



The Ambiguity Triangle: Uncovering Fundamental Patterns of Behavior Under Uncertainty

Daniel R. Burghart
Thomas Epper
Ernst Fehr

CESIFO WORKING PAPER NO. 5420
CATEGORY 13: BEHAVIOURAL ECONOMICS
JUNE 2015

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: www.CESifo-group.org/wp

ISSN 2364-1428

The Ambiguity Triangle: Uncovering Fundamental Patterns of Behavior Under Uncertainty

Abstract

The probability triangle (also called the Marschak-Machina triangle) allows for compact and intuitive depictions of risk preferences. Here, we develop an analogous tool for choice under uncertainty – the ambiguity triangle – and show that indifference curves in this triangle capture preferences for unknown probabilities. In particular, the ambiguity triangle allows us to examine whether subjects adhere to the generalized axiom of revealed preference (GARP) and satisfy a non-parametric test for constant ambiguity attitudes. We find that more than 95% of subjects adhere to GARP and that about 60% satisfy our test for a constant ambiguity attitude. Yet, among these 60% of subjects there is substantial preference heterogeneity. We characterize this heterogeneity with finite-mixture estimates of a one-parameter extension of Expected Utility Theory wherein 48% of subjects are ambiguity averse, 22% are ambiguity seeking, and 30% are close to ambiguity neutral. The ambiguity triangle also highlights how variable ambiguity attitudes arise mainly because indifference curves are 'fanning-in' across the triangle. This fanning-in property implies that aversion to ambiguity increases as the likelihood of receiving a good outcome increases. We capture this behavior with a simple parametric model that also allows for finite mixture characterizations of preference heterogeneity for these subjects. We show that for a substantial share of these subjects (43%) their fanning-in is so strong that, although they are initially ambiguity seeking, they become strongly ambiguity averse as the likelihood of receiving a good outcome increases.

JEL-Code: D810, C910.

Keywords: uncertainty, ambiguity triangle.

*Daniel R. Burghart**
University of Zurich
Department of Economics
Switzerland – 8006 Zurich
dan.burghart@econ.uzh.ch

Thomas Epper
University of St. Gallen / School of
Economics and Political Science
Switzerland – 9000 St. Gallen
thomas.epper@unisg.ch

Ernst Fehr
University of Zurich
Department of Economics
Switzerland – 8006 Zurich
ernst.fehr@econ.uzh.ch

*corresponding author

June 22, 2015

All authors gratefully acknowledge valuable input from Jim Cox, Charles Efferson, Tony Williams, participants in the seminars at UZH, Virginia Tech, and Wesleyan University, and participants at the 2014 Economic Science Association North American meeting held in Fort Lauderdale, Florida. Special thanks go to Aurélien Baillon and Peter Wakker as they gave us detailed feedback on an early draft of this paper, in addition to some key references. Sonja Vogt was instrumental in preparing the urn used to resolve uncertainty. And we also thank the students, faculty and staff in the Department of Economics at UZH for their assistance in composing the urn. Stefan Wehrli, Oliver Brägger and Nadja Jehli provided exceptional technical and practical assistance during data collection. This work was supported by a grant from the European Research Council on the Foundations of Economic Preferences (#295642).

The analysis of choices under risk involves prospects that have fully-known probabilities. Real-world settings, however, rarely involve full knowledge about prospects. And when individuals do not possess full knowledge about prospects they are said to make choices in the face of uncertainty (also called “Knightian uncertainty,” after Knight, 1921). In the analysis of risky choices the probability triangle is a tool that provides valuable insights (Marschak, 1950; Machina, 1982). For example, Expected Utility implies linear and parallel indifference curves in the probability triangle with the slope of these indifference curves determined by an individual’s risk aversion. Furthermore, the famous Allais (1953) common-ratio and common-consequence paradoxes challenge that indifference curves in the probability triangle are parallel. As Machina (1982) points out, theories accommodating these paradoxes predict that risk aversion increases when shifting probability mass from worse to better outcomes. In this way, and others, the probability triangle provides succinct and intuitive representations of risk preferences (Sugden, 2004).

In the domain of uncertainty a similar tool is not available. In view of this shortcoming we develop such a tool – the ambiguity triangle – and demonstrate how it enables insights into choices under uncertainty. We will show, for example, that the ambiguity triangle permits an assessment of whether individuals exhibit constant ambiguity attitudes. Moreover, the ambiguity triangle facilitates this assessment being done non-parametrically. Furthermore, the ambiguity triangle permits an intuitive representation of varying ambiguity attitudes as linear indifference curves that have changing slopes across the triangle. In turn, this enables descriptive insights into the characteristics of preferences with varying ambiguity attitudes.

The ambiguity triangle is based on the concept of *lower envelope lotteries*.¹ Lower envelope lotteries specify lower bounds on probabilities, $\{\underline{p}_1, \dots, \underline{p}_K\}$, for a set of monetary outcomes, $\{z_1, \dots, z_K\}$, and the amount of unassigned probability mass, $y = 1 - \sum_k \underline{p}_k$, an objective quantity we call *ambiguity*. Formally, we denote a lower envelope lottery as $L = (\underline{p}_1, \dots, \underline{p}_K, y)$. As with traditional lotteries, the entries in a lower envelope lottery must be non-negative and sum to one.

For the case of two outcomes, $\{z_1, z_2\}$, with $z_1 > z_2$, the set of lower envelope lotteries can be represented graphically in a figure we call the *ambiguity triangle*. As depicted in Figure 1(a) the vertices of the ambiguity triangle represent three special cases: (i) receiving the good outcome (z_1)

¹The *lower envelope* nomenclature is borrowed from the literature on imprecise probabilities which describes partial knowledge about a probability distribution by the lower (and upper) envelopes of the set of candidate distributions. See Dempster (1967) and Shafer (1976).

with certainty (i.e. $\underline{p}_1 = 1$), (ii) receiving the bad outcome (z_2) with certainty (i.e. $\underline{p}_2 = 1$), and (iii) facing an alternative in which nothing is known about the probability distribution over the two outcomes (i.e. $y = 1$).


The ambiguity triangle restricts attention to lower envelope lotteries of the form $(\underline{p}_1, \underline{p}_2, y)$. And we examine choice-situations in which these types of lower envelope lotteries lie on linear “budget lines” in the ambiguity triangle. As depicted in Figure 1(a) the dots on a budget line indicate the lower envelope lotteries available in that choice-situation. Varying the slope and position of these budget lines allows us to examine properties of preferences in the ambiguity triangle. Of particular interest is the slope of indifference curves as this captures the local ambiguity attitude. Also of interest is how the slope of indifference curves changes, or doesn’t, in response to changes in the location of budget lines in the ambiguity triangle.

To illustrate how the slope of indifference curves captures ambiguity attitude consider the three lower envelope lotteries (A , B , and C) depicted in Figure 1(b). A horizontal movement from A to B involves an increase in ambiguity and an equal sized decrease in the minimum probability of the bad outcome. If an individual is indifferent between B and A he behaves as if he believes that the increase in ambiguity leaves the probability of the good outcome unchanged. This could reflect extremely pessimistic expectations about how uncertainty is being resolved. In contrast, if an individual strictly prefers B over A , then the horizontal movement from A to B is associated with a utility increase. This means that the individual behaves as if he believes that a part of the increase in ambiguity will increase the chances of receiving the good outcome. In order to keep utility constant to A (i.e. to stay on the same indifference curve as A) it would be necessary to offer a lower envelope lottery (C) that has a lower minimum probability for the good outcome than B . In turn this would imply that the indifference curve is negatively sloped. The higher the utility increase associated with the increase in ambiguity (i.e. when moving from A to B) the steeper the indifference curve. Therefore, a relatively flat indifference curve indicates ambiguity aversion. In contrast, a relatively steep indifference curve indicates ambiguity seeking.


We show in Section 1 that the lower envelope lotteries implied by the Ellsberg (1961) experiment can be represented in the ambiguity triangle. Moreover, the typical behavior in Ellsberg’s experiment can be rationalized in the ambiguity triangle by simply constraining the slope of linear and parallel indifference curves. Drawing motivation from Ellsberg’s setup, we designed and con-

Figure 1: The Ambiguity Triangle

(a) The Ambiguity Triangle Can Depict Lower Envelope Lotteries That Have Two Possible Outcomes $\{z_1, z_2\}$, where $z_1 > z_2$



(b) Illustrating Ambiguity Aversion



ducted an experiment that had 25 different budget lines in the ambiguity triangle. These 25 budget lines have varying slopes and locations in the triangle such that they form a web of intersecting budgets. This web of budgets enables us to examine whether choices satisfy the generalized axiom of revealed preference (GARP). We find that more than 95% of subjects obey GARP. This implies that choices can be modeled as if they resulted from maximizing a complete, transitive, continuous, and monotone preference relation (Varian, 1982).

A fundamental question in the domain of risk is whether risk aversion at 'high' utility levels differs from risk aversion at 'low' utility levels. An analogous question arises in the domain of ambiguity. We can address this question by examining how the slope of indifference curves in the ambiguity triangle change, or don't, as we vary the location of budgets. Because the slope of the indifference curves captures ambiguity attitudes, subjects with a constant attitude towards ambiguity should display indifference curves that have the same slope everywhere in the ambiguity triangle. That is, their indifference curves should be linear and parallel to each other. The web of budgets that enabled us to examine GARP-compliance can also be used to investigate, non-parametrically, this aspect of ambiguity attitudes. Our non-parametric test for constant ambiguity attitude reveals that roughly 60% of subjects satisfy this assumption while the other 40%, those who fail the test, must display variable ambiguity attitudes across the triangle. Our non-parametric test for constant ambiguity attitude gives us the confidence to apply a parametric model of constant ambiguity preferences to the 60% of subjects that pass the test. Finite mixture estimates of this parametric model endogenously partition this subject pool into ambiguity averse, neutral, and seeking individuals.

We call the parametric model we estimate Hurwicz Expected Utility (HEU). This model was axiomatized in the general 'sets of lotteries' setting by Olszewski (2007).² The HEU model can be interpreted as a combination of Expected Utility and the Hurwicz criterion (Hurwicz, 1951). The Hurwicz criterion stipulates that in the face of partial or complete uncertainty, behavior should maximize a weighted sum of the worst and the best possible outcomes (Luce and Raiffa (1958), §13.5). The setting of lower envelope lotteries makes an application of the Hurwicz criterion straightforward. Adopting the convention that outcomes are ordered from best to worst,

²Olszewski (2007) did not use the terminology 'Hurwicz Expected Utility' to describe his model. Although his notation for the utility function used to evaluate a set of lotteries, $H_{u,\alpha}$, may have been inspired by the widely-known Hurwicz criterion (see, for example, Arrow and Hurwicz, 1972).

$(z_1 > z_2 > \dots > z_K)$, and given a lower envelope lottery $L = (\underline{p}_1, \dots, \underline{p}_K, y)$, the worst outcome, from an ex-ante perspective, would be a lottery for which all ambiguity is ascribed to the probability of receiving the smallest monetary outcome: $\underline{L} = (\underline{p}_1, \dots, \underline{p}_K + y)$. The best ex-ante outcome would be a lottery for which all of the ambiguity is ascribed to the probability of receiving the largest monetary outcome: $\bar{L} = (\underline{p}_1 + y, \dots, \underline{p}_K)$. The Hurwicz criterion implies that for these worst and best lotteries, individuals construct a mixture between the two: $\alpha \underline{L} + (1 - \alpha) \bar{L}$. This mixture is assumed to be evaluated with Expected Utility (EU). The result is the Hurwicz Expected Utility model:³

$$HEU(L) = EU\left(\alpha \underline{L} + (1 - \alpha) \bar{L}\right). \quad (1)$$

HEU is a parsimonious model – it has just one more parameter than Expected Utility. The additional parameter, α , captures an everywhere-constant attitude towards ambiguity. For example, if $\alpha = 1$ an individual can be considered as completely pessimistic because he evaluates lower envelope lotteries by the worst possible lottery. This implies horizontal indifference curves in the ambiguity triangle. These individuals can be described as ‘ambiguity averse.’ When $\alpha = \frac{1}{2}$ an individual mixes equally between the best and worst possible lotteries and is said to be ‘ambiguity neutral.’ In this case indifference curves have a slope of negative 45-degrees in the ambiguity triangle.

To provide an empirical characterization of the heterogeneity of HEU preferences we employ a finite mixture approach. This approach provides an endogenous classification of individuals to different ambiguity types and characterizes the preferences for each type. Our estimation results show that there is substantial heterogeneity of ambiguity preferences in our experimental participants. This heterogeneity is best characterized by the existence of three distinct types that, coincidentally, correspond to ambiguity neutrality, aversion, and seeking behaviors. We also obtain estimated proportions of participants that belong to each of these distinct preference types. One type is comprised of 30% of the subjects and corresponds to (near) ambiguity neutrality ($\hat{\alpha} = 0.517$).

³ The Hurwicz Expected Utility model resembles, but is distinct from, the α -MaxMin Expected Utility model discussed by Marinacci (2002) and Ghirardato et al. (2004). The HEU model differs in two important ways. First, HEU preferences are defined over sets of lotteries whereas α -MaxMin preferences are defined over acts (bets). Second, in the HEU model the best and the worst ex-ante outcomes (\bar{L} and \underline{L}) are objectively defined entities. In contrast, in the α -MaxMin model the worst and the best probability distribution over outcomes are determined by an individual’s subjective beliefs. Section 4 discusses the implications of these differences for empirical work.

Roughly 48% of subjects exhibited ambiguity aversion ($\hat{\alpha} = 0.583$). The third type, comprising roughly 22% of subjects, display ambiguity seeking behavior ($\hat{\alpha} = 0.395$). Importantly, the finite-mixture approach does not assume the ex-ante existence of these three types. Instead, these three types emerge endogenously from our finite mixture estimates. In addition, almost all of the 120 subjects with constant ambiguity attitudes are unambiguously assigned to one distinct preference type or another (i.e. the subjects' posterior probabilities of belonging to a specific type is almost exclusively close to one).

But what about the other 83 subjects (roughly 40% of our sample) who exhibited varying ambiguity attitudes across the triangle? Our non-parametric test implies that their preferences cannot be characterized with the HEU model. Therefore a more flexible model is required for rationalizing their choices. But what type of flexibility is necessary to rationalize these participant's choices? Does the everywhere-constant ambiguity attitude embedded in the HEU model fail while other properties are retained (specifically, linearity along a single indifference curve)? Or are there broader inconsistencies with HEU such as non-linearities of indifference curves in the ambiguity triangle?

