

Estimating ambiguity aversion in a portfolio choice experiment

DAVID AHN

Department of Economics, University of California, Berkeley

SYNGJOO CHOI

Department of Economics, University College London

DOUGLAS GALE

Department of Economics, New York University

SHACHAR KARIV

Department of Economics, University of California, Berkeley

We report a portfolio-choice experiment that enables us to estimate parametric models of ambiguity aversion at the level of the individual subject. The assets are Arrow securities that correspond to three states of nature, where one state is risky with known probability and two states are ambiguous with unknown probabilities. We estimate two specifications of ambiguity aversion, one kinked and one smooth, that encompass many of the theoretical models in the literature. Each specification includes two parameters: one for ambiguity attitudes and another for risk attitudes. We also estimate a three-parameter specification that includes an additional parameter for pessimism/optimism (underweighting/overweighting the probabilities of different payoffs). The parameter estimates for individual subjects exhibit considerable heterogeneity. We cannot reject the null hypothesis of subjective expected utility for a majority of subjects. Most of the remaining subjects exhibit statistically significant ambiguity aversion or seeking and/or pessimism or optimism.

KEYWORDS. Uncertainty, ambiguity aversion, risk aversion, pessimism/optimism, subjective expected utility, maxmin expected utility, α -maxmin expected utility, Choquet expected utility, contraction expected utility, recursive expected utility, recursive nonexpected utility, rank-dependent utility, experiment.

JEL CLASSIFICATION. C91, D81.

David Ahn: dahn@econ.berkeley.edu

Syngjoo Choi: syngjoo.choi@ucl.ac.uk

Douglas Gale: douglas.gale@nyu.edu

Shachar Kariv: kariv@berkeley.edu

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 James Andreoni, Yoram Halevy, Glenn Harrison, Thomas Palfrey, Chris Shannon, Peter Wakker, and William Zame for helpful discussions. The paper has benefited from suggestions from the participants of seminars at several universities. We acknowledge the National Science Foundation Grants SES-0550224 (Ahn) and SES-0617955 (Gale and Kariv), and the Coleman Fung Risk Management Research Center (Open-Link Fund) at UC Berkeley (Kariv) for financial support.

Copyright © 2014 David Ahn, Syngjoo Choi, Douglas Gale, and Shachar Kariv. Licensed under the [Creative Commons Attribution-NonCommercial License 3.0](http://creativecommons.org/licenses/by-nc/3.0/). Available at <http://www.qeconomics.org>.

DOI: [10.3982/QE243](https://doi.org/10.3982/QE243)

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 confirmed Ellsberg's conjecture, and a large theoretical literature has developed models that are consistent with this behavior.

In this paper, we use experimental data to estimate models of attitudes toward risk and ambiguity. The data are generated by subjects solving a series of portfolio-choice problems. In the experimental design, there are three states of nature, denoted by $s = 1, 2, 3$. For each state s , there is an Arrow security that pays \$1 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 letting 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 nonnegative portfolio $\mathbf{x} \geq \mathbf{0}$ that satisfies the budget constraint. The budget sets are displayed on a computer screen using the graphical interface introduced by Choi, Fisman, Gale, and Kariv (2007a) and exploited by Choi, Fisman, Gale, and Kariv (2007b) for the study of risky decisions.

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 (Section 6.1) that can be derived as a special case of a variety of utility models in the literature, including maxmin expected utility (MEU), Choquet expected utility (CEU), α -maxmin expected utility (α -MEU), and contraction expected utility. The second is a "smooth" specification (Section 6.2) that can be derived from the class of recursive expected utility (REU) models.¹ The kinked and smooth specifications are both characterized by two parameters: one is a measure of ambiguity aversion and the other is the familiar coefficient of risk aversion.

Figure 1 illustrates the indifference curves between the securities that pay off in the ambiguous states x_1 and x_3 for the kinked specification. The indifference curves have a *kink* at the 45-degree line. The shape of the indifference curve on either side of the 45-degree line is determined by the individual's attitude toward risk and the nature of the kink is determined by the individual's attitude toward ambiguity. The left hand diagram depicts a typical indifference curve for an ambiguity-averse individual. The right hand diagram depicts the indifference curve for an ambiguity-seeking individual. Note that an ambiguity-averse individual chooses unambiguous portfolios that

¹See, MEU: Gilboa and Schmeidler (1989); CEU: Schmeidler (1989); α -MEU: Ghirardato, Maccheroni, and Marinacci (2004) and Olszewski (2007); contraction expected utility: Gajdos, Hayashi, Tallon, and Vergnaud (2008); REU: Ergin and Gul (2004), Klibanoff, Marinacci, and Mukerji (2005), Nau (2006), Seo (2009), and related work by Halevy and Feltkamp (2005), Giraud (2005), and Ahn (2008).




FIGURE 1. An illustration of an indifference curve between the securities that pay off in the ambiguous states, x_1 and x_3 , in the kinked specification (equation (1)).

satisfy $x_1 = x_3$ when the security prices p_1 and p_3 are sufficiently similar. In contrast, an ambiguity-seeking individual does not choose such unambiguous portfolios, not even when the security prices are equal. For the smooth specification, in contrast, the indifference curves between any pair of securities are smooth everywhere. The individual's attitude toward ambiguity changes the shape of the indifference curves between the securities that pay off in the ambiguous states x_1 and x_3 in a manner qualitatively similar to increased/decreased risk aversion.

Unlike Ellsberg-type experiments, 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 ambiguous states. When $x_1 = x_3$, there is no effective exposure to ambiguity. Similarly, exposure to ambiguity becomes small as x_1 and x_3 approach one another. Although we find considerable heterogeneity of attitudes toward ambiguity across subjects, one of the striking features of the data is that for some subjects, the chosen portfolios tend to be much more concentrated around the unambiguous portfolios $x_1 = x_3$ than they are around the portfolios $x_1 = x_2$ and $x_2 = x_3$. Since the design is symmetrical in all other respects, one can attribute this greater concentration to ambiguity aversion. For other subjects, the data show evidence of ambiguity seeking, which reveals itself behaviorally by a lesser tendency to equalize demands for x_1 and x_3 than the other pairs of securities.

The tendency to choose unambiguous portfolios $x_1 = x_3$ is not only evidence suggestive of ambiguity aversion; it also allows us to discriminate between the kinked and the smooth specifications. Since the price vectors are continuously distributed, this event should occur with probability close to zero if the subject's preferences can be represented by a smooth utility function. This feature of the data is therefore consistent with the kinked specification, but not with the smooth specification.

We also find there is a strong tendency for some subjects to equalize their demands for the other pairs of securities, x_1 and x_2 or x_2 and x_3 , where only one of the securities

pays off in an ambiguous state. This feature of the data cannot be accommodated by either the kinked or smooth specifications, but is consistent with “pessimism” with respect to probabilities as proposed by Quiggin (1982). The tendency to equate the demands for securities that pay off in the ambiguous states could also result from pessimism, but the tendency to equalize demands is greater for x_1 and x_3 than the other pairs of securities, suggesting an aversion to ambiguity rather than just pessimism.

To distinguish ambiguity aversion from pessimism (overweighting the probabilities of low payoffs and underweighting the probabilities of high payoffs), we are led to consider a *generalized* kinked specification (Section 8) that builds on Quiggin’s (1982) rank-dependent utility (RDU). This three-parameter specification allows for ambiguity aversion or seeking as well as pessimism or optimism. Its key feature is that indifference curves may have kinks at portfolios with equal amounts of any pair of securities. The parameters of the generalized kinked specification provide measures of risk aversion, ambiguity attitudes, and pessimism/optimism.

Table 1 provides a population-level summary of the individual-level estimates by classifying subjects as ambiguity seeking, neutral, or averse and pessimistic, neutral, or optimistic using a 5 percent significance level. The estimation results for the generalized kinked specification reinforce the conclusion from the kinked and smooth specifications that preferences vary widely across subjects. More interestingly, for 64.3 percent of subjects, we cannot reject the null hypothesis of preferences consistent with SEU. Most of the remaining subjects have a statistically significant degree of ambiguity aversion or seeking and/or pessimism or optimism: 13.0 percent are classified as either averse to or seeking ambiguity and 24.7 percent of the subjects are classified as either pessimistic or optimistic.

