, to infinite risk aversion. Although in both the kinked and smooth specifications there is considerable heterogeneity in both parameters, $\hat{\alpha}_n$ and $\hat{\rho}_n$, approximately two-third of our subjects have a positive degree of ambiguity aversion. The remaining one-third are well approximated by preferences consistent with SEU, a much higher proportion than found in previous experimental studies. There is also considerable heterogeneity in subjects' risk preferences in both the ambiguity-averse and

ambiguity-neutral subsamples but they are within the range of recent estimates of risk aversion.

One of the most striking features of the data is the strong tendency to equate the demand for the securities that pay off in the ambiguous states. This feature of the data is accommodated by the kinked specification, but is hard to reconcile with the smooth specification. An additional problem with the smooth specification is the difficulty of interpreting the parameters, which are not truly identified without some auxiliary assumption about cardinal utility. While it is not our intention to run a horse race between the different specifications, these observations may provide food for thought to decision theorists interested in developing more empirically relevant models.

References

- [1] Abdellaouia, M., A. Baillonb, L. Placidoa, and P. Wakker (2008) “The Rich Domain of Uncertainty,” mimo.
- [2] Afriat, S. (1967) “The Construction of a Utility Function from Expenditure Data,” *Econometrica*, 6, pp. 67-77.
- [3] Afriat, S. (1972) “Efficiency Estimates of Production Functions,” *International Economic Review*, 8, pp. 568-598.
- [4] Ahn, D. (2008) “Ambiguity Without a State Space,” *Review of Economic Studies*, 75, pp. 3-28.
- [5] Bossaerts, P., C. Plott, and W. Zame (2007) “Prices and Portfolio Choices in Financial Markets: Theory, Econometrics, Experiments,” *Econometrica*, 75, pp. 993-1038.
- [6] Bossaerts, P., P. Ghirardato, S. Guarnaschelli, and W. Zame (2008) “Prices and Allocations in Asset Markets with Heterogeneous Attitudes Towards Ambiguity,” mimo.
- [7] Bronars, S. (1987) “The power of nonparametric tests of preference maximization,” *Econometrica*, 55, pp. 693-698.
- [8] Camerer, C. (1995) “Individual Decision Making,” in *Handbook of Experimental Economics*. J. Kagel and A. Roth, eds. Princeton U. Press.
- [9] Camerer, C., and M. Weber (1992) “Recent Developments in Modelling Preferences: Uncertainty and Ambiguity,” *Journal of Risk and Uncertainty*, 5, pp. 325-370.

- [10] Choi, S., R. Fisman, D. Gale, and S. Kariv (2007a) “Revealing Preferences Graphically: An Old Method Gets a New Tool Kit,” *American Economic Review*, Papers & Proceedings, 97, pp. 153-158.
- [11] Choi, S., R. Fisman, D. Gale, and S. Kariv (2007b) “Consistency and Heterogeneity of Individual Behavior under Uncertainty,” *American Economic Review*, 97, pp. 1921-1938.
- [12] Ellsberg, D. (1961) “Risk, Ambiguity, and the Savage Axioms,” *Quarterly Journal of Economics*, 75, pp. 643-669.
- [13] Ergin, H., and F. Gul (2004) “A Subjective Theory of Compound Lotteries,” mimo.
- [14] Ghirardato, P., F. Maccheroni, and M. Marinacci (2004) “Differentiating Ambiguity and Ambiguity Attitude,” *Journal of Economic Theory*, 118, pp. 133-173.
- [15] Gajdos, T., T. Hayashi, J.-M. Tallon, and J.-C. Vergnaud (2008) “Attitude toward imprecise information,” *Journal of Economic Theory*, 140, pp. 27-65.
- [16] Gilboa, I., and D. Schmeidler (1989) “Maxmin Expected Utility with a Non-Unique Prior”, *Journal of Mathematical Economics*, 18, pp. 141-153.
- [17] Giraud, R. (2006) “Objective Imprecise Probabilistic Information, Second Order Beliefs and Ambiguity Aversion: An Axiomatization,” mimo.
- [18] Halevy, Y. (2007) “Ellsberg Revisited: An Experimental Study”, *Econometrica*, 75, pp. 503-536.
- [19] Halevy, Y., and V. Feltkamp (2005) “A Bayesian Approach to Uncertainty Aversion,” *Review of Economic Studies*, 72, pp. 449-466.
- [20] Hayashi T., and R. Wada (forthcoming) “Choice with Imprecise Information: An Experimental Approach,” *Theory and Decision*.
- [21] Hey, J., G. Lotito, A. Maffioletti (2007) “Choquet OK?” mimo.
- [22] Hey, J., G. Lotito, A. Maffioletti (2008) “The Descriptive and Predictive Adequacy of Theories of Decision Making Under Uncertainty/Ambiguity,” mimo.