The ambiguity triangle and our set of linear budget lines enables us to examine this question in a non-parametric way. If nearly all choices are at the boundaries of the budget lines (i.e., almost no choices are in the interior of the budgets), it is unlikely that non-linear indifference curves are needed to rationalize choices. This follows from the observation that non-linear indifference curves can lead to interior choices. The choice data indicate that non-linear indifference curves do not play an important role in our setting. More specifically, we find that relaxing the everywhere-constant ambiguity attitude embedded in the HEU model but retaining the assumption of linear indifference curves can serve as a way to rationalize these 83 participant's choices. To parametrically characterize these varying ambiguity attitudes we introduce a model that, in an empirically expedient fashion, generalizes HEU. We call this the β -Hurwicz Expected Utility (β -HEU) model. The β -HEU model has one more parameter than HEU. It retains the linear indifference curves of HEU but permits indifference curves to 'fan-in' ($\beta > 0$) or 'fan-out' ($\beta < 0$) across the lower envelope lottery space. Using finite mixture estimation we show that under a β -HEU representation all 83 participants can be described as exhibiting increasing aversion to ambiguity aversion as the minimum probability of the good outcome (i.e. $\beta > 0$). But again there is substantial heterogeneity

in this group. About 43% of these subjects display a very large variation in their ambiguity preferences ($\hat{\beta} = 0.261$). In fact, these individuals are strongly ambiguity seeking when the minimum probability of z_1 is low but then strongly ambiguity averse when the minimum probability of z_1 is high. For the remaining subjects (i.e. 57% of the β -HEU representable individuals) ambiguity aversion also increases significantly, but not as strongly ($\hat{\beta} = 0.093$).


The remainder of the paper is organized as follows. The next section shows that the Ellsberg paradox can be neatly captured by appropriately sloped linear indifference curves in the ambiguity triangle. Section 2 introduces our experimental design and methods in more details. In Section 3 we present our empirical results and Section 4 discusses the relationship between (i) the HEU model and other theoretical approaches towards ambiguity and (ii) our empirical approach and other empirical studies of ambiguity attitudes. This section also discusses recent empirical studies of ambiguity including Abdellaoui et al. (2011), Ahn et al. (2014), Halevy (2007), Hey et al. (2010), Hey and Pace (2014) and Stahl (2014). Section 5 concludes and provides directions for future work.

1 The Ellsberg Experiment

In this section we show how an experiment proposed by Ellsberg (1961) can be described with lower envelope lotteries. This description allows for an intuitive graphical representation of Ellsberg's experiment in the ambiguity triangle. This description also shows that the so-called Ellsberg paradox can be explained if subjects have linear and parallel indifference curves (that exhibit ambiguity aversion) in the triangle.

Ellsberg (1961) argued that individuals are likely to violate a key postulate of Subjective Expected Utility Theory – the Sure-Thing Principle (postulate P2 in Savage, 1954). The Sure-Thing Principle requires that preferences between acts must be maintained if the state(s) for which the acts pay off are both augmented with the same state. For example, suppose an urn is composed of red, black, and yellow balls in unknown numbers. Denote as $\{R\}$ (and $\{B\}$) the act that pays \$100 if a red (black) ball is drawn from the urn and \$0 otherwise. Supposing that an individual prefers the bet on red to the bet on black, so that $\{R\} \succ \{B\}$, the Sure-Thing Principle requires that the individual must also have $\{R \cup Y\} \succ \{B \cup Y\}$ (where $\{R \cup Y\}$ denotes the bet that pays \$100 if either a red or yellow ball is drawn from the urn, \$0 otherwise; $\{B \cup Y\}$ denotes the bet that pays

Figure 2: The Ellsberg (1961) Three-Color Urn



\$100 if either a black or yellow ball is drawn from the urn, \$0 otherwise).

To illustrate his case, Ellsberg proposed constructing an urn that contained 90 balls. As depicted in Figure 2, exactly 30 of the balls in Ellsberg’s urn were red. Each of the remaining 60 balls were either black or yellow so that the proportion of black to yellow balls was unknown. Ellsberg’s thought experiment had two choice situations. The first choice-situation had participants selecting between a bet on red, $\{R\}$, and a bet on black, $\{B\}$, which, as above, paid \$100 if the color of the ball drawn from the urn matched the color of the bet, and \$0 otherwise. The second choice-situation had participants selecting between a bet on red or yellow $\{R \cup Y\}$, and a bet on black or yellow, $\{B \cup Y\}$. When surveyed, most people preferred the bet on red in the first choice situation (i.e. $\{R\} \succ \{B\}$) and the bet on black or yellow in the second choice situation (i.e. $\{B \cup Y\} \succ \{R \cup Y\}$).⁴ This pattern of choices clearly violates the Sure-Thing Principle. This violation, Ellsberg argued, was motivated by an aversion to unknown probabilities. Specifically, for Ellsberg’s urn, the bets on black, $\{B\}$, and red or yellow, $\{R \cup Y\}$, ascribe interval probabilities, or sets of probability distributions, to the two outcomes, $\{\$100, \$0\}$. In contrast, the preferred bets, those on red, $\{R\}$, and black or yellow, $\{B \cup Y\}$, each imply an exact probability distribution over the outcomes: $(\frac{1}{3}, \frac{2}{3})$ and $(\frac{2}{3}, \frac{1}{3})$, respectively. The term ‘ambiguity aversion’ has been used to describe this aversion toward sets of probabilities.

⁴Slovic and Tversky (1974) collected, but did not report, choices from an Ellsberg three-color urn experiment. The experimental data reported by Dobbs (1991) has the majority of participants choosing the canonical Ellsberg bets (for the bets over gains). Becker and Brownson (1964) and Einhorn and Hogarth (1986) both reported choices consistent with the pattern of behavior hypothesized by Ellsberg (1961) in his two-color, two-urn problem.

1.1 The Ellsberg Experiment as Lower Envelope Lotteries

All of the bets in the Ellsberg experiment can be described with lower envelope lotteries. Adopt the convention $(\underline{p}_1, \underline{p}_2, y)$, where \underline{p}_1 is the minimum probability of receiving \$100, \underline{p}_2 is the minimum probability of receiving \$0, and y is the amount of ambiguity or unassigned probability.

Consider first the bet on red. Exactly 30 of the 90 balls in the urn are red so the probability of receiving \$100 is exactly $\frac{1}{3}$. Exactly 60 of the 90 balls are not red so the probability of receiving \$0 is exactly $\frac{2}{3}$. Denoting this as a lower envelope lottery: $(\frac{1}{3}, \frac{2}{3}, 0) = R$.

Next, consider the bet on black. Of the 90 balls in the urn there could be any number between 0 and 60 that are black. So the probability of receiving \$100 could be anywhere between 0 and $\frac{2}{3}$. This means that the minimum probability of receiving \$100 is 0. That the maximum probability of receiving \$100 is $\frac{2}{3}$ means that the minimum probability of receiving \$0 is $\frac{1}{3}$. Denoting this as a lower envelope lottery: $(0, \frac{1}{3}, \frac{2}{3}) = B$.

The third bet in the Ellsberg experiment is for red or yellow. Because there are exactly 30 red balls in the urn and any number between 0 and 60 yellow balls in the urn, the minimum probability of receiving \$100 is $\frac{1}{3}$ while the maximum is 1. Because the maximum probability of receiving \$100 is 1, the minimum probability of receiving \$0 is 0. Denoting this as a lower envelope lottery: $(\frac{1}{3}, 0, \frac{2}{3}) = RY$.

The fourth bet in Ellsberg's experiment is for black or yellow. There are exactly 60 black or yellow balls in the urn so the probability of receiving \$100 is exactly $\frac{2}{3}$. And exactly 30 of the 90 balls are red so the probability of receiving \$0 is exactly $\frac{1}{3}$. Denoting this as a lower envelope lottery: $(\frac{2}{3}, \frac{1}{3}, 0) = BY$. The right-most column of Table 1 shows the lower envelope lotteries implied by each of the four bets in Ellsberg's experiment.

1.2 A Graphical Illustration of the Ellsberg Experiment in the Ambiguity Triangle

The lower envelope lotteries that describe the Ellsberg experiment can be graphically depicted in the ambiguity triangle. As depicted in Figure 3(a), the lower-left vertex of the ambiguity triangle represents the lower envelope lottery $(0, 1, 0)$ – certainty of receiving \$0. The upper-right vertex represents $(1, 0, 0)$ – certainty of receiving \$100. The lower-right vertex represents $(0, 0, 1)$ – probabilities are completely unknown. Movements from this vertex to the left involve a decrease of

Table 1: The Ellsberg Experiment Described as Lower Envelope Lotteries

Bet on...	Payoff (\$) if the Drawn-Ball is...			Lower Envelope Lotteries
	Red	Black	Yellow	$(\underline{p}_1, \underline{p}_2, y)$
*Red – $\{R\}$	100	0	0	$(\frac{1}{3}, \frac{2}{3}, 0) = R$
Black – $\{B\}$	0	100	0	$(0, \frac{1}{3}, \frac{2}{3}) = B$
Red or Yellow – $\{R \cup Y\}$	100	0	100	$(\frac{1}{3}, 0, \frac{2}{3}) = RY$
*Black or Yellow – $\{B \cup Y\}$	0	100	100	$(\frac{2}{3}, \frac{1}{3}, 0) = BY$

* indicates the modal choice observed in this type of experiment.

ambiguity and an increase in the minimum probability assigned to the \$0 outcome, (\underline{p}_2) – along the horizontal leg of the ambiguity triangle no minimum probability for the \$100 outcome is specified so that $\underline{p}_2 + y = 1$. Movements upward from the lower-right vertex involve a decrease of ambiguity and an increase in the minimum probability assigned to the \$100 outcome, (\underline{p}_1) – along the vertical leg of the triangle no minimum probability for the \$0 outcome is specified so that $\underline{p}_1 + y = 1$.

Two of the lower envelope lotteries implied by bets in Ellsberg’s experiment are fully-specified lotteries: R and BY . As can be seen in Figure 3(b), both of these lotteries lie on the hypotenuse of the ambiguity triangle. The other two lower envelope lotteries implied by the bets in Ellsberg’s experiment, B and RY , lie on the horizontal-leg and vertical-leg of the triangle, respectively. The parallel, dashed lines connecting R to B , and BY to RY , indicate the pairs of lower envelope lotteries that are implied by the two choice-situations in Ellsberg’s experiment.


1.3 Rationalizing Choices in Ellsberg’s Experiment with Hurwicz Expected Utility

The canonical pattern of choices in Ellsberg’s experiment have participants betting on red in the first choice-situation and betting on black or yellow in the second choice-situation. Adopting the lower envelope lottery notation: $R \succ B$ and $BY \succ RY$. Figure 3(b) indicates these choices with blue circles.


Many models have been proposed that can rationalize the canonical pattern of choices in Ellsberg’s experiment. Here, we utilize a parsimonious model that can be considered an amalgam of

Figure 3: Ambiguity Triangles Illustrating Ellsberg's Experiment and how the Hurwicz Expected Utility (HEU) Model can Rationalize Canonical Choices Therein


(a) The ambiguity triangle can be used to graphically depict two-outcome lower envelope lotteries




(b) Ellsberg's bets depicted in the ambiguity triangle. The blue circles show the canonical Ellsberg choices ($R \succ B$ and $BY \succ RY$)



(c) Three *distinct* HEU preferences depicting ambiguity **aversion** ($\alpha = \frac{3}{5}$), **neutrality** ($\alpha = \frac{1}{2}$), and **seeking** ($\alpha = \frac{2}{5}$)



(d) Ambiguity **averse** HEU preferences can rationalize canonical Ellsberg choices



two normative standards for behavior under uncertainty: Expected Utility (EU) and the Hurwicz criterion. This so-called Hurwicz Expected Utility (HEU) model evaluates lower envelope lotteries (L) as the Expected Utility, denoted $EU(\cdot)$, of a mixture between the objectively worst (\underline{L}) and best (\bar{L}) lotteries that can arise from a lower envelope lottery:

$$\text{HEU}(L) = \text{EU}(\alpha \underline{L} + (1 - \alpha) \bar{L}). \quad (2)$$

Adopting the notation that the utility for an outcome is u_k , and because we assume that outcomes are ordered from largest to smallest, we have $u_1 > u_2 > \dots > u_K$. The linearity of Expected Utility permits a rearrangement of Equation 2 into an α -MaxMin Expected Utility-like expression:

$$\begin{aligned} \text{HEU}(L = (\underline{p}_1, \dots, \underline{p}_K, y)) &= \alpha \text{EU}(\underline{L}) + (1 - \alpha) \text{EU}(\bar{L}) \\ &= \alpha \left[(\underline{p}_K + y) u_K + \sum_{k=1}^{K-1} \underline{p}_k u_k \right] + (1 - \alpha) \left[(\underline{p}_1 + y) u_1 + \sum_{k=2}^K \underline{p}_k u_k \right] \end{aligned} \quad (3)$$

The expression in the left-hand brackets of Equation 3 captures total pessimism: all ambiguity (y) is added to the minimum probability of receiving the worst outcome (\underline{p}_K). The resulting lottery, $\underline{L} = (\underline{p}_1, \underline{p}_2, \dots, \underline{p}_K + y)$, is then evaluated with Expected Utility. The expression in the right-hand brackets captures total optimism: all ambiguity (y) is added to the minimum probability of receiving the best outcome (\underline{p}_1). The resulting lottery, $\bar{L} = (\underline{p}_1 + y, \underline{p}_2, \dots, \underline{p}_K)$, is also evaluated with Expected Utility. The preference parameter, $\alpha \in [0, 1]$, thus represents a weighting between total pessimism and total optimism. For the extreme case of $\alpha = 0$ a decision maker is perfectly optimistic. In such cases we will say this decision maker is ambiguity seeking. At the other extreme, when $\alpha = 1$, a decision maker is perfectly pessimistic, and we will say they are ambiguity averse. When $\alpha = \frac{1}{2}$, a decision maker is equally balanced between pessimism and optimism. In such cases we will say that the decision maker is ambiguity neutral. Also note that when a lower envelope lottery is a fully specified lottery (i.e. $y = 0$), Hurwicz Expected Utility reduces to Expected Utility.