The rest of the paper is organized as follows. Section 2 provides a discussion of some related literature. Section 3 describes the experimental design and procedures. Section 4 summarizes some important features of the data. Section 5 examines some implications of revealed preference tests. Section 6 describes the kinked and the smooth specifications, and Section 7 provides the estimation results. Section 8 describes the generalized

TABLE 1. Subjects’ classifications (fractions of subjects).

	Ambiguity			Total
	Seeking	Neutral	Averse	
Pessimism/optimism				
Optimistic	0.006	0.123	0.013	0.143
Neutral	0.006	0.643	0.065	0.714
Pessimistic	0.013	0.065	0.026	0.104
Total	0.026	0.831	0.104	0.961

We estimated the three specifications for a total of 154 subjects. In the generalized kinked specification (equation (3)), we first test, subject-by-subject, for consistency with SEU using Wald test since the null hypothesis involves the joint significance of more than one coefficient. The attitudes toward ambiguity and disappointment of subjects for whom we reject the null hypothesis of SEU are classified using *t*-tests. Only six subjects (3.9 percent) cannot be cleanly classified.

kinked specification and Section 9 provides the results of the estimation. Section 10 summarizes the results and contains some concluding remarks. The paper also uses several data and technical appendices, available for the interested reader in a supplementary file on the journal website, <http://qeconomics.org/supp/243/supplement.pdf>. Replication files are available on the journal website, http://qeconomics.org/supp/243/code_and_data.zip.

2. RELATED LITERATURE

We will not attempt to review the large and growing experimental literature on ambiguity aversion. [Camerer and Weber \(1992\)](#) and [Camerer \(1995\)](#) provided 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 \(2007\)](#) presented a cleverly designed experiment that allows him to distinguish between four models of ambiguity aversion—SEU, MEU, CEU, and REU—and recursive nonexpected utility (RNEU) proposed by [Segal \(1987, 1990\)](#). Subjects are asked their reservation values for four different urns, which represent different types of uncertainty (pure risk, pure ambiguity, and two types of compound lotteries). The different models generate different predictions about how the urns will be ordered, based on whether and how reduction of compound lotteries may fail. For each subject, there will be a unique model that predicts (is consistent with) the subject's reservation values. [Halevy \(2007\)](#) pointed out that there is “a tight association between ambiguity neutrality and reduction of compound lotteries,” which is consistent with SEU. His main finding is that all models are represented in the pool of subjects but only about 20 percent of the subjects behave as if they were ambiguity neutral.

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 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. 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. The parameter estimates allow us to measure individual heterogeneity within as well as across the different specifications.

[Bossaerts, Ghirardato, Guarnaschelli, and Zame \(2010\)](#) studied the impact of ambiguity on portfolio holdings and asset prices in a financial market experiment. The experimental procedures were adapted from those used by [Bossaerts, Plott, and Zame \(2007\)](#) to study markets with pure risk. [Bossaerts et al. \(2010\)](#) pointed out that there is substantial heterogeneity in attitudes to ambiguity and that this heterogeneity has important effects on equilibrium prices and portfolio holdings. One important finding is that traders

who are sufficiently ambiguity averse refuse to hold securities with ambiguous payoffs. In the present paper, we restrict attention to pure (individual) decision problems, thus avoiding the complications of strategic behavior and arbitrage opportunities in laboratory financial markets. This allows us to focus on the estimation of parameters that measure individual attitudes to risk and ambiguity. The use of market-generated data for this purpose is problematic because beliefs and other crucial variables are unobserved.

A few recent studies use variants of Ellsberg-type decision problems to estimate parameter values or functional forms for individual subjects. Abdellaoui, Baillon, Placido, and Wakker (2011) captured attitudes to ambiguity by fitting source functions. A source function converts a subjective (choice-based) probability into a willingness to bet. The authors found considerable heterogeneity in subjects' preferences, both in an Ellsberg-type experiment and in experiments using naturally occurring uncertainties. Finally, Machina (2009) introduced variants of the Ellsberg problem to test the CEU model. Baillon, L'Haridon, and Placido (2011) showed that the thought experiments proposed by Machina (2009) would also lead to violations of other theories, including MEU, α -MEU, and REU, and confirm these conjectures in the laboratory.

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. 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, except for the number of assets and the presence of ambiguity. 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 and chose portfolios through a simple "point and click." Full experimental instructions, including the computer program dialog window, are available in Appendix I.

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.

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. 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 one of the accounts was selected with probability $\frac{1}{3}$, whereas the other two accounts were selected with unknown probabilities that sum to $\frac{2}{3}$.² The unambiguous account was held constant throughout a given experimental session, but its labeling, x , y , or z , was changed across sessions.

During the course of the experiment, subjects were not provided with any information about the account that had been selected in each round. 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.

To describe preferences with some precision at the level of the individual subject, it is necessary to generate many observations per subject over a wide range of budget sets. Our subjects made decisions over 50 budget sets and one decision round was selected at random to be used for payoffs. If we paid for all rounds, subjects would have faced 50 independent “three-way coin flips,” and could easily hedge against both risk and ambiguity, like Samuelson (1963) reports his colleague did. Bade (2013) showed that random incentive mechanisms prevent hedging only under stringent independence conditions, and she provided a simple hedging opportunity example involving two budget sets defined by the prices $\mathbf{p}^1$ and $\mathbf{p}^2$. However, the example relies heavily on the fact that the individual knows the second budget constraint $\mathbf{p}^2$ when he chooses a portfolio subject to the first budget constraint $\mathbf{p}^1$. In our experiment, subjects face a large menu of highly heterogeneous budget sets, and are only informed about the prices' random generating process, making it greatly more difficult to hedge.

4. DATA DESCRIPTION

4.1 Aggregate behavior

A subject can completely avoid ambiguity 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

²In practice, the probability of one of the “ambiguous” states 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.

reduces exposure to ambiguity without eliminating it altogether. A tendency to equate x_1 and x_3 could also result from simple risk aversion or pessimism, but this tendency would apply to all pairs of securities. We use this fact to detect the presence of ambiguity aversion in the aggregate data. We can also detect the presence of ambiguity seeking in the aggregate data. If the concentration of aggregate choices around the ray $x_1 = x_3$ is greater than the concentration around $x_1 = x_2$ and $x_2 = x_3$, this could be taken as evidence of ambiguity-averse behavior. Conversely, ambiguity-seeking individuals will avoid the ray $x_1 = x_3$. If the aggregate data contain a mixture of ambiguity-averse and ambiguity-seeking behavior, this will reduce the concentration of choices around the ray $x_1 = x_3$.

For any portfolio $\mathbf{x} = (x_1, x_2, x_3)$ and any pair of states s and $s' \neq s$, we define the *relative demand* for the security that pays off in state s to be the ratio $x_s/(x_s + x_{s'})$. Figure 2 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 $x_3/(x_2 + x_3)$. The latter two ratios measure the extent to which subjects equalize payoffs in two states, exactly one of which is ambiguous. Before calculating these densities, we remove data points that correspond to safe ($x_1 = x_2 = x_3$) and risk neutral ($x_s = x_{s'} = 0$ for some $s \neq s'$) portfolios. These portfolios could be explained by infinite risk aversion and pure risk neutrality, respectively, so including these portfolios would inflate our measures of ambiguity aversion. To account for small mistakes that result from the slight imprecision of subjects' handling of the mouse, we allow




FIGURE 2. The distribution of relative demands (kernel density estimates). Kernel density estimates of the relative demands. *Black*: The relative demand for the securities that pay off in ambiguous states, 1 and 3. *Gray*: The relative demands for two securities, exactly one of which pays off in an ambiguous state. The numbers are calculated after screening out safe and risk neutral portfolios using a narrow confidence interval of two tokens. The safe and risk neutral portfolios account for 20.0 and 6.5 percent of all portfolios, respectively. Including these portfolios in the sample would have biased our relative-demand measures of ambiguity aversion and disappointment aversion.

for a two-token neighborhood in defining these portfolios.³ The safe and risk-neutral portfolios account for 20.0 and 6.5 percent of all portfolios, respectively.

The distributions are nearly symmetric and concentrated around the midpoint $\frac{1}{2}$.⁴ More interestingly, the mode is more pronounced in the distribution of relative demand for the securities that pay off in ambiguous states, x_1 and x_3 . This provides suggestive evidence of ambiguity aversion as a motivation for hedging, in addition to risk aversion and pessimism. 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 that lie 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. Note that if our subjects were leaning overall toward ambiguity seeking, then the mode in the distribution of relative demands for the securities that pay off in ambiguous states would have been less pronounced. The distributions of relative demands suggest that the aggregate data lean overall toward ambiguity aversion rather than ambiguity seeking.