- [23] Houtman, M., and J. Maks (1985) “Determining all Maximal Data Subsets Consistent with Revealed Preference.” *Kwantitatieve Methoden*, 19, pp. 89-104.
- [24] Klibanoff, P., M. Marinacci, and S. Mukerji (2005) “A Smooth Model of Decision Making under Ambiguity,” *Econometrica*, 73, pp. 1849-1892.
- [25] Knight, F. (1921) “Risk, Uncertainty, and Profit.”
- [26] Nau, R. (2006) “Uncertainty Aversion with Second-Order Utilities and Probabilities,” *Management Science*, 52, pp. 136-145.
- [27] Olszewski, W. (2007) “Preferences over Sets of Lotteries,” *Review of Economic Studies*, 74, pp. 567-595.
- [28] Savage, L. J. (1954) “The Foundations of Statistics.”
- [29] Segal, U. (1987) “The Ellsberg Paradox and Risk Aversion: An Anticipated Utility Approach.” *International Economic Review*, 28, pp. 175–202.
- [30] Segal, U. (1990) “Two-Stage Lotteries Without the Reduction Axiom.” *Econometrica*, 58, pp. 349-377.
- [31] Seo, K. (2008) “Ambiguity and Second-Order Belief,” mimo.
- [32] Varian, H. (1982) “The Nonparametric Approach to Demand Analysis.” *Econometrica*, 50, pp. 945-972.
- [33] Varian, H. (1983) “Non-Parametric Tests of Consumer Behavior.” *Review of Economic Studies*, 50, pp. 99-110.
- [34] Varian, H. (1991) “Goodness-of-Fit for Revealed Preference Tests,” mimo.

Figure 1: The distribution of relative demands






Figure 2: The number of diagonal portfolios by subject




Vertical axis: the number of unambiguous portfolios for which $0.45 \leq x_1 / (x_1 + x_3) \leq 0.55$. Horizontal axis: the average number of portfolios for which $x_1 / (x_1 + x_2)$ or $x_2 / (x_2 + x_3)$ lies between these bounds. These numbers are calculated after screening out safe and boundary portfolios using a narrow confidence interval of two tokens. We indicate the IDs of the subjects with many unambiguous portfolios.

Figure 3: Individual-level data

A: ID 11



B: ID 20



C: ID 31






Figure 3 (cont.)

D: ID 23



E: ID 12



F: ID 37




Figure 4: The distribution of CCEI scores




Figure 5: An illustration of the relationships between log-price ratio $\ln(p_1 / p_3)$ and optimal token share $x_1^* / (x_1^* + x_3^*)$

A: Kinked specification (Equation 1)




Figure 5 (cont.)







Figure 5 (cont.)

B: Smooth specification (Equation 2)




Figure 5 (cont.)







Figure 6: Scatterplot of the estimated parameters $\hat{\alpha}_n$ and $\hat{\rho}_n$
A: Kinked specification (Equation 1)




Figure 6 (cont.)

B: Smooth specification (Equation 2) $\hat{\alpha}_n \leq 2$




Figure 6 (cont.)

C: Smooth specification (Equation 2) $\hat{\alpha}_n \leq 1/2$






Figure 7: The relationship between log-price ratio $\ln(p_1 / p_3)$ and estimated token share $\hat{x}_1 / (\hat{x}_1 + \hat{x}_3)$

B: ID 20



C: ID 31







Figure 7 (cont.)

D: ID 23



E: ID 12





Figure 7 (cont.)

F: ID 37