The Hurwicz Expected Utility for the Ellsberg two-outcome setting can be written as:

$$\text{HEU}(\underline{p}_1, \underline{p}_2, y) = \alpha \left[\underline{p}_1 u_1 + (\underline{p}_2 + y) u_2 \right] + (1 - \alpha) \left[(\underline{p}_1 + y) u_1 + \underline{p}_2 u_2 \right]. \quad (4)$$

In the ambiguity triangle the Hurwicz Expected Utility model produces indifference curves that are parallel straight-lines. The slope of these indifference curves is determined by the ambiguity preference parameter $\alpha \in [0, 1]$. The following definition elucidates the relationship between α and commonly used descriptions of behavior.

Definition 1. A Hurwicz Expected Utility type is ambiguity *averse*, *neutral*, or *seeking* whenever:

$$\alpha > \frac{1}{2}, \quad \alpha = \frac{1}{2}, \quad \alpha < \frac{1}{2} \quad (5)$$

The ambiguity triangle can be used to provide a geometric interpretation for these Hurwicz Expected Utility preference types. To facilitate this interpretation consider a choice between 100% ambiguity $(0, 0, 1)$ and a 50/50 lottery $(\frac{1}{2}, \frac{1}{2}, 0)$. Figure 3(c) plots these two alternatives in the ambiguity triangle. The 100% ambiguous alternative is at the lower-right vertex of the triangle and the 50/50 lottery is halfway between the endpoints of the hypotenuse. Because there is a one-to-one tradeoff between the ambiguity and the probabilities attached to each outcome an ambiguity *averse* decision maker would strictly prefer the 50/50 lottery, so that

$$\text{HEU}\left(\left(\frac{1}{2}, \frac{1}{2}, 0\right)\right) = \frac{1}{2}u_1 + \frac{1}{2}u_2 > \text{HEU}\left((0, 0, 1)\right) = (1 - \alpha)u_1 + \alpha u_2 \Leftrightarrow \alpha > \frac{1}{2}. \quad (6)$$

An ambiguity *neutral* decision maker would be indifferent, such that

$$\frac{1}{2}u_1 + \frac{1}{2}u_2 = (1 - \alpha)u_1 + \alpha u_2 \Leftrightarrow \alpha = \frac{1}{2}. \quad (7)$$

And an ambiguity *seeking* decision maker would strictly prefer the ambiguous option so that

$$\frac{1}{2}u_1 + \frac{1}{2}u_2 < (1 - \alpha)u_1 + \alpha u_2 \Leftrightarrow \alpha < \frac{1}{2}. \quad (8)$$

Figure 3(c) depicts three indifference curves, one for each of these three distinct types of preferences. The pink indifference curve depicts an ambiguity *seeking* preference (specifically, $\alpha = \frac{2}{5}$). The green indifference curve depicts an ambiguity *neutral* preference ($\alpha = \frac{1}{2}$). And the blue indifference curve depicts an ambiguity *averse* preference ($\alpha = \frac{3}{5}$). Figure 3(d) illustrates how these ambiguity *averse* preferences can rationalize the canonical pattern of choices in the Ellsberg ex-

periment. Similar to how Expected Utility produces parallel and straight-line indifference curves in the probability triangle, HEU produces indifference curves that are parallel straight-lines in the ambiguity triangle. So, an ambiguity **averse** indifference curve that has $R \succ B$ must also have $BY \succ RY$, which is the canonical pattern of choices in Ellsberg’s experiment.

2 Experimental Design And Methods


To study preferences towards objective ambiguity we expanded upon the (implied) design of Ellsberg’s one-urn experiment in lower envelope lottery space. Similar to the Ellsberg experiment, we used only two monetary outcomes: 60 or 20 Swiss Francs (CHF).⁵ In contrast to the Ellsberg experiment, our design had many more choice-situations and each choice-situation had several alternatives. Broadly speaking, our experimental design was a traditional demand elicitation mechanism for minimum probabilities and ambiguity. We systematically varied (i) the relative tradeoff between minimum probabilities and ambiguity, a value we can interpret as a price (i.e. the slope of a linear budget), and (ii) the probability assigned to the CHF 60 outcome for the fully-specified lottery available in a choice-situation, a value we call the *likelihood* ascribed to a specific choice-situation.

To illustrate how our design can be viewed as traditional demand elicitation, consider the three example budgets (line segments) in Figure 4(a). The thin, solid line segment represents a choice situation in which the tradeoff between ambiguity and the minimum probabilities of receiving each outcome is one-to-one. Specifically, if ambiguity is reduced by two units this results in a one unit increase in each of the minimum probabilities. This is easy to see when moving from the end of the budget at the lower-right vertex (100% ambiguity), to the end of the budget at the midpoint on the hypotenuse (the 50/50 lottery). This movement completely reduces the 100% ambiguity into an equal division of probabilities. Next, consider the thick line segment (budget) in Figure 4(a) that connects the point (0.4, 0, 0.6) on the vertical leg of the triangle to the 70/30 lottery on the hypotenuse. This line retains the tradeoff between ambiguity and minimum probabilities, but increases the probability for the CHF 60 outcome in the fully-specified lottery available. We call this type of movement a ‘tradeoff-constant increase in likelihood,’ or just ‘likelihood increase.’


⁵At the time of the experiment, one CHF was worth approximately \$1.10.

Figure 4: The Design of our Experiment

(a) Three budgets illustrating (i) a tradeoff-constant increase in likelihood (thin line to thick line) and (ii) a likelihood-constant tradeoff change (thick line to dashed line)



(b) The 25 budgets used in our experiment, each featuring six alternatives



Finally, consider the dashed line in Figure 4(a). The dashed line and thick line have the same probability for the CHF 60 outcome in the fully-specified lottery available from those budgets (i.e. both budgets intersect the hypotenuse of the ambiguity triangle at the same point). The dashed line, however, changes the tradeoff between minimum probabilities and ambiguity such that a two unit reduction in ambiguity results in a more-than-one unit increase in the minimum probability of CHF 60, and a smaller-than-one-unit increase in the minimum probability of CHF 20. Put simply, for the dashed line segment in Figure 4(a), the tradeoffs are such that the minimum probability for CHF 60 is cheap.

The budgets used in our experiment are plotted in Figure 4(b). The black lines depict the budgets. The dots on each line depict the alternatives available on that budget. This set of budgets provide a wide range of variation that permits a close examination of preferences towards ambiguity. Table 4 in Appendix C provides a detailed summary of these budgets. Also in Appendix C, Table 5 provides a comprehensive list of the lower envelope lotteries available on each budget.

2.1 Elicitation Software


Participant choices were elicited using software programmed with the Psychtoolbox Matlab libraries freely available at <http://www.psychtoolbox.org>. And Figures 5(a - d) show screenshots from our elicitation software for one choice situation (budget). The row of boxes at the bottom of each screenshot, each with a thumbnail image, provide a visual depiction of all the lower envelope lotteries available in that choice situation. When participants moved their mouse over a box, the corresponding thumbnail image was drawn as a large image in the center of the screen. The large image provided detailed outcome, minimum probability, and ambiguity information for the alternative depicted by the highlighted thumbnail.

Figures 5(a - d) depict alternatives available from one choice situation (budget) in the experiment. More specifically, Figures 5(a - d) depict alternatives available from the budget depicted by the thin solid line segment in Figure 4(a). In the row of boxes at the bottom of each screenshot in Figures 5(a - d) the left-most alternative is the 50/50 lottery (i.e. the lower envelope lottery $(\frac{1}{2}, \frac{1}{2}, 0)$ which is represented as the midpoint on the hypotenuse in the ambiguity triangle). The right-most alternative is the 100% ambiguous alternative (i.e. the lower envelope lottery $(0, 0, 1)$ which is represented as the lower-right vertex in the ambiguity triangle). Figure 5(a) illustrates when the mouse was moved over the thumbnail depicting the 100% ambiguous alternative – a green outline surrounds the right-most thumbnail and the large image shows the outcome information for this alternative and that no minimum probability information is provided. Figure 5(b) depicts when the mouse highlighted an alternative from the interior of the budget – a green outline surrounds the thumbnail and the specific minimum probability and ambiguity information are shown in the large image. Figure 5(c) depicts when the mouse was used to highlight the fully-specified lottery available from this budget – a green outline surrounds the left-most thumbnail and the probability information is shown in the large image.


Participants were instructed to select the alternative they wanted most in each choice-situation. Participants indicated their choice by moving the mouse over the relevant thumbnail, such that it was depicted as the large image in the middle of the screen, and then clicking the thumbnail. After clicking the thumbnail the green outline was locked around that thumbnail and “Confirm” and “Cancel” buttons appeared (Figure 5(d) shows these buttons when the 50/50 lottery was selected).

Figure 5: Screenshots from our Elicitation Software


(a) The Lower Envelope Lottery $(0, 0, 1)$




(b) The Lower Envelope Lottery $(\frac{2}{10}, \frac{2}{10}, \frac{6}{10})$



(c) The Lower Envelope Lottery $(\frac{1}{2}, \frac{1}{2}, 0)$
(i.e. the 50/50 Lottery)




(d) Confirm and Cancel Buttons Appeared After a Thumbnail was Clicked




Panels (a) - (c) depict different alternatives from the same budget that corresponds to the thin solid line-segment in Figure 4(a) (i.e. the budget with endpoints $(0, 0, 1)$ and $(\frac{1}{2}, \frac{1}{2}, 0)$). Panel (d) depicts “Confirm” and “Cancel” buttons that appeared after participants clicked one of the thumbnails at the bottom of the screen.

Figure 6: A Zoomed in View of the Rows of Thumbnails Depicting the Alternatives Available from the Three Line Segments (Budgets) in Figure 4(a)


(a) Thumbnails representing the alternatives available from the thin line segment in Figure 4(a)



(b) A Tradeoff Constant Likelihood Increase – Thumbnails representing the alternatives available from the thick line segment in Figure 4(a)



(c) A Likelihood Constant Tradeoff Change – Thumbnails representing the alternatives available from the dashed line segment in Figure 4(a)



If participants clicked the “Cancel” button, the “Confirm” and “Cancel” buttons vanished, and the green outline was unlocked so as to allow a different alternative to be selected in that choice situation. If participants clicked the “Confirm” button the software recorded the participant’s choice and moved on to the next choice-situation.

The row of thumbnail images at the bottom of the screen was designed to minimize participant confusion. Specifically, the row of thumbnails provide a straightforward and easy-to-see representation of the tradeoff and likelihood in a given budget. To see how this works consider Figures 6(a - c). Figure 6(a) is a closeup of the thumbnail images in the screen shots in Figure 5(a - d). It shows a one-to-one tradeoff between ambiguity and minimum probabilities. Figure 6(b) illustrates how the thumbnails made it easy to see a tradeoff constant likelihood shift – these thumbnails correspond to alternatives available from the thick line segment in Figure 4(a). Figure 6(c) illustrates how the thumbnails made it easy to see a likelihood constant tradeoff shift – these thumbnails correspond to alternatives available from the dashed line segment in Figure 4(a).

2.2 Resolution of Uncertainty

The choice-situations in our experiment made available many distinct lower envelope lotteries. To be able to resolve the uncertainty for each of these lower envelope lotteries, with a single source, and in one stage, we constructed an urn with 100 balls in it. To illustrate how this worked consider the lower envelope lottery $(\frac{2}{10}, \frac{3}{10}, \frac{5}{10})$ depicted in Figure 7. Relevant for this case the urn contained 30 balls numbered 1 to 30 and 20 balls numbered 81 to 100. The remaining 50 balls were in the interval 31 to 80 but it was unknown how many of them had an odd-number on them and how many had an even-number on them. If the ball drawn from the urn had any number 1 to 30 on it, the lottery paid out CHF 20. If the ball drawn from the urn had any number 81 to 100 on it, the lottery paid out CHF 60. If the ball was an odd-number between 31 and 80, the lottery paid out CHF 20. And if the ball was an even-number between 31 and 80 the lottery paid out CHF 60. Thus, the key was that odd-numbered and even-numbered balls in the ambiguity range were used for the resolution of ambiguity. This method worked because subjects did not know how many odd- and even-numbered balls were actually in the range of ambiguity.

Figure 8 illustrates how we constructed the urn for the ambiguity range. The figure assumes, for simplicity, that the range of ambiguity is given by the numbers from 1 to 10. This range was subdivided into pairs of subsequent numbers. For any pair of subsequent numbers it was possible that the two numbers were both odd-numbered (e.g., 7-7) or both even-numbered (e.g., 8-8) or a mix of an odd and an even number (e.g., 7-8). Thus, in the extreme, it was possible that all balls in the interval 1 to 10 were even-numbered or, that all of them were odd-numbered. Of course any other combination of odd- and even-numbered balls was also possible. In this way our urn allowed us to resolve the uncertainty in one stage for any of the lower envelope lotteries in our experiment.⁶

⁶Following Hey et al. (2010), and the re-assessment of their data by Kothiyal et al. (2014), we wanted to insure that participants were not concerned about experimenters 'stacking the urn' against the participants. Thus, experimenters were blind to the exact composition of the urn. This was accomplished by giving various students, faculty, and staff in the Department of Economics at the University of Zürich envelopes with two balls, a pen, and instructions telling them to, for example, write 7 on both balls, write 8 on both balls, or 7 on one ball and 8 on the other ball, and then to seal the envelope. The balls were placed in the urn, and quality control for readability and consistency with the instructions, was done by a person not involved with the experiment. Experimental participants were told how the urn was constructed during the instructions, and informed that experimenters were blinded to the exact composition of the urn.

Figure 7: Illustrating How the 100-ball Urn Was Used to Resolve Uncertainty For the Lower Envelope Lottery $(\frac{2}{10}, \frac{3}{10}, \frac{5}{10})$





Figure 8: The 100-Ball Urn Used In Our Experiment




2.3 Making Draws From The Urn

To make draws from our urn we used a commercially available bingo blower. As can be seen in Figure 9, we occluded the windows on our bingo blower with purple masking tape so that balls could be seen mixing inside the blower but the numbers on the balls could not be seen. Before participants entered the lab they were told that a bingo blower would be used as a randomization device to determine their payment in the experiment. To familiarize participants with the bingo blower it was turned on and placed in the corridor where participants entered the laboratory.