We continue the analysis of the experimental data by looking for more signs of ambiguity aversion, as well as ambiguity seeking, in the aggregate data. 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 $x_s/(x_1 + x_2 + x_3)$, that is, the number of tokens payable in state s as a fraction of the sum of tokens payable in all three states. Simulations of the kinked specification show that, other things being equal, ambiguity-seeking individuals will, on average, allocate a smaller token share to the cheapest security when that security pays off in the unambiguous state than when it pays off in one of the ambiguous states. For individuals who exhibit ambiguity aversion, the prediction is less clear but intuition and simulations suggest that they should, on average, allocate a larger token share to the cheapest security when that security pays off in the unambiguous state than when it pays off in one of the ambiguous states.⁵

Figure 3 below depicts kernel density estimates of the token share of the cheapest security. We distinguish between portfolios where the cheapest security pays off in one of the ambiguous states (black) and portfolios where the cheapest security pays off in the unambiguous state (gray). The distributions in Figure 3 are quite similar: both have a mode near $\frac{1}{3}$ and the distribution falls off sharply away from $\frac{1}{3}$, with mass concentrated to the right.⁶ However, the mode around $\frac{1}{3}$ is much more pronounced when the cheap-

³A portfolio $\mathbf{x}$ is defined as safe if $|x_s - x_{s'}| \leq 2$ for all s, s' and as boundary if $\min\{x_s, x_{s'}\} \leq 2$ for some $s \neq s'$. We generate virtually identical results when we allow for a one-token neighborhood.

⁴The distribution of relative demand for the securities that pay off in ambiguous states, x_1 and x_3 , is significantly different using a two-sample Kolmogorov–Smirnov test (p -value = 0.000).

⁵The kinked specification in equation (1) considers an individual who evaluates each portfolio by a convex combination of its worst and best expected utilities using the weights α and $1 - \alpha$, respectively. Hence, $\alpha \in [0, 1]$ serves as a parameter that reflects attitudes toward ambiguity. The individual is ambiguity averse if $\alpha > \frac{1}{2}$, ambiguity neutral if $\alpha = \frac{1}{2}$, and ambiguity seeking if $\alpha < \frac{1}{2}$. Assuming that the set of subjective probabilities is consistent with the objective information, when the cheapest security pays off in an ambiguous state, the individual is assumed to assign the probability belief $\frac{2}{3}(1 - \alpha)$ to that state, which is thus perceived to be higher (lower) than $\frac{1}{3}$ by an ambiguity-seeking (ambiguity-averse) individual. It is this skewing of the probabilities that causes the individual to allocate more or less to the cheapest security depending on whether it pays off in an ambiguous or unambiguous state.

⁶The two distributions are significantly different using a two-sample Kolmogorov–Smirnov test (p -value = 0.000).




FIGURE 3. The distribution of the fraction of tokens allocated to the cheapest security (kernel density estimates). Kernel density estimates of the token share of the cheapest security (the security with the lowest price). *Black*: Portfolios where the cheapest security pays off in one of the *ambiguous* states, 1 or 3. *Gray*: Portfolios where the cheapest security pays off in the *unambiguous* state.

est security pays off in one of the ambiguous states, which provides more evidence of ambiguity aversion rather than ambiguity seeking.

To make this suggestion more precise, we generate a benchmark against which to compare our results using the choices of ambiguity-neutral and ambiguity-averse simulated subjects who maximize the kinked specification in equation (1) below using a range of parameter values.⁷

We present the results in Appendix II. Each figure shows the kernel density estimates of the token share of the cheapest security depicted in Figure 3 with the data generated by a sample of simulated subjects that make choices from the same set of budget sets as the human subjects do (each panel assumes a different ambiguity parameter). Using the choices of ambiguity-neutral simulated subjects, the two distributions are virtually identical. As the simulated subjects become more ambiguity averse, the concentration around $\frac{1}{3}$ becomes more pronounced when the cheapest security pays off in one of the ambiguous states.

4.2 Individual behavior

The aggregate data above tell us little about the choice behavior of individual subjects. The primary innovation of our experimental design is that it allows us to analyze behavior at the level of the individual subject. For that purpose, Figure 4 presents within-subject comparisons of the number of unambiguous portfolios for which $x_1/(x_1 + x_3)$

⁷We assume that risk preferences are represented by a von Neumann–Morgenstern utility function 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.




FIGURE 4. Scatterplot of 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. Again, the numbers are calculated after screening out safe and risk neutral portfolios using a narrow confidence interval of two tokens. This results in many fewer observations for a small number of subjects who often choose these portfolios.

lies between 0.45 and 0.55 (vertical axis) 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 (horizontal axis). Before calculating these numbers, we again exclude safe and risk-neutral portfolios. This results in many fewer observations for a small number of subjects who often choose these portfolios.

Figure 4 confirms the heterogeneity of individual behavior and the prevalence of portfolios for which $0.45 \leq x_s/(x_s + x_{s'}) \leq 0.55$ for any s and $s' \neq s$ in the individual-level data. Of the 154 subjects, 87 subjects (56.5 percent) choose portfolios in these bounds in at least half of the 50 decision rounds. Of these, 19 subjects (12.3 percent) choose such portfolios in at least 40 rounds and seven subjects (4.5 percent) in 45 or more rounds. Although there is a large amount of heterogeneity, the most notable feature of the scatter plot in Figure 4 is that the data are concentrated above the diagonal and skewed to the upper left (the vertical axis corresponds to the number of unambiguous portfolios). This provides evidence on both the prominence and the heterogeneity of subjects' attitudes toward ambiguity. Also note that there is little evidence of ambiguity-seeking subjects, as those would fall close to the horizontal axis (many more portfolios for which $x_1/(x_1 + x_2)$ or $x_2/(x_2 + x_3)$ is close to $\frac{1}{2}$ than unambiguous portfolios for which $x_1/(x_1 + x_3)$ is close to $\frac{1}{2}$).

Finally, Figure 5 presents within-subject comparisons of the average token shares of the cheapest security, $x_s/(x_1 + x_2 + x_3)$, where $p_s < p_{s'}$ for any $s' \neq s$. We again distinguish between portfolios where the cheapest security pays off in one of the ambiguous states (vertical axis) and portfolios where the cheapest security pays off in the unambiguous state (horizontal axis). Again, ambiguity aversion (seeking) can reveal itself behaviorally




FIGURE 5. Scatterplot of the average fraction of tokens allocated to the cheapest security by subject. *Vertical axis*: The average fraction of tokens allocated to the cheapest security (the security with the lowest price) when it pays off in an *ambiguous* state. *Horizontal axis*: The average fraction of tokens allocated to the cheapest when it pays off in the *unambiguous* state.

by allocating, on average, more (less) tokens to the cheapest security when it pays off in the unambiguous state. Note that since there are two ambiguous states and only one unambiguous state, the cheapest price conditional on the state being ambiguous will tend to be lower than the cheapest price conditional on the state being unambiguous. The bias that results from the asymmetry of the two conditional distributions of the cheapest price can skew the data above the diagonal, incorrectly identifying subjects as ambiguity seeking.⁸

Although there is considerable heterogeneity, the data in Figure 5 are concentrated close to the diagonal and skewed to the lower right. Of the 154 subjects, 69 subjects (44.8 percent) are located within a narrow 0.05 band around the diagonal and this number increases to 112 subjects (72.7 percent) if we double the bandwidth. Of the remaining subjects on the graph, 30 subjects (32.7 percent) are located below the diagonal, which reflects aversion to ambiguity. On the other hand, only 12 subjects (7.8 percent) are located above the diagonal. Overall, the data in Figure 5 clearly illustrate the extent to which subjects lean toward ambiguity-neutral and ambiguity-averse behavior.