The bingo-blower we used can create a queue of balls in the drawing tube. To avoid ball-queuing in the drawing tube, we inserted a stopper. As can be seen in Figure 9, the stopper was a wooden dowel, cut to the same length as the drawing tube, with a large wooden ball attached to the top. When participants were at the cashier's desk, they removed the stopper so that one ball was "drawn" by the bingo blower. After the participant left the cashier's desk, the drawn-ball and stopper were replaced in the bingo blower by the experimenter.

Figure 9: The Bingo Blower and Stopper Used to Make Draws from the Urn



3 Analyzing Choices

In this section we analyze choices made by 203 participants in experimental sessions conducted at the Decision Science Laboratory, located in the ETH Zürich. All experimental sessions were conducted during one week in July, 2014. All experimental procedures were consistent with a protocol approved by the ethics committee at the University of Zürich.

Our analysis proceeds in three steps. The first step assesses whether participant's choices adhere to the generalized axiom of revealed preference (GARP). Choices that adhere to GARP can be rationalized with a preference relation that is complete, transitive, continuous, and monotone, foundational assumptions which underlie the majority of theories of choice under uncertainty. Because our experiment was designed as a set of intersecting linear budgets, it has excellent power for detecting violations of GARP. In general, we find very high rates of GARP compliance.

Our next analytical step uses a non-parametric test to assess whether individual choices can be rationalized with a Hurwicz Expected Utility representation. Using this non-parametric test, we identify a large subset of participants (120 out of 203) whose choices can be represented with the parsimonious HEU representation. We then take this subset of HEU-representable participants and calibrate ambiguity preferences using parametric estimation techniques. We start with just one type of preferences (i.e. a homogeneous preference model) and, one-by-one, increase the number of preference types. Three qualitatively distinct preference types emerge from our analysis. These three types correspond to ambiguity seeking, ambiguity aversion, and, roughly speaking, ambiguity neutrality.

Our two-step analytical procedure, that combines non-parametric and parametric assessments, is distinct from the purely parametric approach used by Stahl (2014) and Hey and Pace (2014). For example, Hey and Pace (2014) used experimental data in an assessment of five single-stage models of decision-making under ambiguity in which beliefs were taken to be endogenous. Their experimental design had participants betting on which color of ball (with three possible colors) would be drawn from an urn that was mixed by a bingo blower. Participants were shown the actively mixing bingo blower that contained the urn in a transparent plexiglass box.⁷ Hey and

⁷The transparency of the bingo blower is an important consideration when interpreting the results in Hey and Pace (2014). Transparency means that some information regarding the composition of the urn would be available to participants through visual inspection. For example, the extreme possibility of the urn being comprised of just one color of ball could be easily ruled out.

Pace (2014) then used parametric estimation techniques to describe and predict choices. They found support for α -MaxMin Expected Utility (Marinacci, 2002) as a representation for behavior in their experimental data, in addition to Vector Expected Utility (Siniscalchi, 2009).

The final step in our analysis examines choices made by the 83 participants which, arguably, could not be rationalized with a Hurwicz Expected Utility representation. We closely examine the choice data to understand why HEU failed as a parsimonious model. We propose an extension, β -Hurwicz Expected Utility, that provides a good fit for this sub-sample of participants. The β -HEU model retains the linear indifference curves of HEU but, its additional parameter, β , allows for fanning-in or fanning-out across the ambiguity triangle.

3.1 Assessing the Generalized Axiom of Revealed Preference (GARP)


Choices from a set of linear budgets (with strictly positive prices) are consistent with the maximization of a preference relation that is complete, transitive, continuous, and monotone if and only if they adhere to the generalized axiom of revealed preference (Afriat, 1967; Varian, 1982). In an experimental setting such as ours, the generalized axiom is an easy-to-test condition and Varian (1982) provides an easy-to-implement algorithm. Technical details are in Appendix A.1.

Adherence to an axiomatic criterion is binary: Either choices conform perfectly to GARP or they do not. While taking a binary perspective on GARP-compliance can be informative, a widely accepted approach is to determine, conditional on observing a violation, just “how badly” choices departed from perfect GARP-compliance. A common measure used in this regard is Houtman-Maks (HM). Houtman-Maks is the largest subset of choices which are GARP-compliant (Houtman and Maks, 1985). So, for example, if we see that a full set of 25 choices violate GARP but, by removing one offending choice, the remaining choices are GARP-compliant, this would represent a Houtman-Maks of 24.

To determine whether a participant’s choices are GARP-compliant, we compare their Houtman-Maks to a critical value derived from a Monte Carlo simulation. Our simulation uses choices for 5,000 synthetic experimental participants, each with a uniform random choice rule. We calculated Houtman-Maks for each synthetic. We interpret the distribution of all 5,000 synthetics’ Houtman-Maks as a sampling distribution and use it to determine critical values. This sampling distribution

Figure 10: Houtman-Maks for the Generalized Axiom of Revealed Preference (GARP)

(a) A High Rate of GARP-compliance Illustrated by the Distribution of Houtman-maks for Experimental Participants (red) and a Sampling Distribution Implied by 5,000 Synthetics With a Uniform Random Choice Rule (Blue)



(b) Supporting Data

Houtman-Maks	Humans			Synthetics		
	Freq.	%	Cum.	Freq.	%	Cum.
12	0	0	0	2	0.04	0.04
13	0	0	0	10	0.2	0.24
14	0	0	0	86	1.72	1.96
15	0	0	0	309	6.18	8.14
16	0	0	0	728	14.56	22.7
17	0	0	0	1,093	21.86	44.56
18	0	0	0	1,158	23.16	67.72
19	1	0.49	0.49	885	17.7	85.42
20	4	1.97	2.46	489	9.78	95.2
21	11 [†]	5.42	7.88	188	3.76	98.96
22	21 [†]	10.34	18.23	43	0.86	99.82
23	39 [†]	19.21	37.44	7	0.14	99.96
24	45 [†]	22.17	59.61	1	0.02	99.98
25	82 [†]	40.39	100	1	0.02	100
Total	203	100	-	5,000	100	-

[†] – Exceeds the critical value for the 95% confidence level established by 5,000 synthetics with a uniform random choice rule.


is illustrated with the blue bars in Figure 10(a), and exact values are reported in Table 10(b). The critical value for the 95% confidence level implied by this sampling distribution is 21. Relative to this critical value, 198 of our 203 experimental participants (98%) had a Houtman-Maks that met or exceeded this threshold (i.e. $HM \geq 21$). If the cutoff for GARP-compliance is raised to the critical value for the 99% confidence level of 22, 187 participants (92%) had a Houtman-Maks that met or exceeded this threshold (i.e. $HM \geq 22$). Generally speaking, we find very high rates of GARP-compliance.

3.2 Assessing The Hurwicz Expected Utility Representation

In the ambiguity triangle, the Hurwicz Expected Utility model exhibits indifference curves that are linear and parallel. This geometry necessitates a particular pattern of choices in an experiment like ours. Consider, for example, the three budgets, each with a distinct slope, depicted in Figure 11. Suppose that the 50/50 lottery was chosen from the middle budget (i.e. the budget connecting the 100% ambiguous alternative to the 50/50 lottery). Any linear indifference curve that rationalizes this choice must lie in the shaded area. And because HEU indifference curves must be parallel, the choice from any steeper budget, like the upper-most budget in Figure 11, must also be the lottery on the hypotenuse. When a budget is flatter than the middle budget, however, no such prediction can be placed on choices. If the original choice from the middle budget was the 100% ambiguous alternative, then the reverse structure would be required. Flatter budgets would require the most ambiguous alternative available to be selected, while no such requirement could be placed on choices from steeper budgets. If the original choice was an alternative from the strict interior of the middle budget, any HEU indifference curve that rationalizes this choice must lie on top of that budget line. So, for any steeper budget the lottery must be selected while for any flatter budget the most ambiguous alternative must be selected. The comprehensive technical details for testing the HEU representation in this way, and algorithmic details, can be found in Appendix A.2.

Paralleling the logic of assessing “how badly” choices can depart from the generalized axiom, we adapt the logic of Houtman-Maks (HM) to our test of the Hurwicz Expected Utility representation. We define “Houtman-Maks HEU” (HM-HEU) as the largest subset of choices that conform to our test. Figure 12(a) is a histogram of HM-HEU for the 203 participants in our experiment (red

Figure 11: Testing for the Linear and Parallel Indifference Curves Required Under Hurwicz Expected Utility



bars). The supporting data are in Table 12(b). The blue bars in Figure 12(a) show the sampling distribution for HM-HEU implied by 5,000 synthetic experimental participants who had a 50/50 choice rule for just the endpoints of the budgets – these 5,000 synthetics’ choices were restricted to be either the fully-specified lottery or the most ambiguous alternative. The critical value for the 95% confidence level for this sampling distribution was 20. Of the 203 participants in our experiment, 120 (59.1%) had an HM-HEU that met or exceeded this critical value. Considering the critical value for the 99% confidence level of 21, 103 participants (50.7%) had an HM-HEU that met or exceeded this value. Generally speaking, Hurwicz Expected Utility is a good representational assumption for the majority of our experimental participants.


3.2.1 Estimating Hurwicz Expected Utility Parameters

This section describes the procedures we used to estimate Hurwicz Expected Utility preference parameters. We assume random utility (RU) with additively separable choice noise (McFadden, 1981). Denoting L^{jk} as the k^{th} lower envelope lottery available in the j^{th} choice-situation, the random utility from choosing L^{jk} is given by

$$RU(L^{jk}) = HEU(L^{jk}) + \varepsilon^{jk}, \quad (9)$$

Figure 12: Houtman-Maks for the Hurwicz Expected Utility Representation

(a) Histogram



(b) Supporting Data

Houtman-Maks	Humans			Synthetics		
	Freq.	%	Cum.	Freq.	%	Cum.
5	1	0.49	0.49	0	0	0
6	2	0.99	1.48	0	0	0
7	2	0.99	2.46	0	0	0
9	3	1.48	3.94	0	0	0
10	0	0	3.94	0	0	0
11	1	0.49	4.43	0	0	0
12	3	1.48	5.91	0	0	0
13	4	1.97	7.88	156	3.12	3.12
14	4	1.97	9.85	562	11.24	14.36
15	15	7.39	17.24	983	19.66	34.02
16	7	3.45	20.69	1,098	21.96	55.98
17	15	7.39	28.08	972	19.44	75.42
18	10	4.93	33	665	13.3	88.72
19	16	7.88	40.89	371	7.42	96.14
20	17 [‡]	8.37	49.26	126	2.52	98.66
21	23 [‡]	11.33	60.59	45	0.9	99.56
22	17 [‡]	8.37	68.97	18	0.36	99.92
23	17 [‡]	8.37	77.34	3	0.06	99.98
24	13 [‡]	6.4	83.74	1	0.02	100
25	33 [‡]	16.26	100	0	0	100
Total	203	100	-	5,000	100	-

‡ – Exceeds the critical value for the 95% confidence level established by 5,000 synthetics with a 50/50 choice rule for only the endpoints of a budget.

where ε^{jk} is an independent and identically distributed random variable. The systematic, HEU-component of $RU(L^{jk})$, for the CHF 60 and CHF 20 outcomes available in our experiment, can be written as

$$\text{HEU}(\underline{p}_{60}, \underline{p}_{20}, y) = \alpha \left[\underline{p}_{60} u_{60} + (\underline{p}_{20} + y) u_{20} \right] + (1 - \alpha) \left[(\underline{p}_{60} + y) u_{60} + \underline{p}_{20} u_{20} \right]. \quad (10)$$

Normalizing outcome utilities such that $u_{20} = 0$ and $u_{60} = 1$, Expression 10 simplifies to

$$\text{HEU}(\underline{p}_{60}, \underline{p}_{20}, y) = \underline{p}_{60} + (1 - \alpha)y, \quad (11)$$

an expression with just one parameter (α) that captures an individual's everywhere-constant ambiguity attitude. A participant chooses the k^{th} lower envelope lottery (i.e. $L^{j*} = L^{jk}$) if and only if $RU(L^{j*}) \geq RU(L^{jk})$. Thus, the probability that the subject chooses the k^{th} lower envelope lottery in choice situation j is given by

$$\text{Prob} \left[L^{j*} = L^{jk} \right] = \text{Prob} \left[\text{HEU}(L^{j*}) - \text{HEU}(L^{jk}) \geq \varepsilon^{jk} - \varepsilon^{j*} \right]. \quad (12)$$

We assume that ε^{jk} follows a type I extreme value distribution with scale parameter σ . So the choice probabilities take the standard form:

$$\text{Prob} \left[L^{j*} = L^{jk} \right] = \frac{\exp\left(\frac{\text{HEU}(L^{jk})}{\sigma}\right)}{\sum_{k=1}^6 \exp\left(\frac{\text{HEU}(L^{jk})}{\sigma}\right)}. \quad (13)$$

From these choice probabilities we obtain individual i 's contribution to the likelihood function

$$f(\mathbb{L}; \theta) = \prod_{j=1}^{25} \prod_{k=1}^6 \text{Prob} \left[L^{jk} \right] \mathbb{1} \left[L_i^{j*} = L_i^{jk} \right], \quad (14)$$

where \mathbb{L} is the set of all lower envelope lotteries in the choice experiment and the function $\mathbb{1}[\cdot]$ is an indicator that returns one if the condition in the bracket is true and zero otherwise.

Under the assumption that all participants have the same ambiguity preferences (i.e. homoge-

neous preferences), the log-likelihood function is

$$\mathcal{L}_{\text{HEU}}^1 = \sum_{i=1}^n \ln f(\mathbb{L}; \alpha, \sigma), \quad (15)$$

where the subscript i denotes the i^{th} of the $n = 120$ individuals with an HEU representation. The HEU model involves estimation of two parameters, the index of ambiguity aversion α , and the dispersion parameter σ .

We then relax the assumption of homogenous preferences and explore the possibility that our data were generated by multiple and distinct preference types. To do so, we employ a finite mixture approach.⁸ The principal idea of such models is assigning each subject to one of C different preference types. Each type is endogenously characterized by a distinct vector of parameters, (α_c, σ_c) , with $c \in \{1, \dots, C\}$. In addition, the estimation procedure yields estimates of the proportion of the sample that belongs to each type, π_c . Summing over all C behavioral types yields the complete log-likelihood function

$$\mathcal{L}_{\text{HEU}}^C = \sum_{i=1}^n \ln \sum_{c=1}^C \pi_c f(\mathbb{L}; \alpha_c, \sigma_c). \quad (16)$$

In finite mixture models the number of behavioral types needs to be fixed prior to estimation. There is a large literature discussing the optimal number of types but there is no common agreement on which measure is optimal (for a summary see McLachlan and Peel, 2000). Here, we follow a similar route as Bruhin et al. (2010) and stop incrementing the number of types as soon as adding more types does not generate novel qualitative insights. This means, for example, that if a model with only one type indicates that this type is ambiguity averse while a model with two types shows that there is also a substantial minority of ambiguity seeking types, the characterization of the population with two types is preferable because otherwise we would overlook the substantial minority of subjects whose behavior is qualitatively different (i.e., ambiguity seeking).

To estimate parameters we use the iterative expectation maximization (EM) algorithm, details of which can be found in Dempster et al. (1977) and the supplementary material of Bruhin et al. (2010). Reported standard errors are obtained by the bootstrap method with 1,000 replications

⁸Applications of finite mixture models are discussed in El Gamal and Grether (1995), Stahl and Wilson (1995), Houser et al. (2004), Bruhin et al. (2010), Fehr-Duda et al. (2010), Conte et al. (2011), and Fehr-Duda and Epper (2012).

(Efron and Tibshirani, 1993). Resampling is done at the level of participants. We prevent the issue that the order of types in the parameter vector might change during resampling (i.e. label switching) by first calculating the Euclidean distance between the replication types' parameter vectors and the original types' parameter vectors, and then reordering the full parameter vector accordingly.

3.2.2 Estimation Results: Examining Ambiguity Preference Types

Estimation results are reported in Table 2. Model 1 assumes homogeneous preferences across the sub-sample of 120 participants that are HEU representable. The parameter estimates and bootstrapped 95% confidence intervals for ambiguity preference parameter (α) are plotted in **blue** in the top portion of Figure 13. That the confidence interval lies above 0.5 indicates that, on average, participants exhibited ambiguity aversion.

Model 2 relaxes the restriction of homogeneous preferences and allows for two preference types. Fitted ambiguity parameters, α_1 and α_2 , and their standard errors (s.e.) are shown in the middle column of Table 2. The middle portion of Figure 13 shows the bootstrapped 95% confidence intervals for these two preference types. The confidence interval for Type 1 preferences, colored in **blue**, is greater than 0.5, consistent with ambiguity aversion. The confidence interval for Type 2 preferences, colored in **red**, is less than 0.5, consistent with ambiguity seeking. The finite mixture model also yields an estimate for the parameter π_1 , the proportion of the sample that is associated with the Type 1 preferences. This estimate, shown in the first row for the Type 1 preferences in Table 2, indicates that about 80% of participants are ambiguity averse. The remaining 20% of participants belong to the Type 2 preference type, or ambiguity seeking.


The estimates for Model 3, which has three preference types, are in the right-most column of Table 2 and the bootstrapped 95% confidence intervals for the three preference types are shown in the lower portion of Figure 13. As with Model 2, the **blue** and **red** confidence intervals correspond to ambiguity aversion and neutrality. The **green** confidence interval, from a strictly statistical perspective, also corresponds to ambiguity aversion. However, the difference between the fitted parameter and 0.5 is negligible – it is only 3% larger than 0.5. We therefore interpret this type as exhibiting (near) ambiguity neutrality. In terms of the estimated proportions for each of

Table 2: Model Estimates for the 120 Participants Representable With Hurwicz Expected Utility

		Model 1 (One Type)	Model 2 (Two Types)	Model 3 (Three Types)
Type 1	π_1 (s.e.)	1.000 (-)	0.795* (0.096)	0.483* (0.084)
	α_1 (s.e.)	0.529 * ^{1/2} (0.008)	0.562 * ^{1/2} (0.017)	0.583 * ^{1/2} (0.025)
	σ_1 (s.e.)	0.025* (0.002)	0.020* (0.002)	0.026* (0.002)
Type 2	π_2 (s.e.)	-	0.205 (-)	0.302* (0.067)
	α_2 (s.e.)	-	0.390 * ^{1/2} (0.042)	0.517 * ^{1/2} (0.005)
	σ_2 (s.e.)	-	0.029* (0.005)	0.005* (0.002)
Type 3	π_3	-	-	0.215 (-)
	α_3 (s.e.)	-	-	0.395 * ^{1/2} (0.038)
	σ_3 (s.e.)	-	-	0.029* (0.004)
Individuals		120	120	120
Observations		3,000	3,000	3,000
\mathcal{L}		-3,190	-2,882	-2,681
BIC		6,396	5,804	5,425

Remarks: *^{1/2} denotes that α_j is different from 1/2 (ambiguity neutrality) at the 95% confidence level. * denotes difference from 0 at the 95% confidence level. Standard errors are calculated using 1,000 bootstrap replications. For further details see Appendix B.

Figure 13: Bootstrapped 95% Confidence Intervals for the Ambiguity Preference Parameters (α) in Table 2




these three preference types, 48% are ambiguity averse, 30% are nearly ambiguity neutral, and 22% are ambiguity seeking. If we further increase the number of types to four or five, no new qualitative insights emerge. There is always a type best characterized by near ambiguity neutrality and comparatively little choice noise and the overall share of ambiguity averse types remains relatively stable. Thus, the main consequence of increasing the number of types to four or five is the emergence of subdivisions among the ambiguity averse and the ambiguity seeking types.

3.3 Examining why Hurwicz Expected Utility was Violated

Section 3.2 reveals that a significant proportion of participants made choices that cannot be rationalized with Hurwicz Expected Utility (41%, or 83 of 203 participants; see Figure 12 for details). Two explanations for why HEU failed as a parsimonious model seem plausible. First, it could be that the everywhere-constant ambiguity attitude embedded in the HEU representation was violated, while the linearity of indifference curves was retained. That indifference curves would still be linear means that choices would be predominantly at the endpoints (corners) of the budgets in our experiment. In this case, the fraction of choices between the fully-specified lottery endpoint and the 'most ambiguous' endpoint should vary across the ambiguity triangle. A second explanation for why HEU failed as a parsimonious model would be that indifference curves in the ambiguity triangle were non-linear. A pattern of choices consistent with this explanation would mean that choices were, to a large extent, on the interior of the budgets in our experiment.

For this sub-sample of 83 participants we observe very low proportions of choices on the interior of the budgets. And the bubble plot in Figure 14 is illustrative of this general pattern. It displays choices made from five budgets, each with the same tradeoffs but differing likelihoods. As can be seen in the figure, only a small proportions of choices are on the interior of these budgets. Put another way, there are a large proportion of choices at the endpoints of the budgets. The ratio of choices between the fully-specified lottery endpoint and the most ambiguous endpoint, however, varies across the ambiguity triangle. This pattern of choices is roughly consistent with the first explanation above. We therefore consider a generalization of Hurwicz Expected Utility that relaxes the everywhere constant ambiguity attitude only.

Figure 14: A Bubble Plot of the Choices Made by 83 Participants on Five Budgets With a Constant Price Illustrating Non-constant Ambiguity Attitudes Across the Ambiguity Triangle



3.3.1 The β -Hurwicz Expected Utility Model

As an empirical expedient for accommodating a non-constant ambiguity attitude we introduce a generalization of the Hurwicz Expected Utility model. We call this generalization the β -Hurwicz Expected Utility (β -HEU) model. The β -HEU model weakens the everywhere-constant ambiguity attitude embedded in HEU by allowing indifference curves to fan-in or fan-out across the ambiguity triangle. The additional parameter, β , controls the extent to which indifference curves fan-in or fan-out across the ambiguity triangle. Formally, the β -HEU of a lower envelope lottery in our two-outcome experiment is given by

$$\begin{aligned} \beta\text{-HEU} \left(\underline{p}_{60}, \underline{p}_{20}, y \right) &= (\alpha + \beta l_j) \left[\left(\underline{p}_{20} + y \right) u_{20} + \underline{p}_{60} u_{60} \right] \\ &+ (1 - \alpha - \beta l_j) \left[\underline{p}_{20} u_{20} + \left(\underline{p}_{60} + y \right) u_{60} \right], \end{aligned} \quad (17)$$

where l_j is the likelihood for choice-situation j (i.e. the probability of receiving the CHF 60 outcome for the fully specified lottery available from the budget; see Section 2 for more details).

It is possible to interpret the β -HEU model as introducing a systematically varying ambiguity attitude, where the parameter β captures whether ambiguity preference is increasing ($\beta > 0$) or decreasing ($\beta < 0$) as the likelihood increases. To facilitate this interpretation, and make the link between β -HEU and HEU clearer, we utilize the following normalization for the likelihood term:

$$\tilde{l}_j = 10 \left(l_j - \frac{1}{2} \right). \quad (18)$$

Adopting this normalization facilitates a more intuitive interpretation for β . Specifically, the multiplicative term (10) will mean that β now measures the effect of an increase of 0.1 for the probability assigned to the good outcome for the fully specified lottery available from the budget. The subtractive term ($\frac{1}{2}$) will mean that α determines the slope of the β -HEU indifference curve that goes through the 50/50 lottery. Adopting this normalization into Equation 17 gives us the expression for β -HEU that we will use during estimation:


$$\begin{aligned} \beta\text{-HEU}(\underline{p}_{60}, \underline{p}_{20}, y) = & \left(\alpha + 10\beta \left(l_j - \frac{1}{2} \right) \right) \left[(\underline{p}_{20} + y) u_{20} + \underline{p}_{60} u_{60} \right] \\ & + \left(1 - \alpha - 10\beta \left(l_j - \frac{1}{2} \right) \right) \left[\underline{p}_{20} u_{20} + (\underline{p}_{60} + y) u_{60} \right]. \end{aligned} \quad (19)$$

Figure 15 shows example indifference curves for two parameterizations of β -HEU that adopt our likelihood normalization and are plotted in the ambiguity triangle. Both examples have $\alpha = \frac{1}{2}$ so that the indifference curve intersecting the ambiguity triangle at the 50/50 lottery (i.e. the midpoint of the hypotenuse) has a slope equivalent to ambiguity neutrality. Panel (a) shows a positive value of β which produces indifference curves that ‘fan-in’ across the ambiguity triangle. Put another way, positive values of β mean that aversion to ambiguity is increasing as the likelihood increases (i.e. as l_j , or \tilde{l}_j , increases). Negative values of β , like those depicted in panel (b) of Figure 15, have the opposite pattern – Indifference curves ‘fan-out’ across the ambiguity triangle. Fanning-out of indifference curves is consistent with an increasing preference for ambiguity as the likelihood increases.

Figure 15: Indifference Curves for the β -Hurwicz Expected Utility Model

(a) Increasing aversion to ambiguity: $\alpha = \frac{1}{2}, \beta > 0$
('Fanning In')

(b) Decreasing aversion to ambiguity: $\alpha = \frac{1}{2}, \beta < 0$
('Fanning Out')



3.3.2 Estimating β -Hurwicz Expected Utility Parameters

In general we follow the same procedures as those in Section 3.2.1. The only differences being that we now consider the $n = 83$ participants whose choices could not be rationalized with a HEU representation. Accordingly, we replace the systematic component of random utility with β -HEU and estimate three parameters per type: $(\alpha_c, \beta_c, \sigma_c)$. Normalizing the outcome utilities for the systematic β -HEU component of random utility, so that $u_{20} = 0$ and $u_{60} = 1$, we have

$$\beta\text{-HEU} \left(\underline{p}_{60}, \underline{p}_{20}, y \right) = \underline{p}_{60} + \left(1 - \alpha - 10\beta \left(l_j - \frac{1}{2} \right) \right) y. \quad (20)$$

The number of behavioral types (C) has to be fixed prior to estimation. We follow the same procedure as in Section 3.2.1 by incrementing the number of types until additional types provide no new qualitative insights. Here, we end up with $C = 2$ types, both showing considerable fanning across the simplex. Interestingly, one type exhibits ambiguity seeking when likelihood is low and ambiguity aversion when likelihood is high. Perforce, these preferences appear ambiguity neutral near the 'middle' of the ambiguity triangle. The other type reveals an average behavior which is everywhere ambiguity averse, and increasingly so as likelihood increases.⁹

⁹We also estimated models with $C = 3$ types. The additional type can be broadly described as a subgroup of the

3.3.3 Estimation Results: Examining β -Hurwicz Expected Utility Preference Types

Estimation results are reported in Table 3. Model 1 assumes homogeneous preferences across the sub-sample of 83 participants that are β -HEU-representable. The parameter estimates and a bootstrapped 95% confidence region for preference parameters α and β are plotted in **blue** in Figure 16(a). The confidence intervals indicate that α is larger than 0.5 and β is positive. In general, these homogeneous parameter estimates exhibit a large degree of ambiguity aversion. This is easy to see in Figure 17(a) which plots indifference curves for Model 1 as **blue** lines in the ambiguity triangle.

Model 2 estimates two types of preferences. Fitted preference parameters, α_1 , β_1 , α_2 , and β_2 , and their standard errors (s.e.), are shown in the right column of Table 3. Bootstrapped 95% confidence regions for both preference types are shown in panel (b) of Figure 16. The confidence region for Type 1 preferences is colored **green** while the confidence region for Type 2 preferences is colored **blue**. Indifference curves for these two types are plotted in panel (b) of Figure 17. Interestingly, Type 1 preferences exhibit ambiguity seeking in the lower portion of the ambiguity triangle and ambiguity aversion in the upper portion of the triangle. Type 2 preferences are everywhere ambiguity averse with levels of ambiguity aversion approaching extreme pessimism (i.e. $\alpha = 1$ under an HEU representation) in the upper portion of the ambiguity triangle. Approximately 43% of this sample is estimated to be associated with the Type 1 preferences (see π_1 for Model 2 in Table 3). Broadly speaking, the parameter estimates for Model 2 are consistent with the bubble plots discussed above – aversion to ambiguity increases as the likelihood for a good outcome increases.

¹‘everywhere ambiguity averse’ type when $C = 2$.