We again generate a benchmark level using the choices of simulated subjects who maximize the kinked specification in equation (1) using various parameter values for attitudes toward risk and ambiguity. We present the results in Appendix II. The data are generated by a sample of simulated subjects that make choices from the same set of

⁸The number of budget sets for which the cheapest asset pays off in the unambiguous state varies widely across subjects and ranges from 9 to 26 budget sets. Pooling the data, in 2543 budget sets (33.0 percent), the cheapest asset pays off in the unambiguous state, and in 4457 budget sets, it pays off in one of the ambiguous states. Given our large and rich menu of budget sets, the aggregate distributions of the cheapest price are virtually identical.

budget sets as the human subjects do (each panel assumes a different ambiguity parameter). The scatter plots provide a clear graphical illustration that, given the randomness in security prices, behaviors in our narrow bandwidth around the diagonal are consistent with ambiguity-neutral behavior. We thus conclude that both ambiguity-neutral and ambiguity-averse behaviors are well represented among our subjects, whereas there is little evidence of ambiguity-seeking behavior.

In Appendix III we depict the portfolios chosen by individual subjects in terms of token shares $x_s/(x_1 + x_2 + x_3)$ and *budget shares* $p_s x_s$ (prices are normalized by income so that $\mathbf{p} \cdot \mathbf{x} = 1$) for the three securities as points in the unit simplex. In addition, we 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 $x_1/(x_1 + x_2)$. These scatter plots illustrate the sensitivity of portfolio decisions to changes in relative prices. For many subjects, we can easily pick out the portfolios that satisfy $x_s = x_{s'}$ for intermediate relative prices corresponding to $\log(p_s/p_{s'})$ in a neighborhood of zero. This is what ambiguity-averse and/or pessimistic individuals would do. In contrast, we cannot readily detect subjects who choose portfolios that satisfy $x_1 = x_2$ or $x_2 = x_3$ for values of $\log(p_1/p_2)$ or $\log(p_2/p_3)$ close to zero, but do not choose unambiguous portfolios that satisfy $x_1 = x_3$ when $\log(p_1/p_3)$ is close to zero, which is what optimistic and ambiguity-seeking individuals would do.

5. REVEALED PREFERENCES

The most basic question to ask about choice data is whether they are consistent with individual utility maximization. In principle, the presence of ambiguity could cause not just a departure from SEU, but a more fundamental departure from rationality in the sense of a complete, transitive preference ordering. Thus, before postulating particular utility functions, we first test whether choices can be utility-generated. Afriat (1967) and Varian (1982, 1983) established the generalized axiom of revealed preference (GARP) as a direct test for whether the finite set of observed price and quantity data that our experiment generated may be rationalized by a utility function.

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 so as 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, closer to perfect consistency with utility maximization.

Over all subjects, the CCEI scores averaged 0.945. Out of the 154 subjects, 127 subjects (82.5 percent) had CCEI scores above the 0.90 threshold, and of those, 93 subjects (60.4 percent) were above the 0.95 threshold. Choi et al. (2007a) demonstrated that if utility maximization is not in fact the correct model, then the experiment is sufficiently powerful to detect it. We thus interpret these numbers as a confirmation that subject

choices are generally consistent with utility maximization.^{9,10} Throughout the remainder of the paper, we present results for all subjects, and for those with CCEI scores above 0.9 and 0.95 in parallel. Appendix IV lists, by subject, the number of violations of GARP, and also reports the values of different goodness-of-fit measures. In practice, all these measures yield similar conclusions.

6. THE KINKED AND SMOOTH SPECIFICATIONS

In this section, we introduce two parsimonious specifications of utility functions. 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. Each of our specifications is characterized by two parameters, one of which is a measure of ambiguity aversion and the other of which is the familiar coefficient of risk aversion. We skip the models’ development and analysis, and instead focus on the restrictions—in the form of specific functional form assumptions—that we place on the general models so as to generate specific parametric formulations amenable to analysis. We refer the interested reader to Appendix V for more details.

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.

6.1 The kinked specification

The kinked specification is so called because the indifference curves have a “kink” at all unambiguous portfolios $x_1 = x_3$. Following the α -MEU model, we assume that the unknown probabilities π_1 and π_3 are skewed using the weights $0 \geq \alpha \geq 1$ and $1 - \alpha$ for the low and high payoffs, respectively. That is, the lower of the two payoffs $x_{\min} = \min\{x_1, x_3\}$

⁹We refer the interested reader to Choi et al. (2007a, 2007b) for more details on testing for consistency with GARP, and for other measures that have been proposed for this purpose by Varian (1990, 1991) and Houtman and Maks (1985). The papers by Afriat (2012), Diewert (2012), Varian (2012), and Vermeulen (2012), published in a special volume of the *Economic Journal*, provide an excellent overview and a discussion of some recent developments in the literature.

¹⁰Diewert (2012) started developing Afriat-type inequalities for testing for expected utility maximization. We would have liked to use revealed preference methods to test whether the data are consistent with a utility function with some special structure, particularly the kinked and smooth specifications, but we came to the conclusion that the answer is not a simple adjustment to the usual tests, which are all computationally intensive for large data sets like our own. This is an interesting avenue for future work, but is beyond the scope of this paper.

is given a probability weight $\frac{2}{3}\alpha$ and the higher of the two payoffs $x_{\max} = \max\{x_1, x_3\}$ is given probability weight $\frac{2}{3}(1 - \alpha)$. The utility of a portfolio $\mathbf{x} = (x_1, x_2, x_3)$ takes the form

$$U(\mathbf{x}; \alpha, \rho) = -\frac{2}{3}\alpha \exp\{-\rho x_{\min}\} - \frac{1}{3} \exp\{-\rho x_2\} - \frac{2}{3}(1 - \alpha) \exp\{-\rho x_{\max}\}, \quad (1)$$

where the parameter α represents the attitude toward ambiguity: $\frac{1}{2} < \alpha \leq 1$ indicates preferences that are ambiguity averse, $0 \leq \alpha < \frac{1}{2}$ indicates preferences that are ambiguity seeking, and if $\alpha = \frac{1}{2}$, we have the standard SEU representation.

As we show in Appendix V, the kinked specification in equation (1) can be generated by different classes of preferences. In all of these theories, the parameter α depends on the set of prior beliefs (or the capacity, in the case of CEU). Unless the set of prior beliefs is objectively known, knowledge of the estimated parameter α does not allow us to characterize the degree of ambiguity aversion independently of the degree of ambiguity in the decision problem. In any case, the lack of identification is inherent to these theoretical models, rather than a feature of our data. To simplify the exposition and facilitate comparisons, in the sequel we adopt the α -MEU interpretation by fixing the set of priors and allowing α to vary.

6.2 The 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

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

where the parameter α measures the attitude toward ambiguity: any $\alpha > 0$ indicates preferences that are ambiguity averse, $\alpha < 0$ indicates preferences that are ambiguity seeking, and as $\alpha \rightarrow 0$, we approach the standard SEU representation.

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. The general form of the REU model and the derivation of the smooth specification in equation (2) are relegated to Appendix V.

While the kinked specification in equation (1) can be interpreted using a variety of different models, the smooth specification is motivated by a single model. We follow [Halevy \(2007\)](#) in referring to this model as REU, owing to its recursive double expectation. One of the crucial features of the REU model is its reliance on a cardinal utility indicator. Unlike the kinked specification, which is invariant to affine transformations of the utility function $u(\cdot)$, the smooth specification is not independent of a change in

the scale of utility. Consequently, the parameters of the smooth specification are not identified without some auxiliary assumption about cardinal utility.

To clarify, if we introduce a scale parameter A and set $u(x) = -Ae^{-\rho x}$, the ranking of portfolios will not be invariant to changes in A under REU. 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, 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 $u(\cdot)$ and, hence, will reduce the ambiguity to which the agent is exposed. Generally, comparison of ambiguity aversion across individuals is a delicate matter. In the REU model, if the concavity of the ambiguity aggregator is naively measured by α , its values can only be sensibly related when risk attitudes are similar. This is because the range of utilities $[-A, -Ae^{-100\rho}]$ is a nonlinear contraction of $[0, 100]$, the range of tokens. We will return to this subject in the next section when we discuss the parameter estimates of the smooth specification.

6.3 Properties of demand

Before presenting the estimation of the kinked and smooth specifications, it is important to understand the implications of these specifications for individual behavior. Figure 1 illustrates the indifference curves between the securities that pay off in the ambiguous states, x_1 and x_3 , for the kinked specification in equation (1). Figure 6 illustrates the relationship between the log-price ratio $\ln(p_1/p_3)$ and the optimal relative demand $x_1^*/(x_1^* + x_3^*)$ for the kinked specification. We distinguish between preferences that are ambiguity averse ($\frac{1}{2} < \alpha \leq 1$) and preferences that are ambiguity seeking ($0 \leq \alpha < \frac{1}{2}$).