Table 3: Model Estimates for the 83 Participants Representable With β -Hurwicz Expected Utility

		Model 1 (One Type)		Model 2 (Two Types)	
Type 1	π_1 (s.e.)	1.000	(-)	0.428*	(0.102)
	α_1 (s.e.)	0.719 ^{*1/2}	(0.024)	0.494	(0.109)
	β_1 (s.e.)	0.136 *	(0.015)	0.261 *	(0.063)
	σ_1 (s.e.)	0.099 *	(0.008)	0.137 *	(0.026)
Type 2	π_2 (s.e.)	-		0.572	(-)
	α_2 (s.e.)	-		0.785 ^{*1/2}	(0.025)
	β_2 (s.e.)	-		0.093 *	(0.019)
	σ_2 (s.e.)	-		0.065 *	(0.010)
Individuals		83		83	
Observations		2,075		2,075	
\mathcal{L}		-2,816		-2,720	
BIC		5,654		5,494	

Remarks: ^{*1/2} denotes that α_j is different from 1/2 at the 95% confidence level. * denotes difference from 0 at the 95% confidence level. Standard errors are calculated using 1,000 bootstrap replications. For further details see Appendix B.

Figure 16: Bootstrapped 95% Confidence Regions Based on 1,000 Replications for the β -Hurwicz Expected Utility Preference Parameters Reported in Table 3






Figure 17: Indifference Curves for the β -Hurwicz Expected Utility Preference Parameters Reported in Table 3

(a) Model 1 – Homogeneous Preferences



(b) Model 2 – Two Preference Types



The dashed lines are for reference and represent ambiguity neutrality. Note that we only observe choices in a subdomain of the ambiguity triangle – Indifference curves in the upper-right and lower-left corner are therefore extrapolations.

4 Related Literature

We do not attempt to review the large literature on ambiguity in this section. For this purpose, the interested reader is referred to the excellent papers by Camerer and Weber (1992), Camerer (1995), and Etner et al. (2010). We limit ourselves here to a short discussion of the more recent wave of empirical research and refer to relevant theories where appropriate.

Motivated by the so-called ‘Ellsberg paradox,’ there has been extensive theoretical work focused on models that could accommodate ambiguity aversion and, as a result, canonical Ellsberg behavior. For example, Schmeidler (1989) put forward Choquet Expected Utility. This model assumes unique beliefs that can be represented with *capacities*. Unlike probabilistic beliefs, capacities need not be additive which makes it possible to rationalize Ellsberg-type behavior.

In contrast to the unique beliefs assumed in the Subjective and Choquet Expected Utility models, Gilboa and Schmeidler (1989) proposed a model in the multiple priors class. The so-called MaxMin Expected Utility model assumes that bets induce sets of probabilistic beliefs. A bet is then evaluated by the belief that has the lowest Expected Utility. A natural counterpart to the MaxMin Expected Utility model is the MaxMax Expected Utility model – a bet is evaluated by the belief that has the highest Expected Utility. Marinacci (2002) and Ghirardato et al. (2004) proposed an amalgam of the MaxMin and MaxMax Expected Utility models. The so-called α -MaxMin Expected Utility model represents a mixture of MaxMin and MaxMax evaluations of bets.¹⁰ In this way α -MaxMin Expected Utility is based on both beliefs and ambiguity attitudes (as captured by α). From an empirical perspective, this introduces serious challenges when attempting to measure both beliefs and preference parameters (Hey et al., 2010; Kothiyal et al., 2014).

At first blush, α -MaxMin and HEU appear quite similar. They are distinguished, however, in two important ways. First, HEU preferences are defined over sets of lotteries whereas α -MaxMin preferences are defined over acts (bets). Second, in the HEU model the best and the worst ex ante outcomes (\bar{L} and \underline{L}) are objectively defined entities. In contrast, in the α -MaxMin model the worst and the best probability distribution over outcomes are determined by an individual’s subjective beliefs.¹¹ As mentioned before, empirical identification of these subjective components

¹⁰The theory introduced in Ghirardato et al. (2004) considers beliefs as a component of the theory (i.e. beliefs are endogenous). An α -MaxMin Expected Utility model without endogenous beliefs was axiomatized by Jaffray (1994).

¹¹Putting aside the difference in domain (i.e. lower envelope lotteries vs. acts), these two models can make identical behavioral predictions if an α -MaxMin Expected Utility maximizer considers as his set of priors the full range of

is very challenging and requires comprehensive data or additional, potentially critical, assumptions. Thus, from an empirical perspective, the HEU approach is advantageous because the slope of the indifference curve fully captures the relevant information about ambiguity attitudes. Moreover, this slope can be easily measured with appropriately designed budgets in the space of lower envelope lotteries.

Our use of lower envelope lotteries to study ambiguity attitudes contributes to the small, but growing, literature that explores preferences over ‘sets of lotteries.’ While lower envelope lotteries are simpler than the more general ‘sets of lotteries’ setting investigated theoretically by Olszewski (2007), Ahn (2008) and Dumav and Stinchcombe (2013), this simplicity is advantageous for the empirical study of ambiguity attitudes. It allows, for example, applications of non-parametric revealed preference assessments of model classes. More specifically, we can use the simplicity of lower envelope lotteries to, roughly speaking, assign a price to ambiguity. This, in turn, permits non-parametric tests of (i) the generalized axiom of revealed preference and (ii) whether preferences can be characterized by linear and parallel indifference curves. The parsimony of lower envelope lotteries also permits a tractable and intuitive way to describe the pattern of choices that results from a violation of constant ambiguity attitudes such as when indifference curves “fan-in” or “fan-out”. In addition, Section 1 documents how Ellsberg’s experiment can be described using lower envelope lotteries, plotted in the ambiguity triangle, and canonical behavior therein rationalized with a simple model that has linear and parallel indifference curves.

Our paper also contributes to a growing literature that examines ambiguity attitudes in experimental settings (e.g. Abdellaoui et al., 2011; Ahn et al., 2014; Halevy, 2007; Hey et al., 2010; Hey and Pace, 2014; Stahl, 2014). For example, Abdellaoui et al. (2011) show that ambiguity attitudes can vary significantly depending on the source of ambiguity. The goal of their study was, however, different from ours as they were not interested in examining the GARP-compliance of subjects nor the representation of heterogeneity in ambiguity preferences for their choice data. Nor were they interested in the endogenous assignment of individuals to different preference classes. Instead, they developed a new tool (source functions) for translating events (i.e. collections of states) into a willingness to bet. By comparing source functions elicited for urns with known and unknown compositions, these authors provide a rich characterization of attitudes towards ambiguity. Al-

probabilities $[p, p + y]$.

though they focussed on different questions, their aggregate data also suggest that varying attitudes towards ambiguity are playing a role. Specifically, they found that for large probabilities ($p > 0.5$) source functions were significantly higher for an urn with a known composition (urn K) than for an urn with a completely unknown composition (urn U). In contrast, when probabilities were small ($p < 0.5$) there was no meaningful difference in source functions between the urns. Because the difference between the source functions for K and U provides a quantitative measure of ambiguity aversion, this finding suggests that ambiguity aversion increases with the likelihood of the good outcome. Our non-parametric test for constant ambiguity aversion suggests, however, that increasing ambiguity aversion is not a property of all subjects; 60% of our subjects exhibit constant ambiguity attitudes and within this group we find individuals who exhibit ambiguity aversion, ambiguity seeking and ambiguity neutrality. But we also find that a substantial share of subjects (40%) exhibit increasing ambiguity aversion. This property of our data can be parsimoniously captured using a β -Hurwicz Expected Utility model with indifference curves that 'fan-in' across the ambiguity triangle.

Ahn et al. (2014) examined ambiguity preferences by asking subjects to allocate a budget between three Arrow-securities. Each security paid unity in one of three mutually exclusive states, and zero otherwise. For one of the states the probability that it would be realized was known to be $\frac{1}{3}$. For the other two states the exact probabilities were unknown. The approach employed by Ahn et al. (2014) highlights a natural complementarity between examining ambiguity preferences with Arrow-securities and the lower envelope lottery approach employed here. Ahn et al. (2014) relied on varying monetary outcomes under a fixed set of probability distributions. In contrast, the lower envelope lottery approach used here varies minimum probabilities and ambiguity (i.e. the set of lotteries) while holding outcomes fixed. Our paper also differs from Ahn et al. (2014) because we test, non-parametrically, for constant ambiguity attitudes and show that roughly 40% of our subjects display varying ambiguity attitudes - a finding that suggests that the assumption of constant ambiguity attitudes made in Ahn et al. (2014) may perhaps not be appropriate for all of their subjects. In addition, we find a much higher percentage of subjects that display ambiguity aversion: Ahn et al. (2014) found that only 10% of their subjects showed ambiguity aversion. In our data 29% of subjects exhibit constant ambiguity aversion and 23% display varying ambiguity aversion that increases from moderate to high levels as the likelihood for the good outcome in-

creases. Future research should clarify whether these differences are methodological in origin or due to differences in subject pools.

In a setup distinct from ours, Halevy (2007) examined different explanations for non-neutrality towards ambiguity. His experiment distinguished between theories that model ambiguity aversion as violations of probabilistic sophistication and theories that model choice under uncertainty using two-stage prospects. He found substantial support for the latter explanation. While our paper asks different questions, by focusing on GARP compliance and parsimonious empirical characterizations of heterogeneity in ambiguity attitudes, we believe that Halevy (2007)'s finding suggests a promising extension of the present study. In particular, in our setup uncertainty is objectively resolved in one stage. So, violations of the reduction of compound lottery axiom are unlikely to play a role in our observed non-neutral ambiguity attitudes. We can, however, easily introduce a two-stage resolution of uncertainty in our setup. By comparing behavior under one-stage and two-stage resolution of uncertainty it would be straightforward to measure the change in non-neutral ambiguity attitudes that is due to reducing compound lotteries.

Hey et al. (2010) and Hey and Pace (2014) presented data from experiments involving bets on balls mixing in a transparent bingo blower. Both of these papers employed parametric estimation techniques to examine the differential fits for a range of models at the individual level. As demonstrated in these ambitious studies (see also Kothiyal et al., 2014), identification of preference parameters and beliefs using, for example, an α -MaxMin-Expected Utility specification, can be challenging. Due to the large number of parameters that have to be estimated, identification requires either a large number of observations or strong auxiliary parametric assumptions. The examination of indifference curves in the ambiguity triangle obviates these challenges. This is because the slope of indifference curves reflects ambiguity attitudes under a Hurwicz Expected Utility representation (or simple extensions of HEU in case of varying ambiguity attitudes). This does not require disentangling preferences from beliefs.

Stahl (2014) conducted an experiment in which participants made a series of choices in settings similar to Ellsberg's two- and three-color configurations. Stahl (2014)'s primary focus was on characterizing heterogeneity in ambiguity attitudes. To accomplish this he estimated (i) individual-level models, (ii) a representative agent model, and (iii) finite mixture models for a representation which characterized agents/types by two parameters: an index of ambiguity aversion and

an error dispersion parameter. Note that this means his setup did not permit an examination of varying ambiguity attitudes. Stahl (2014) found a comparatively small proportion of subjects that exhibited ambiguity aversion ($\approx 12\%$) and a rather large group of subjects ($> 60\%$) that revealed a behavior that “... suggests that the majority of participants found the choice tasks profoundly confusing” (Stahl, 2014, p.617). We believe that participants in our experiment were not confused because (i) our choice situations were very easy to understand (see Section 2), (ii) the rate of GARP-compliance was extremely high, (iii) we found no types with a random choice rule, and (iv) the assignment of individuals to different preference classes is very clean.

5 Discussion and Conclusion

This paper introduced lower envelope lotteries as a simple framework for examining ambiguity preferences. This simplicity can be advantageous in the experimental settings frequently explored in studies of choices involving ambiguity. For example, as we documented in Section 1, the Ellsberg experiment can be described with lower envelope lotteries and plotted in a figure we call the ambiguity triangle. Moreover, canonical choices in Ellsberg’s experiment can be rationalized using a parsimonious model that is an amalgam of Expected Utility and the Hurwicz criterion.

Because the slope of indifference curves captures ambiguity attitudes the ambiguity triangle constitutes a useful tool for examining basic, yet important, empirical patterns of behavior under uncertainty. It enables the examination of GARP-compliance, allows for a non-parametric assessment of whether ambiguity attitudes are constant or vary across the triangle and, in combination with finite mixture methods, one can parsimoniously characterize the heterogeneity of ambiguity attitudes in the triangle. Moreover, subjects’ response to variations in the slope and the position of linear budget lines on which the lower envelope lotteries are located also provides descriptive insights into the nature of varying ambiguity attitudes.

For future research the setting of lower envelope lotteries can be used as an input toward a richer characterization of choices made in the face of uncertainty. One potential direction could be an exploration of whether individuals that exhibit increasing ambiguity aversion also commit Allais-type violations of Expected Utility (Allais, 1953; Burghart et al., 2014). For Allais-type violations it is widely taken that risk aversion is increasing as the likelihood for the good outcome

increases (see Machina (1982), Hypothesis II and Figure 5b). Here, we show that many people exhibit increasing ambiguity aversion as likelihood increases. This raises the question of whether individuals who exhibit increasing ambiguity aversion will also exhibit increasing risk aversion.

Another direction that seems promising is using lower envelope lotteries to examine ambiguity preferences when a risky alternative is not available in the choice set. If behavior differs when removing the zero-ambiguity alternative this could challenge existing theories of choice under uncertainty. Further, in a setting where lower envelope lotteries are defined on complementary events, our setup could also shed further light on violations of probabilistic sophistication. Previous experimental work by Baillon and Bleichrodt (2015) indicates that the additivity of disjoint events required under probabilistic sophistication fails when eliciting matching probabilities for naturally occurring uncertainties.

Another prospective direction could be an exploration of whether the two-stage resolution of lower envelope lotteries has an effect on the results documented in this paper. Such an extension seems natural given that Halevy (2007) documents a relationship between ambiguity aversion, as elicited with Ellsberg-style urns, and violations of the reduction of compound lotteries. Lower envelope lotteries provide a simple and intuitive setting for deeper explorations of these questions, and others, in the domain of uncertainty.

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, 695–723.
- AFRIAT, S. N. (1967): "The Construction of Utility Functions from Expenditure Data," *International Economic Review*, 8, 67–77.
- AHN, D., S. CHOI, D. GALE, AND S. KARIV (2014): "Estimating Ambiguity Aversion in a Portfolio Choice Experiment," *Quantitative Economics*, 5, 195–223.
- AHN, D. S. (2008): "Ambiguity Without a State Space," *Review of Economic Studies*, 75, 3–28.
- ALLAIS, M. (1953): "Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'école américaine," *Econometrica*, 21, 503–546.
- ARROW, K. J. AND L. HURWICZ (1972): "An Optimality Criterion for Decision-Making under Ignorance," in *Uncertainty and Expectations in Economics: Essays in Honour of G.L.S. Shackle*, ed. by C. F. Carter and J. L. Ford, Oxford: Basil Blackwell, 1–11.
- BAILLON, A. AND H. BLEICHRODT (2015): "Testing Ambiguity Models through the Measurement of Probabilities for Gains and Losses," *American Economic Journal: Microeconomics*, forthcoming.
- BECKER, S. W. AND F. O. BROWNSON (1964): "What Price Ambiguity? Or the Role of Ambiguity in Decision-Making," *Journal of Political Economy*, 72, 62–73.
- BRUHIN, A., H. FEHR-DUDA, AND T. EPPER (2010): "Risk and Rationality: Uncovering Heterogeneity in Probability Distortion," *Econometrica*, 78, 1375–1412.
- BURGHART, D. R., T. EPPER, AND E. FEHR (2014): "The Two Faces of Independence: Betweenness and Homotheticity," *mimeo*.
- CAMERER, C. (1995): "Individual Decision Making," in *The Handbook of Experimental Economics*, ed. by A. Kagel, J. & Roth, Princeton, NJ: Princeton University Press, 587–683.
- CAMERER, C. AND M. WEBER (1992): "Recent Developments in Modeling Preferences: Uncertainty and Ambiguity," *Journal of Risk and Uncertainty*, 5, 325–370.
- CONTE, A., J. D. HEY, AND P. G. MOFFATT (2011): "Mixture Models of Choice Under Risk," *Journal of Econometrics*, 162, 79–88.
- DEMPSTER, A., N. LAIRD, AND D. RUBIN (1977): "Maximum Likelihood from Incomplete Data via the EM Algorithm," *Journal of the Royal Statistical Society, Series B*, 39, 1–38.
- DEMPSTER, A. P. (1967): "Upper and Lower Probabilities Induced by a Multivalued Mapping," *The Annals of Mathematical Statistics*, 38, 325–339.
- DOBBS, I. M. (1991): "A Bayesian Approach to Decision-Making under Ambiguity," *Economica*, 58, 417–440.
- DUMAV, M. AND M. STINCHCOMBE (2013): "The Von Neumann/Morgenstern Approach To Ambiguity," *Mimeo*, 1–37.

- EFRON, B. AND R. J. TIBSHIRANI (1993): *An Introduction to the Bootstrap*, New York: Chapman & Hall.
- EINHORN, H. J. AND R. M. HOGARTH (1986): "Decision Making under Ambiguity," *Journal of Business*, 59, S225–S250.
- EL GAMAL, M. A. AND D. M. GREYER (1995): "Are People Bayesian? Uncovering Behavioral Strategies," *Journal of the American Statistical Association*, 90, 1137–1145.
- ELLSBERG, D. (1961): "Risk, Ambiguity, and the Savage Axioms," *Quarterly Journal of Economics*, 75, 643–669.
- ETNER, J., M. JELEVA, AND J.-M. TALLON (2010): "Decision Theory Under Ambiguity," *Journal of Economic Surveys*, 26, 234–270.
- FEHR-DUDA, H., A. BRUHIN, T. EPPER, AND R. SCHUBERT (2010): "Rationality on the Rise: Why Relative Risk Aversion Increases With Stake Size," *Journal of Risk and Uncertainty*, 40, 147–180.
- FEHR-DUDA, H. AND T. EPPER (2012): "Probability and Risk: Foundations and Economic Implications of Probability-Dependent Risk Preferences," *Annual Review of Economics*, 4, 567–593.
- GHIRARDATO, P., F. MACCHERONI, AND M. MARINACCI (2004): "Differentiating ambiguity and ambiguity attitude," *Journal of Economic Theory*, 118, 133–173.
- GILBOA, I. AND D. SCHMEIDLER (1989): "Maxmin Expected Utility With Non-Unique Prior," *Journal of Mathematical Economics*, 18, 141–153.
- HALEVY, Y. (2007): "Ellsberg Revisited: An Experimental Study," *Econometrica*, 75, 503–536.
- HEY, J. D., G. LOTITO, AND A. MAFFIOLETTI (2010): "The Descriptive and Predictive Adequacy of Theories of Decision Making Under Uncertainty / Ambiguity," *Journal of Risk and Uncertainty*, 41, 81–111.
- HEY, J. D. AND N. PACE (2014): "The explanatory and predictive power of non two-stage-probability theories of decision making under ambiguity," *Journal of Risk and Uncertainty*, 49, 1–29.
- HOUSER, D., M. KEANE, AND K. MCCABE (2004): "Behavior in a Dynamic Decision Problem: An Analysis of Experimental Evidence Using a Bayesian Type Classification Algorithm," *Econometrica*, 72, 781–822.
- HOUTMAN, M. AND J. A. H. MAK (1985): "Determining all Maximal Data Subsets Consistent with Revealed Preference," *Kwantitatieve Methoden*, 19, 89–104.
- HURWICZ, L. (1951): "Some Specification Problems and Applications to Econometric Models," *Econometrica*, 19, 343–344.
- JAFFRAY, J.-Y. (1994): "Dynamic Decision Making With Belief Functions," in *Advances in the Dempster-Shafer Theory of Evidence*, ed. by R. R. Yager, M. Fedrizzi, and J. Kacprzyk, New York: Wiley, 331–352.
- KNIGHT, F. H. (1921): "Risk, Uncertainty and Profit," Boston, MA: Hart, Schaffner & Marx; Houghton Mifflin Co.

- KOTHİYAL, A., V. SPINU, AND P. P. WAKKER (2014): "An experimental test of prospect theory for predicting choice under ambiguity," *Journal of Risk and Uncertainty*, 48, 1–17.
- LUCE, R. D. AND H. RAIFFA (1958): *Games and Decisions*, New York: John Wiley & Sons, 3rd ed.
- MACHINA, M. (1982): "'Expected Utility' Analysis without the Independence Axiom," *Econometrica*, 50, 277–323.
- MARINACCI, M. (2002): "Probabilistic Sophistication and Multiple Priors," *Econometrica*, 70, 755–764.
- MARSCHAK, J. (1950): "Rational Behavior, Uncertain Prospects, and Measurable Utility," *Econometrica*, 18, 111–141.
- MCFADDEN, D. (1981): "Econometric Models of Probabilistic Choice," in *Structural Analysis of Discrete Data with Econometric Applications*, ed. by C. Manski, MIT Press, Cambridge MA.
- MCLACHLAN, G. AND D. PEEL (2000): *Finite Mixture Models*, Wiley Series in Probabilities and Statistics. New York: Wiley.
- OLSZEWSKI, W. (2007): "Preferences Over Sets of Lotteries," *Review of Economic Studies*, 74, 567–595.
- SAVAGE, L. J. (1954): *The Foundations of Statistics*, New York: John Wiley & Sons.
- SCHMEIDLER, D. (1989): "Subjective Probability and Expected Utility without Additivity," *Econometrica*, 57, 571–587.
- SHAFER, G. (1976): *A Mathematical Theory of Evidence*, Princeton, NJ: Princeton University Press.
- SINISCALCHI, M. (2009): "Vector Expected Utility and Attitudes Toward Variation," *Econometrica*, 77, 801–855.
- SLOVIC, P. AND A. TVERSKY (1974): "Who Accepts Savage's Axiom?" *Behavioral Science*, 19, 368–373.
- STAHL, D. O. (2014): "Heterogeneity of Ambiguity Preferences," *Review of Economics and Statistics*, 96, 609–617.
- STAHL, D. O. AND P. W. WILSON (1995): "On Players' Models of Other Players: Theory and Experimental Evidence," *Games and Economic Behavior*, 10, 218–254.
- SUGDEN, R. (2004): "Alternatives to Expected Utility: Foundations," in *Handbook of Utility Theory*, ed. by S. Barbera, P. Hammond, and C. Seidl, Springer US, 685–755.
- VARIAN, H. R. (1982): "The Nonparametric Approach to Demand Analysis," *Econometrica*, 50, 945–973.

Online Appendix

A Analytical Details

A.1 Testing For GARP

Varian (1982) provides a convenient algorithm for testing the Generalized Axiom of Revealed Preferences (GARP) and, in general, we adopt his notation and methods here. Following Burghart et al. (2014), testing GARP in this setting is computationally less intensive, and more intuitive, if we interpret lower envelope lotteries as two-dimensional demand vectors, relative to the worst outcome. To do this we transform lower envelope lotteries into two dimensions. Denoting the lower envelope lottery selected from the i^{th} budget as L^{i*} , we use the following transformation:

$$L^{i*} = (p_{60}^{i*}, p_{20}^{i*}, y^{i*}) \rightarrow (p_{60}^{i*} + y^{i*}, p_{60}^{i*}) = (x_1^{i*}, x_2^{i*}) = \mathbf{x}^{i*} \quad (21)$$

This transformation means that \mathbf{x}^{i*} can be interpreted as a two-dimensional demand vector, relative to the worst outcome. Denote as \mathbf{q}^i the two-dimensional vector of prices that describes the linear budget from which \mathbf{x}^{i*} was selected.

The starting point for testing GARP is to construct a directly revealed preferred graph. For the 25 budgets in our experiment, and each participants choices $(\mathbf{x}^{1*}, \dots, \mathbf{x}^{25*})$, we construct a 25 by 25 matrix M (the directly revealed graph) whose ij^{th} entry is given by

$$m_{ij} = \begin{cases} 1 & \text{if } \mathbf{q}^i \cdot \mathbf{x}^{i*} \geq \mathbf{q}^i \cdot \mathbf{x}^{j*} \\ 0 & \text{otherwise} \end{cases} \quad (22)$$

We then construct the indirectly revealed preferred graph MT , which is the transitive closure of M . The ij^{th} entry of the closure is given by

$$mt_{ij} = \begin{cases} 1 & \text{if } m_{ij}^{25} > 0 \\ 0 & \text{otherwise} \end{cases} \quad (23)$$

where m_{ij}^{25} is the ij^{th} entry of the matrix $M^{25} = MM \cdots M$. If $mt_{ij} = 1$ and $\mathbf{q}^j \cdot \mathbf{x}^{j*} > \mathbf{q}^j \cdot \mathbf{x}^{i*}$ for

some i and j (with $\mathbf{x}^{i*} \neq \mathbf{x}^{j*}$) there is a GARP violation.

To calculate Houtman-Maks, conditional on observing at least one violation of GARP, we use a brute force approach. We check whether all subsets of size 24 are consistent with GARP (using the above algorithm). If no subset of size 24 is GARP compliant we take all subsets of size 23 and check whether any of these subsets are GARP compliant. We proceed in such a manner until we find at least one subset of the data that is GARP compliant. The cardinality of that subset is the Houtman-Maks.

A.2 Testing For Hurwicz Expected Utility

The Hurwicz Expected Utility model gives rise to indifference curves that linear and parallel in the ambiguity triangle. So, given a choice from one budget, the HEU indifference curve structure lets us make predictions about choices from other budgets. To make predictions, we need to compare the steepness of the i^{th} budget, denoted by the ratio of its prices $\left(\frac{q_1}{q_2}\right)^i$, to the steepness of the j^{th} budget, denoted by the ratio of its prices $\left(\frac{q_1}{q_2}\right)^j$.¹² We say:

1. The j^{th} budget is flatter than the i^{th} budget whenever $\left(\frac{q_1}{q_2}\right)^j < \left(\frac{q_1}{q_2}\right)^i$
2. The j^{th} budget is steeper than the i^{th} budget whenever $\left(\frac{q_1}{q_2}\right)^j > \left(\frac{q_1}{q_2}\right)^i$

We can use these steepness comparisons to make an exhaustive list of predictions. Using \mathbf{x}^{i*} to denote the choice from the i^{th} budget:

- **Case 1:** The choice from the i^{th} budget is the most ambiguous alternative available (i.e. the alternative was on the horizontal or vertical leg of the ambiguity triangle). Denote this “corner” alternative by $\lfloor \mathbf{x} \rfloor^i$

$$\mathbf{x}^{i*} = \lfloor \mathbf{x} \rfloor^i \rightarrow \begin{cases} \mathbf{x}^j = \lfloor \mathbf{x} \rfloor^j & \text{when } \left(\frac{q_1}{q_2}\right)^j < \left(\frac{q_1}{q_2}\right)^i \\ \text{no prediction} & \text{when } \left(\frac{q_1}{q_2}\right)^j > \left(\frac{q_1}{q_2}\right)^i \\ \text{no prediction} & \text{when } \left(\frac{q_1}{q_2}\right)^j = \left(\frac{q_1}{q_2}\right)^i \end{cases}$$

- **Case 2:** The choice from i^{th} budget fully-specified lottery (i.e. the alternative selected was on the hypotenuse in the ambiguity triangle). Denote this “corner” alternative by $\lceil \mathbf{x} \rceil^i$

$$\mathbf{x}^{i*} = \lceil \mathbf{x} \rceil^i \rightarrow \begin{cases} \text{no prediction} & \text{when } \left(\frac{q_1}{q_2}\right)^j < \left(\frac{q_1}{q_2}\right)^i \\ \mathbf{x}^j = \lceil \mathbf{x} \rceil^j & \text{when } \left(\frac{q_1}{q_2}\right)^j > \left(\frac{q_1}{q_2}\right)^i \\ \text{no prediction} & \text{when } \left(\frac{q_1}{q_2}\right)^j = \left(\frac{q_1}{q_2}\right)^i \end{cases}$$

- **Case 3:** The choice from the i^{th} budget is on the “interior” (i.e. not Case 1 or Case 2). Denote

¹² See Burghart et al. (2014) for a similar test in the domain of choice under risk.

this set of alternatives by $[\mathbf{x}]^i$

$$\mathbf{x}^{i*} \in [\mathbf{x}]^i \rightarrow \begin{cases} \mathbf{x}^j = \lfloor \mathbf{x} \rfloor^j & \text{when } \left(\frac{q_1}{q_2}\right)^j < \left(\frac{q_1}{q_2}\right)^i \\ \mathbf{x}^j = \lceil \mathbf{x} \rceil^j & \text{when } \left(\frac{q_1}{q_2}\right)^j > \left(\frac{q_1}{q_2}\right)^i \\ \text{no prediction} & \text{when } \left(\frac{q_1}{q_2}\right)^j = \left(\frac{q_1}{q_2}\right)^i \end{cases}$$

This provides an exhaustive list of predictions under an HEU representation.