If preferences are ambiguity averse and 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. Thus, an ambiguity averse individual (black line) some of the time chooses unambiguous portfolios that satisfy $x_1^* = x_3^*$ and he does this for intermediate relative prices that correspond to $\ln(p_1/p_3)$ around zero. In contrast, an ambiguity-seeking individual (gray line) does not choose unambiguous portfolios that satisfy $x_1^* = x_3^*$, not even when $\ln(p_1/p_3)$ equals zero.

For the smooth specification in equation (2), the relationship between the log-price ratio $\ln(p_1/p_3)$ and the optimal relative demand $x_1^*/(x_1^* + x_3^*)$ is smooth for all price ratios. The ambiguity aversion parameter α flattens (bends) the $x_1^*/(x_1^* + x_3^*)$ curves in a manner qualitatively similar to increased (decreased) risk aversion. In contrast to the kinked specification, the choice of a portfolio without ambiguity $x_1^* = x_3^*$ occurs for a negligible set of prices in the smooth specification.

The relative demand functions for different parameter values are illustrated further in Appendix VI (the figures are difficult to see in the small black and white format in a printed version). For each specification, kinked and smooth, we show the relationships between the log-price ratio $\ln(p_1/p_3)$ and the optimal relative demand $x_1^*/(x_1^* + x_3^*)$ (left




FIGURE 6. An illustration of the relationships between log-price ratio and the optimal token share of the securities that pay off in the ambiguous states in the kinked specification (equation (1)).

panels), and between $\ln(p_1/p_2)$ and $x_1^*/(x_1^* + x_2^*)$ (right panels) using a range of parameter values (each panel assumes a different value for α). For reasons explained below, we restrict the parameters so that preferences are risk averse ($\rho \geq 0$) in both specifications and ambiguity averse ($\alpha \geq 0$) in the smooth specification. A comparison of the two relative demands illustrates the different choices individuals make under the two specifications. These differences are important in understanding how these specifications fit the data in the econometric analysis presented in the next section.

7. ESTIMATION I

7.1 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 $\hat{\alpha}_n$ and $\hat{\rho}_n$ using nonlinear least squares (NLLS). These estimates are chosen to minimize

$$\sum_{i=1}^{50} \|\mathbf{x}^i - \mathbf{x}^*(\mathbf{p}^i; \alpha_n, \rho_n)\|^2,$$

where $\|\cdot\|$ denotes the three-dimensional Euclidean norm and $\mathbf{x}^*(\mathbf{p}^i; \alpha_n, \rho_n)$ denotes the optimal portfolio subject to the budget constraint $\mathbf{p}^i \cdot \mathbf{x} = 1$.¹¹

Before proceeding to estimate the parameters, we make an important remark about the econometric specification. Our subjects are close enough to satisfying GARP that they can be considered utility maximizers, and Afriat's theorem tells us that the underlying utility function that rationalizes the data can be chosen to be well behaved (that is, piecewise linear, continuous, increasing, and concave). The reason for this is that choices subject to linear budget constraints will never be made at points where the underlying utility function is not quasi-concave. Consequently, choices that satisfy GARP can always be treated "as if" they have been generated by risk-averse (or risk-neutral) and ambiguity-averse (or ambiguity-neutral) behavior. In particular, we cannot distinguish between risk- or ambiguity-seeking behavior, on the one hand, and risk- or ambiguity-neutral behavior, on the other, except by exploiting functional form assumptions.¹²

We are therefore forced to restrict the parameters so that preferences are always risk averse in both specifications ($\rho \geq 0$), as otherwise the underlying utility function is quasi-convex everywhere. In the smooth specification in equation (2), we also restrict the parameters so that preferences are also ambiguity averse ($\alpha \geq 0$). The intuitive reason for this is that in the smooth specification, ambiguity seeking and risk seeking are qualitatively identical. Because of computational difficulties when α is large, we also impose the restriction $\alpha \leq 2$ in the smooth specification. This involves minimal loss of fit, since the predicted choices that correspond to $\alpha > 2$ are virtually indistinguishable. We do not restrict the attitudes toward ambiguity in the kinked specification in equation (1) ($0 \leq \alpha \leq 1$). The indifference curves between the securities that pay off in the ambiguous states, x_1 and x_3 for the kinked specification are illustrated in Figure 1.

7.2 Econometric results

7.2.1 The kinked specification To economize on space, the individual-level estimates are relegated to Appendix VII. Table 2 provides a population-level summary of the individual-level estimation results for the kinked specification in equation (1) by reporting summary statistics and percentile values (top panel), and classifies subjects as ambiguity seeking, neutral, or averse using a 5 percent significance level (bottom panel). We present the results for all subjects, as well as for the subsamples of subjects with CCEI scores above 0.90 and 0.95. The patterns are very similar across different CCEI cutoff thresholds, indicating that the results are not driven by inconsistent subjects.

¹¹For simplicity, the estimation technique for both specifications is NLLS, rather than a structural model using maximum likelihood (ML). We favor the NLLS approach because it provides a good fit and offers straightforward interpretation. The NLLS estimation is still computationally intensive for even moderately large data sets.

¹²In the absence of ambiguity, risk-seeking individuals always allocate all their tokens to the cheapest asset. This is also the behavior that would be implied by risk neutrality, so the attitude toward risk is immaterial and, hence, cannot be estimated. In the presence of ambiguity aversion, the implications of risk-seeking behavior are not quite so stark, but the difficulty of identifying the underlying risk preferences remains.

TABLE 2. The kinked specification (equation (1)) estimation results.

	Panel A. Summary Statistics					
	All Subjects		CCEI ≥ 0.9		CCEI ≥ 0.95	
	α	ρ	α	ρ	α	ρ
Mean	0.515	0.195	0.516	0.223	0.508	0.271
s.d.	0.115	0.668	0.123	0.733	0.117	0.850
Percentiles						
5	0.397	0.007	0.365	0.007	0.400	0.007
10	0.420	0.016	0.420	0.016	0.420	0.021
25	0.467	0.031	0.463	0.033	0.462	0.038
50	0.506	0.057	0.510	0.060	0.507	0.062
75	0.562	0.114	0.556	0.129	0.548	0.129
90	0.632	0.223	0.653	0.271	0.617	0.282
95	0.706	0.457	0.717	0.458	0.696	1.235
	Panel B. Subjects' Classification					
	All Subjects		CCEI ≥ 0.9		CCEI ≥ 0.95	
Seeking	0.097		0.118		0.140	
Neutral	0.656		0.630		0.591	
Averse	0.247		0.252		0.269	
No. of obs.	154		127		93	

Of the 154 subjects listed in Appendix VII, we cannot reject the hypothesis that $\alpha_n = \frac{1}{2}$ for a total of 101 subjects (65.6 percent) at the 5 percent significance level.¹³ Of the 53 remaining subjects, 38 subjects (24.7 percent) exhibit significant degrees of ambiguity aversion $\hat{\alpha}_n > \frac{1}{2}$. No subject displayed the maximal ambiguity aversion, $\hat{\alpha} = 1$, the behavior consistent with MEU. We reject the hypothesis that $\alpha_n = 1$ for all ambiguity-averse subjects. The remaining 15 subjects (9.7 percent) exhibit significant degrees of ambiguity seeking $\hat{\alpha}_n < \frac{1}{2}$. Figure 7 shows a scatter plot of the estimates of $\hat{\alpha}_n$ and $\hat{\rho}_n$ from the kinked specification, and illustrates the heterogeneity of preferences that we find.

Finally, a significant fraction of our subjects have moderate coefficients of absolute risk aversion $\hat{\rho}_n$, which are within the range of estimates reported in Choi et al. (2007b). For 69 of the 154 subjects listed in Appendix VII (38.3 percent), $\hat{\rho}_n$ is in the 0–0.05 range, and this increases to a total of 113 subjects (73.4 percent) if we consider the bounds 0–0.1. Our levels of ρ are much lower than standard single-parameter estimates. The reason is that we have estimated two-parameter specifications so that the parameter

¹³In comparison, Halevy (2007) reported that only 28 of his 142 subjects (19.7 percent) behave as if they were ambiguity neutral. This stands in stark contrast to our experiment in which the majority of subjects do not have a significant degree of ambiguity aversion. The experimental designs and analyses are different in a couple of ways. In particular, Halevy (2007) assumed no randomness in individual choices and thus no room for individual-level statistical tests, whereas we estimate functional forms for individual subjects and compare SEU to non-SEU alternatives while allowing for randomness.