Given a set of predictions, it seems reasonable to assess two conditions:

- **Condition 1:** Predictions about choices are internally consistent.
- **Condition 2:** Choices are consistent with the predictions.

The first condition requires that choices from distinct budgets do not generate conflicting predictions about the choice from a third budget (also distinct). The second condition simply requires that choices are consistent with the set of (internally consistent) predictions. We say that choices are consistent with an HEU representation if these two conditions are met.

Algorithmically, we test the Condition 1 by constructing three square matrices, $\lfloor D \rfloor$, $\lceil D \rceil$, and D , with entries defined, respectively:

$$\lfloor d \rfloor_{ij} = \begin{cases} -1 & \text{if } \mathbf{x}^{i*} = \lfloor \mathbf{x} \rfloor^i \text{ and } \left(\frac{q_1}{q_2}\right)^j < \left(\frac{q_1}{q_2}\right)^i \\ 0 & \text{otherwise} \end{cases} \quad (24)$$

$$\lceil d \rceil_{ij} = \begin{cases} +1 & \text{if } \mathbf{x}^{i*} = \lceil \mathbf{x} \rceil^i \text{ and } \left(\frac{q_1}{q_2}\right)^j > \left(\frac{q_1}{q_2}\right)^i \\ 0 & \text{otherwise} \end{cases} \quad (25)$$

$$d_{ij} = \begin{cases} -1 & \text{if } \mathbf{x}^{i*} \in [\mathbf{x}]^i \text{ and } \left(\frac{q_1}{q_2}\right)^j < \left(\frac{q_1}{q_2}\right)^i \\ +1 & \text{if } \mathbf{x}^{i*} \in [\mathbf{x}]^i \text{ and } \left(\frac{q_1}{q_2}\right)^j > \left(\frac{q_1}{q_2}\right)^i \\ 0 & \text{otherwise} \end{cases} \quad (26)$$

Essentially, these three matrices document whether a choice from the j^{th} budget is predicted to be the most ambiguous alternative available (-1), the fully-specified lottery ($+1$), or neither (0),

based upon some original choice (\mathbf{x}^i) and steepness of that original budget $\left(\frac{q_1}{q_2}\right)^i$.

Because $[D]$ is the only matrix that has both $+1$ and -1 entries (i.e. only when $\mathbf{x}^{i*} \in [\mathbf{x}]^i$ can we get predictions of both $[\mathbf{x}]^j$ and $[\mathbf{x}]^j$) we first check it for internal consistency of its predictions.

We construct two row vectors, $\text{Max}[D]$ and $\text{Min}[D]$, where

$$\begin{aligned}\text{Max}[D] &= [\text{Max}\{[d]_{.1}, 0\}, \text{Max}\{[d]_{.2}, 0\}, \dots, \text{Max}\{[d]_{.J}, 0\}] \\ \text{Min}[D] &= [\text{Min}\{[d]_{.1}, 0\}, \text{Min}\{[d]_{.2}, 0\}, \dots, \text{Min}\{[d]_{.J}, 0\}]\end{aligned}\tag{27}$$

where $[d]_{.j}$ indicates the collection of entries in the j^{th} column of $[D]$. Notice that for the element-by-element multiplication, $\text{Max}[d]_j * \text{Min}[d]_j \geq 0, \forall j = 1, \dots, J$ if and only if predictions are internally consistent in $[D]$.

Conditional on $[D]$ exhibiting internal prediction consistency, all that remains is to verify that all three of the following conditions hold:

$$\begin{aligned}\text{Max}[d]_j \cdot \text{Min}[d]_j &\geq 0, \forall j = 1, \dots, J \\ \text{Max}[d]_j \cdot \text{Min}[d]_j &\geq 0, \forall j = 1, \dots, J \\ \text{Max}[d]_j \cdot \text{Min}[d]_j &\geq 0, \forall j = 1, \dots, J\end{aligned}\tag{28}$$

where

$$\begin{aligned}\text{Max}[D] &= [\text{Max}\{[d]_{.1}\}, \text{Max}\{[d]_{.2}\}, \dots, \text{Max}\{[d]_{.J}\}] \\ \text{Min}[D] &= [\text{Min}\{[d]_{.1}\}, \text{Min}\{[d]_{.2}\}, \dots, \text{Min}\{[d]_{.J}\}]\end{aligned}\tag{29}$$

If predictions are internally consistent (i.e. choices satisfy condition 1) it is a straightforward exercise to verify that choices are consistent with predictions (i.e. choices satisfy condition 2). If choices satisfy both condition 1 and condition 2, they pass our test for a Hurwicz Expected Utility representation.

B Classification

B.1 Hurwicz Expected Utility Parameter Estimates

Figures 18 and 19 depict histograms of the posterior probabilities for the two models presented in the main text. The posterior probability that an individual i choosing from the set of lower envelope lotteries \mathbb{L}_i belongs to type c is defined as (see Section 3.2.1 for notational definitions):

$$\tau_{ic} = \frac{\pi_c f(\mathbb{L}_i; \alpha_c, \sigma_c)}{\sum_{c=1}^C \pi_c f(\mathbb{L}_i; \alpha_c, \sigma_c)}. \quad (30)$$

The histograms in Figures 18 and 19 provide a positive impression regarding how well our classification procedure works. Specifically, the vast majority of participants are clearly assigned to one type because posterior probabilities are near $\tau = 1$ or $\tau = 0$. This clear assignment to one type or another is an important validation for finite mixture methods – if assignment to types is unclear, such that posterior probabilities are away from the bounds, then using finite mixture methods would be inappropriate.

Figure 18: Posterior Probabilities for HEU






Figure 19: Posterior Probabilities for HEU




B.2 β -Hurwicz Expected Utility

The posterior probability that an individual i choosing from the set of lower envelope lotteries \mathbb{L}_i belongs to type c is defined as:

$$\tau_{ic} = \frac{\pi_c f(\mathbb{L}_i; \alpha_c, \beta_c, \sigma_c)}{\sum_{c=1}^C \pi_c f(\mathbb{L}_i; \alpha_c, \beta_c, \sigma_c)}. \quad (31)$$

The histograms in Figures 20 show how well our classification procedure works. As with the posteriors for HEU, we find that the majority of participants are clearly assigned to one type because posterior probabilities are near $\tau = 1$ or $\tau = 0$.

Figure 20: Posterior Probabilities for β -HEU



C Experimental Details

C.1 Choice Situation Details

Table 4 lists normalized prices q_1 , q_2 , their ratio $\frac{q_1}{q_2}$, and likelihood l for each budget.

Table 4: Choice Situation (CS) Budget Details

CS	q_1	q_2	l	$\frac{q_1}{q_2}$
1	0.2	0.8	0.2	0.25
2	0.3	0.7	0.3	0.43
3	0.4	0.6	0.4	0.67
4	0.5	0.5	0.5	1.00
5	0.6	0.4	0.6	1.50
6	0.7	0.3	0.7	2.33
7	0.8	0.2	0.8	4.00
8	0.2	0.6	0.2	0.33
9	0.3	0.5	0.3	0.60
10	0.4	0.4	0.4	1.00
11	0.5	0.3	0.5	1.67
12	0.6	0.2	0.6	3.00
13	0.2	0.6	0.4	0.33
14	0.3	0.5	0.5	0.60
15	0.4	0.4	0.6	1.00
16	0.5	0.3	0.7	1.67
17	0.6	0.2	0.8	3.00
18	0.2	0.4	0.2	0.50
19	0.3	0.3	0.3	1.00
20	0.4	0.2	0.4	2.00
21	0.5	0.1	0.5	5.00
22	0.1	0.5	0.5	0.20
23	0.2	0.4	0.6	0.50
24	0.3	0.3	0.7	1.00
25	0.4	0.2	0.8	2.00

C.2 Exhaustive List of Each Lower Envelope Lottery

Table 5: An Exhaustive List of the Lower Envelope Lotteries Available in Each Choice Situation

CS	Risky End					Ambiguous End
	$\delta = 1.0$	$\delta = 0.8$	$\delta = 0.6$	$\delta = 0.4$	$\delta = 0.2$	$\delta = 0.0$
1	(0.2, 0.8, 0)	(0.16, 0.64, 0.2)	(0.12, 0.48, 0.4)	(0.08, 0.32, 0.6)	(0.04, 0.16, 0.8)	(0, 0, 1)
2	(0.3, 0.7, 0)	(0.24, 0.56, 0.2)	(0.18, 0.42, 0.4)	(0.12, 0.28, 0.6)	(0.06, 0.14, 0.8)	(0, 0, 1)
3	(0.4, 0.6, 0)	(0.32, 0.48, 0.2)	(0.24, 0.36, 0.4)	(0.16, 0.24, 0.6)	(0.08, 0.12, 0.8)	(0, 0, 1)
4	(0.5, 0.5, 0)	(0.4, 0.4, 0.2)	(0.3, 0.3, 0.4)	(0.2, 0.2, 0.6)	(0.1, 0.1, 0.8)	(0, 0, 1)
5	(0.6, 0.4, 0)	(0.48, 0.32, 0.2)	(0.36, 0.24, 0.4)	(0.24, 0.16, 0.6)	(0.12, 0.08, 0.8)	(0, 0, 1)
6	(0.7, 0.3, 0)	(0.56, 0.24, 0.2)	(0.42, 0.18, 0.4)	(0.28, 0.12, 0.6)	(0.14, 0.06, 0.8)	(0, 0, 1)
7	(0.8, 0.2, 0)	(0.64, 0.16, 0.2)	(0.48, 0.12, 0.4)	(0.32, 0.08, 0.6)	(0.16, 0.04, 0.8)	(0, 0, 1)
8	(0.2, 0.8, 0)	(0.16, 0.68, 0.16)	(0.12, 0.56, 0.32)	(0.08, 0.44, 0.48)	(0.04, 0.32, 0.64)	(0, 0.2, 0.8)
9	(0.3, 0.7, 0)	(0.24, 0.6, 0.16)	(0.18, 0.5, 0.32)	(0.12, 0.4, 0.48)	(0.06, 0.3, 0.64)	(0, 0.2, 0.8)
10	(0.4, 0.6, 0)	(0.32, 0.52, 0.16)	(0.24, 0.44, 0.32)	(0.16, 0.36, 0.48)	(0.08, 0.28, 0.64)	(0, 0.2, 0.8)
11	(0.5, 0.5, 0)	(0.4, 0.44, 0.16)	(0.3, 0.38, 0.32)	(0.2, 0.32, 0.48)	(0.1, 0.26, 0.64)	(0, 0.2, 0.8)
12	(0.6, 0.4, 0)	(0.48, 0.36, 0.16)	(0.36, 0.32, 0.32)	(0.24, 0.28, 0.48)	(0.12, 0.24, 0.64)	(0, 0.2, 0.8)
13	(0.4, 0.6, 0)	(0.36, 0.48, 0.16)	(0.32, 0.36, 0.32)	(0.28, 0.24, 0.48)	(0.24, 0.12, 0.64)	(0.2, 0, 0.8)
14	(0.5, 0.5, 0)	(0.44, 0.4, 0.16)	(0.38, 0.3, 0.32)	(0.32, 0.2, 0.48)	(0.26, 0.1, 0.64)	(0.2, 0, 0.8)
15	(0.6, 0.4, 0)	(0.52, 0.32, 0.16)	(0.44, 0.24, 0.32)	(0.36, 0.16, 0.48)	(0.28, 0.08, 0.64)	(0.2, 0, 0.8)
16	(0.7, 0.3, 0)	(0.6, 0.24, 0.16)	(0.5, 0.18, 0.32)	(0.4, 0.12, 0.48)	(0.3, 0.06, 0.64)	(0.2, 0, 0.8)
17	(0.8, 0.2, 0)	(0.68, 0.16, 0.16)	(0.56, 0.12, 0.32)	(0.44, 0.08, 0.48)	(0.32, 0.04, 0.64)	(0.2, 0, 0.8)
18	(0.2, 0.8, 0)	(0.16, 0.72, 0.12)	(0.12, 0.64, 0.24)	(0.08, 0.56, 0.36)	(0.04, 0.48, 0.48)	(0, 0.4, 0.6)
19	(0.3, 0.7, 0)	(0.24, 0.64, 0.12)	(0.18, 0.58, 0.24)	(0.12, 0.52, 0.36)	(0.06, 0.46, 0.48)	(0, 0.4, 0.6)
20	(0.4, 0.6, 0)	(0.32, 0.56, 0.12)	(0.24, 0.52, 0.24)	(0.16, 0.48, 0.36)	(0.08, 0.44, 0.48)	(0, 0.4, 0.6)
21	(0.5, 0.5, 0)	(0.4, 0.48, 0.12)	(0.3, 0.46, 0.24)	(0.2, 0.44, 0.36)	(0.1, 0.42, 0.48)	(0, 0.4, 0.6)
22	(0.5, 0.5, 0)	(0.48, 0.4, 0.12)	(0.46, 0.3, 0.24)	(0.44, 0.2, 0.36)	(0.42, 0.1, 0.48)	(0.4, 0, 0.6)
23	(0.6, 0.4, 0)	(0.56, 0.32, 0.12)	(0.52, 0.24, 0.24)	(0.48, 0.16, 0.36)	(0.44, 0.08, 0.48)	(0.4, 0, 0.6)
24	(0.7, 0.3, 0)	(0.64, 0.24, 0.12)	(0.58, 0.18, 0.24)	(0.52, 0.12, 0.36)	(0.46, 0.06, 0.48)	(0.4, 0, 0.6)
25	(0.8, 0.2, 0)	(0.72, 0.16, 0.12)	(0.64, 0.12, 0.24)	(0.56, 0.08, 0.36)	(0.48, 0.04, 0.48)	(0.4, 0, 0.6)

Each lower envelope lottery listed above can be constructed as the convex combination $\delta R + (1 - \delta)Y$, where R represents the lower envelope lottery at the “Risky End” of the budget, and Y represents the lower envelope lottery at the “Ambiguous End” of the budget.