FIGURE 7. Scatterplot of the estimated parameters $\hat{\alpha}_n$ (horizontal axis) and $\hat{\rho}_n$ (vertical axis) in the kinked specification (equation (1)).

that measures ambiguity aversion $\hat{\alpha}_n$ “absorbs” aversion to risk.¹⁴ Finally, we note that there is considerable heterogeneity in both parameters $\hat{\alpha}_n$ and $\hat{\rho}_n$, and that their values in the kinked specification are not correlated ($r^2 = -0.029$).

7.2.2 The smooth specification The individual-level estimates of the smooth specification are also relegated to Appendix VII. As noted above, the smooth specification in equation (2) is not invariant to affine utility transformations. Table 3 provides a population-level summary of the individual-level estimation results. We remind the reader that we restrict the parameters in the smooth specification so that preferences are always risk and ambiguity averse ($\rho \geq 0$ and $\alpha \geq 0$). We present the statistics for the estimated raw ambiguity parameter $\hat{\alpha}$ as well as the statistics for different normalized ambiguity parameters $\hat{\alpha}^t = t\hat{\alpha}/(1 - e^t\hat{\rho})$. 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. As alluded to earlier in the paper, the $\hat{\alpha}_n$ coefficients that come out of the smooth specification are directly comparable only across subjects with similar $\hat{\rho}_n$ coefficients. Given the difficulty of interpreting the parameters in the smooth specification, we do not discuss the estimated parameters in further detail.

¹⁴Choi et al. (2007b) estimated two-parameter utility functions based on Gul (1991): one parameter is a coefficient of risk aversion and the other is a measure of disappointment aversion. Because the model has two parameters, to summarize the risk attitudes of their subjects by a single univariate measure, Choi et al. (2007b) computed the risk premium based on the difference between the expected value of a gamble and its certainty equivalent. We refer the interested reader to Choi et al. (2007b) for details on how the two parameters affect the risk premium.

TABLE 3. The smooth specification (equation (2)) estimation results.

	All Subjects		CCEI ≥ 0.9		CCEI ≥ 0.95	
	α	ρ	α	ρ	α	ρ
Mean	0.207	0.099	0.234	0.106	0.238	0.108
s.d.	0.483	0.124	0.524	0.129	0.533	0.131
Percentiles						
5	0.000	0.007	0.000	0.007	0.000	0.006
10	0.000	0.014	0.000	0.014	0.000	0.020
25	0.000	0.030	0.000	0.034	0.000	0.035
50	0.020	0.053	0.021	0.056	0.023	0.059
75	0.123	0.103	0.148	0.106	0.148	0.105
90	0.502	0.258	0.558	0.268	0.502	0.264
95	1.899	0.474	2.000	0.474	2.000	0.490
No. of obs.	154		127		93	

8. THE GENERALIZED KINKED SPECIFICATION

The main evidence for ambiguity aversion in the data is the strong tendency for subjects to equate their demands for the securities that pay off in ambiguous states, x_1 and x_3 . This feature of the data is accommodated by the kinked specification, but is hard to reconcile with the smooth specification. Furthermore, the similar, though weaker, tendency to equate the demand for the securities that pay off in the two states where one of the states is ambiguous and the other is not, x_1 and x_2 or x_2 and x_3 , is also difficult to reconcile with a kinked specification that only incorporates ambiguity aversion. The tendency to equate the demand for all pairs of securities suggests a role for pessimism as proposed by [Quiggin \(1982\)](#). If both ambiguity aversion/seeking and pessimism/optimism are present in the data, we need a structural model so as to disentangle the two effects.

We make use of the RDU model of [Quiggin \(1982\)](#), which is a generalization of the expected utility model that replaces probabilities with *decision weights* when calculating the value of expected utility. We refer the interested reader to Appendix VIII for precise details.¹⁵ To provide some intuition, suppose that the probabilities of all states are objectively known and equal ($\pi_s = \frac{1}{3}$ for all s), and consider a rank-ordered portfolio $\mathbf{x}$ with payoffs $x_L \leq x_M \leq x_H$. The cumulative distribution function of the induced lottery assigns to each payoff the probability of receiving that payoff or anything less, with zero probability assigned to anything less than x_L and probability 1 assigned to anything more than x_H . In the RDU model, a weighting function $w: [0, 1] \rightarrow [0, 1]$ transforms the distribution function into decision weights. The weighting function is assumed to be increasing and satisfies $w(0) = 0$ and $w(1) = 1$. With pure risk, the decision weight of each payoff depends only on its (known) probability and its ranking, and can be expressed in

¹⁵[Diecidue and Wakker \(2001\)](#) explained the intuition of the RDU model for decision under risk [Quiggin \(1982\)](#) and ambiguity [Schmeidler \(1989\)](#).

terms of w as

$$\begin{aligned}\beta_L &= w\left(\frac{1}{3}\right), \\ \beta_M &= w\left(\frac{2}{3}\right) - w\left(\frac{1}{3}\right), \\ \beta_H &= 1 - w\left(\frac{2}{3}\right).\end{aligned}$$

Then the RDU of the rank-ordered portfolio $\mathbf{x}$ takes the form

$$U(\mathbf{x}) = \beta_L u(x_L) + \beta_M u(x_M) + \beta_H u(x_H).$$

To introduce ambiguity, we have to take account of the identity of the ambiguous states as well as the ranking of the payoffs. We again assume that state 2 has an objectively known probability $\pi_2 = \frac{1}{3}$, whereas states 1 and 3 occur with unknown probabilities π_1 and π_3 . We again assume that the unknown probabilities π_1 and π_3 are skewed using the weights $0 \leq \alpha \leq 1$ and $1 - \alpha$ for the low $x_{\min} = \min\{x_1, x_3\}$ and high $x_{\max} = \max\{x_1, x_3\}$ payoffs, respectively. With both risk and ambiguity, the decision weight of each payoff depends on whether its probability is unknown and its ranking can be expressed in terms of w as

$$\begin{aligned}\beta_1 &= w\left(\frac{1}{3}\right), \\ \beta_2 &= w\left(\frac{2}{3}\alpha + \frac{1}{3}\right), \\ \beta_3 &= w\left(\frac{2}{3}\alpha\right), \\ \beta_4 &= w\left(\frac{2}{3}\right).\end{aligned}$$

By substituting these four parameters into the preceding formula and adopting the CARA utility function $u(x) = -e^{-\rho x}$, the utility of a portfolio $\mathbf{x} = (x_1, x_2, x_3)$ takes the form

$$\begin{aligned}\text{I. } x_2 &\leq x_{\min} \\ &\quad -\beta_1 \exp\{-\rho x_2\} - (\beta_2 - \beta_1) \exp\{-\rho x_{\min}\} - (1 - \beta_2) \exp\{-\rho x_{\max}\}, \\ \text{II. } x_{\min} &\leq x_2 \leq x_{\max} \\ &\quad -\beta_3 \exp\{-\rho x_{\min}\} - (\beta_2 - \beta_3) \exp\{-\rho x_2\} - (1 - \beta_2) \exp\{-\rho x_{\max}\}, \\ \text{III. } x_{\max} &\leq x_2 \\ &\quad -\beta_3 \exp\{-\rho x_{\min}\} - (\beta_4 - \beta_3) \exp\{-\rho x_{\max}\} - (1 - \beta_4) \exp\{-\rho x_2\}.\end{aligned}\tag{3}$$

As we show in Appendix VIII, the utility function is defined on three regions that correspond to different rankings of the state-contingent payoffs. Without loss of generality, we can normalize the β coefficients to sum to 1 in each region. The continuity of the utility function at the boundaries between regions imposes additional restrictions, so that the nine potential coefficients are reduced to four independent parameters.

We could estimate the five parameters of the RDU model defined above, but apart from the computational difficulty of such an undertaking, the interpretation of the resulting coefficients would also present problems as ambiguity aversion/seeking and pessimism/optimism cannot be identified with any single parameter. For this reason, we adopt a simpler three-parameter model, in which the parameter δ measures the ambiguity attitudes, the parameter γ measures pessimism/optimism, and ρ is the coefficient of absolute risk aversion. The mapping from the two parameters δ and γ to the four parameters $\beta_1, \dots, \beta_4$ is given by the equations

$$\begin{aligned}\beta_1 &= \frac{1}{3} + \gamma, \\ \beta_2 &= \frac{2}{3} + \gamma + \delta, \\ \beta_3 &= \frac{1}{3} + \gamma + \delta, \\ \beta_4 &= \frac{2}{3} + \gamma,\end{aligned}$$

where $-\frac{1}{3} \leq \delta$, $\gamma \leq \frac{1}{3}$, and $-\frac{1}{3} \leq \delta + \gamma \leq \frac{1}{3}$ so that the decision weight attached to each payoff in equation (3) is nonnegative.

By substituting these values into the generalized kinked specification in equation (3), we obtain our third utility specification $U(\mathbf{x}; \delta, \gamma, \rho)$. The “generalized” kinked utility function is so called because the indifference curves have a kink at all portfolios where $x_s = x_{s'}$ for some $s \neq s'$. The kinked specification in equation (1) is a special case when $\gamma = 0$. Conversely, when $\delta = 0$, we obtain the RDU model of Quiggin (1982). When $\delta = 0$ and $\gamma = 0$, we have the standard SEU representation. Otherwise, the indifference curves will have kinks where $x_s = x_{s'}$ for some $s \neq s'$ and individuals will choose portfolios that satisfy $x_s = x_{s'}$ for a nonnegligible set of price vectors.¹⁶

With the generalized kinked specification, maximizing the utility function subject to the budget constraint yields a relationship between the log-price ratio $\ln(p_s/p_{s'})$ and the optimal relative demand $x_s^*/(x_s^* + x_{s'}^*)$ similar to the relationship illustrated in Figure 6, for any pair of securities s and s' . Estimates of the parameters δ and γ will identify the kinks. For example, if $\gamma > 0$, there is a kink at all portfolios where $x_s = x_{s'}$ for any s and s' , and if $\delta > 0$, the kink at portfolios where $x_1 = x_3$ is sharper than the kink at portfolios where $x_1 = x_2$ and $x_2 = x_3$. The simulated demand functions for the generalized kinked specification are illustrated in Appendix IX. We again show 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),

¹⁶Through a suitable change of variables, we can also interpret the generalized kinked specification as reflecting RNEU where the ambiguity is modeled as an *equal* probability that $\pi_1 = \frac{2}{3}$ or $\pi_3 = \frac{2}{3}$. The derivation is included in Appendix VIII.

and between $\ln(p_1/p_2)$ and $x_1^*/(x_1^* + x_2^*)$ (right panels) using a range of parameter values (each panel assumes different value for δ and γ).

9. ESTIMATION II

For the reasons given above, we again restrict the parameters so that preferences are always risk averse ($\rho \geq 0$). We also restrict $-\frac{1}{3} \leq \delta$, $\gamma \leq \frac{1}{3}$, and $-\frac{1}{3} \leq \delta + \gamma \leq \frac{1}{3}$. The estimation method is a direct extension of the procedure for the kinked specification in equation (1). The individual-level estimates are relegated to Appendix X. Table 4 provides a population-level summary of the individual-level estimation results for the generalized kinked specification in equation (3) (top panel), and classifies subjects' attitudes toward ambiguity and levels of pessimism/optimism using a 5 percent significance level (bottom panel). We first test, subject by subject, for consistency with SEU using the Wald test, since the null hypothesis involves the joint significance of the two parameters, $\hat{\delta}_n$ and $\hat{\gamma}_n$. The attitudes toward ambiguity and pessimism/optimism of subjects for whom we reject the null hypothesis of SEU are classified using t -tests. We again present the results for all subjects, as well as for the subsamples of subjects with CCEI scores above 0.90 and 0.95, and find that the patterns are very similar across different CCEI cutoff thresholds.

TABLE 4. The generalized kinked specification (equation (3)) estimation results.

	Panel A. Summary Statistics								
	All Subjects			CCEI ≥ 0.9			CCEI ≥ 0.95		
	δ	γ	ρ	δ	γ	ρ	δ	γ	ρ
Mean	0.014	0.006	0.157	0.014	0.010	0.177	0.006	0.011	0.210
s.d.	0.070	0.136	0.579	0.071	0.140	0.635	0.061	0.146	0.736
Percentiles									
5	-0.077	-0.188	0.005	-0.077	-0.185	0.006	-0.088	-0.188	0.005
10	-0.054	-0.141	0.011	-0.041	-0.141	0.011	-0.026	-0.138	0.010
25	-0.012	-0.073	0.022	-0.010	-0.073	0.022	-0.010	-0.076	0.030
50	0.006	-0.009	0.051	0.006	-0.009	0.057	0.003	-0.013	0.068
75	0.037	0.063	0.100	0.036	0.063	0.105	0.024	0.043	0.107
90	0.092	0.217	0.215	0.082	0.225	0.220	0.058	0.286	0.220
95	0.133	0.323	0.385	0.133	0.323	0.388	0.098	0.323	0.388
	Panel B. Subjects' Classification I								
	All Subjects			CCEI ≥ 0.9			CCEI ≥ 0.95		
	Ambiguity			Ambiguity			Ambiguity		
	Seeking	Neutral	Averse	Seeking	Neutral	Averse	Seeking	Neutral	Averse
Pessimism/optimism									
Optimistic	0.006	0.123	0.013	0.008	0.142	0.008	0.011	0.172	0.000
Neutral	0.006	0.643	0.065	0.008	0.622	0.055	0.011	0.602	0.043
Pessimistic	0.013	0.065	0.026	0.016	0.063	0.031	0.022	0.086	0.022
No. of obs.	154			127			93		

Of the 154 subjects, we cannot reject the joint hypothesis that $\hat{\delta}_n = 0$ and $\hat{\gamma}_n = 0$ for a total of 99 subjects (64.3 percent) at the 5 percent significance level. Of these subjects, 76 subjects (74.5 percent) are also well approximated by preferences consistent with SEU according to the estimation of the kinked specification.¹⁷ Of the 52 remaining subjects, 19 subjects (12.3 percent) only have significant degrees of optimism $\hat{\gamma}_n < 0$, and 10 subjects (6.5 percent) only have significant degrees of pessimism $\hat{\gamma}_n > 0$. Additionally, 10 subjects (6.5 percent) have significant degrees of ambiguity aversion $\hat{\delta}_n > 0$, and only 1 subject (0.6 percent) has a significant degree of ambiguity seeking $\hat{\delta}_n < 0$. Finally, only four subjects (2.6 percent) have significant degrees of ambiguity aversion and pessimism $\hat{\delta}_n > 0$ and $\hat{\gamma}_n > 0$, one subject (0.6 percent) has significant degrees of ambiguity seeking and optimism $\hat{\delta}_n < 0$ and $\hat{\gamma}_n < 0$, and four other subjects (2.6 percent) have significant degrees of ambiguity aversion and pessimism or vice versa. Six additional subjects (3.9 percent) cannot be classified by our statistical tests.

Finally, Figure 8 presents the data from Appendix X graphically in the form of scatter plots of the estimates $\hat{\delta}_n$ or $\hat{\gamma}_n$. The most notable features of the data in Figure 8 are the considerable heterogeneity in both $\hat{\delta}_n$ or $\hat{\gamma}_n$ and that their values are negatively cor-




FIGURE 8. Scatterplot of the estimated parameters $\hat{\delta}_n$ (*horizontal axis*) and $\hat{\gamma}_n$ (*vertical axis*) generalized kinked specification (equation (3)).

¹⁷The kinked specification in equation (1) is a special case of the generalized kinked specification in equation (3) when $\gamma = 0$. Since these specifications involve nonlinear demand functions, the distributional results of standard loss of fit tests based on the sums of squared residuals from the restricted and unrestricted regressions do not hold, but constructing test statistics with exact limit distributions goes beyond the scope of this paper. Using a t -test, of the 154 subjects, we cannot reject the hypothesis that $\hat{\gamma}_n = 0$ for a total of 93 subjects (60.4 percent) at the 5 percent significance level. The hypothesis that $\hat{\gamma}_n = 0$ cannot be rejected for a total of 109 subjects (70.8 percent) using the counterpart to the Wald-type statistic in the linear framework. Of those, 89 subjects (81.7 percent) are consistent with the kinked specification also according to the t -test. We note that the Wald test and the t -test are equivalent in large samples under standard assumptions.

related ($r^2 = -0.224$). Finally, we note that a significant majority of subjects have coefficients of absolute risk aversion $\hat{\rho}_n$ that are very similar to the coefficients of risk aversion from the kinked specification, and that there is a very strong correlation between the estimated $\hat{\rho}_n$ parameters that come from the kinked and generalized kinked specifications ($r^2 = 0.970$).

10. CONCLUSION

We have used a rich data set to estimate parametric models of attitudes toward risk and ambiguity, and pessimism/optimism. Our estimates of preference parameters confirm the heterogeneity of individual attitudes to uncertainty. Despite this heterogeneity, we cannot reject SEU preference for over 60 percent of subjects. While the remainder display either statistically significant ambiguity aversion or pessimism, only a very few subjects exhibit both ambiguity aversion and pessimism. We emphasize again that our population-level conclusions are based on individual-level statistical tests, which are not usually possible in the literature. We also reemphasize that there is little evidence that ambiguity-seeking preferences are present in the subject pool.

The main evidence for ambiguity aversion is the tendency for some subjects to equalize their demands for the securities that pay off in the ambiguous states. The kinked specification associated with MEU, CEU, α -MEU, and contraction expected utility can explain this tendency, whereas the smooth specification associated with the REU cannot. These findings should offer food for thought to theorists interested in developing empirically relevant models. There is also a strong tendency to equate the demand for the securities that pay off in the other pairs of states, which could be explained by pessimism. To distinguish these two effects, we estimated a generalized kinked specification and were able to classify subjects as being seeking, neutral, or averse to ambiguity, and pessimistic, neutral, or optimistic. This specification confirmed the heterogeneity of individual attitudes and also confirmed that the majority of subjects cannot be distinguished from SEU at a statistically significant level.

REFERENCES

- Abdellaoui, M., A. Baillon, L. Placido, and P. P. Wakker (2011), "The rich domain of uncertainty: Source functions and their experimental implementation." *American Economic Review*, 101 (2), 695–723. [200]
- Afriat, S. (1967), "The construction of a utility function from expenditure data." *Econometrica*, 6, 67–77. [207]
- Afriat, S. (1972), "Efficiency estimates of production functions." *International Economic Review*, 13, 568–598. [207]
- Afriat, S. (2012), "Afriat's theorem and the index number problem." *Economic Journal*, 122, 295–304. [208]
- Ahn, D. (2008), "Ambiguity without a state space." *Review of Economic Studies*, 75, 3–28. [196]

- Bade, S. (2013), “Randomization devices and the elicitation of ambiguity averse preferences.” Report. [201]
- Baillon, A., O. L’Haridon, and L. Placido (2011), “Ambiguity models and the Machina paradoxes.” *American Economic Review*, 101 (4), 1547–1560. [200]
- Bossaerts, P., P. Ghirardato, S. Guarnaschelli, and W. Zame (2010), “Ambiguity in asset markets: Theory and experiment.” *Review of Financial Studies*, 23, 1325–1359. [199]
- Bossaerts, P., C. Plott, and W. Zame (2007), “Prices and portfolio choices in financial markets: Theory, econometrics, experiments.” *Econometrica*, 75, 993–1038. [199]
- Camerer, C. (1995), “Individual decision making.” In *Handbook of Experimental Economics* (J. Kagel and A. Roth, eds.), Princeton University Press, Princeton. [199]
- Camerer, C. and M. Weber (1992), “Recent developments in modelling preferences: Uncertainty and ambiguity.” *Journal of Risk and Uncertainty*, 5, 325–370. [199]
- 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 and Proceedings*, 97, 153–158. [196, 200, 207, 208]
- Choi, S., R. Fisman, D. Gale, and S. Kariv (2007b), “Consistency and heterogeneity of individual behavior under uncertainty.” *American Economic Review*, 97, 1921–1938. [196, 200, 208, 213, 214]
- Diecidue, E. and P. Wakker (2001), “On the intuition of rank-dependent utility.” *Journal of Risk and Uncertainty*, 23, 281–298. [215]
- Diewert, W. (2012), “Afriat’s theorem and some extensions to choice under uncertainty.” *Economic Journal*, 122, 305–331. [208]
- Ellsberg, D. (1961), “Risk, ambiguity, and the Savage axioms.” *Quarterly Journal of Economics*, 75, 643–669. [196]
- Ergin, H. and F. Gul (2004), “A subjective theory of compound lotteries.” Report. [196]
- Gajdos, T., T. Hayashi, J.-M. Tallon, and J.-C. Vergnaud (2008), “Attitude toward imprecise information.” *Journal of Economic Theory*, 140, 27–65. [196]
- Ghirardato, P., F. Maccheroni, and M. Marinacci (2004), “Differentiating ambiguity and ambiguity attitude.” *Journal of Economic Theory*, 118, 133–173. [196]
- Gilboa, I. and D. Schmeidler (1989), “Maxmin expected utility with a non-unique prior.” *Journal of Mathematical Economics*, 18, 141–153. [196]
- Giraud, R. (2005), “Objective imprecise probabilistic information, second order beliefs and ambiguity aversion: An axiomatization.” In *ISIPTA’05: Proceedings of the 4th International Symposium on Imprecise Probabilities and Their Applications*, Pittsburgh, PA. [196]
- Gul, F. (2001), “A theory of disappointment aversion.” *Econometrica*, 59, 667–686. [214]

- Halevy, Y. (2007), "Ellsberg revisited: An experimental study." *Econometrica*, 75, 503–536. [199, 209, 213]
- Halevy, Y. and V. Feltkamp (2005), "A Bayesian approach to uncertainty aversion." *Review of Economic Studies*, 72, 449–466. [196]
- Houtman, M. and J. Maks (1985), "Determining all maximal data subsets consistent with revealed preference." *Kwantitatieve Methoden*, 19, 89–104. [208]
- Klibanoff, P., M. Marinacci, and S. Mukerji (2005), "A smooth model of decision making under ambiguity." *Econometrica*, 73, 1849–1892. [196]
- Machina, M. (2009), "Risk, ambiguity, and the rank-dependence axioms." *American Economic Review*, 99, 385–392. [200]
- Nau, R. (2006), "Uncertainty aversion with second-order utilities and probabilities." *Management Science*, 52, 136–145. [196]
- Olszewski, W. (2007), "Preferences over sets of lotteries." *Review of Economic Studies*, 74, 567–595. [196]
- Quiggin, J. (1982), "A theory of anticipated utility." *Journal of Economic Behavior and Organization*, 3, 323–343. [198, 215, 217]
- Samuelson, P. (1963), "Risk and uncertainty: A fallacy of large numbers." *Scientia*, 98, 108–113. [201]
- Savage, L. J. (1954), *The Foundations of Statistics*. Wiley, New York. [196]
- Schmeidler, D. (1989), "Subjective probability and expected utility without additivity." *Econometrica*, 57, 571–587. [196, 215]
- Segal, U. (1987), "The Ellsberg paradox and risk aversion: An anticipated utility approach." *International Economic Review*, 28, 175–202. [199]
- Segal, U. (1990), "Two-stage lotteries without the reduction axiom." *Econometrica*, 58, 349–377. [199]
- Seo, K. (2009), "Ambiguity and second-order belief." *Econometrica*, 77, 1575–1605. [196]
- Varian, H. (1982), "The nonparametric approach to demand analysis." *Econometrica*, 50, 945–973. [207]
- Varian, H. (1983), "Non-parametric tests of consumer behaviour." *Review of Economic Studies*, 50, 99–110. [207]
- Varian, H. (1990), "Goodness-of-fit in optimizing models." *Journal of Econometrics*, 46, 125–140. [208]
- Varian, H. (1991), "Goodness-of-fit for revealed preference tests." Report. [208]
- Varian, H. (2012), "Revealed preference and its applications." *Economic Journal*, 122, 332–338. [208]

Vermeulen, F. (2012), “Foundations of revealed preference: Introduction.” *Economic Journal*, 122, 287–294. [208]

Submitted January, 2012. Final version accepted September, 2013.