Econometrica, Vol. 91, No. 4 (July, 2023), 1393–1416

TESTING HURWICZ EXPECTED UTILITY

HAN BLEICHRODT
Department of Economics (FAE), University of Alicante

SIMON GRANT
Research School of Economics, Australian National University

JINGNI YANG
Research School of Economics, Australian National University

Gul and Pesendorfer (2015) propose a promising theory of decision under uncertainty, they dub Hurwicz expected utility (HEU). HEU is a special case of α-maxmin EU that allows for preferences over sources of uncertainty. It is consistent with most of the available empirical evidence on decision under risk and uncertainty. We show that HEU is also tractable and can readily be measured and tested. We do this by deriving a new two-parameter functional form for the probability weighting function, which fits our data well and which offers a clean separation between ambiguity perception and ambiguity aversion. In two experiments, we find support for HEU's predictions that ambiguity aversion is constant across sources of uncertainty and that ambiguity aversion and first order risk aversion are positively correlated.

KEYWORDS: Hurwicz expected utility, probability weighting functions, source preferences.

1. INTRODUCTION

THE MODAL CHOICES first proposed by Ellsberg (1961) show a dislike of ambiguity and an inconsistency with beliefs that can be represented by a (unique) subjective probability measure.[1] While Ellsberg's paradox suggests ambiguity aversion, subsequent research has found that ambiguity attitudes are richer than simply aversion. For gains on unlikely events and for losses, people are generally ambiguity seeking (Kocher, Lahno, and Trautmann (2018)). Moreover, many experiments have shown that they have preferences over the sources of uncertainty that go beyond a preference for risk (objective probabilistic information) over uncertainty (vaguely specified probabilities). People prefer, for example, to bet on more familiar sources (French and Poterba (1991), Abdellaoui, Baillon, Placido, and Wakker (2011)) and on sources about which they consider themselves competent (Heath and Tversky (1991)). A rich variety of models has been developed to account for Ellsberg-type behavior (see Gilboa and Marinacci (2016) for a review). These models tend to concentrate on ambiguity aversion, and with a few exceptions, do not allow for ambiguity seeking, let alone source preference. Moreover, most of these models

Han Bleichrodt: hanbleichrodt@ua.es
Simon Grant: simon.grant@anu.edu.au
Jingni Yang: jingni.yang@anu.edu.au

We would like to thank four anonymous reviewers for their helpful comments on a previous version of this paper. We also benefit from discussions with Aurélien Baillon and Peter P. Wakker. We would also like to thank Evan Calford and Emma Schultz for making it possible for us to run the additional sessions during Covid. Han Bleichrodt acknowledges financial support from the Generalitat Valenciana (Prometeo/2021/073). Open access publishing facilitated by Australian National University (ANU), as part of the Wiley-Australian National University (ANU) agreement via the Council of Australian University Librarians.

[1]See Trautmann and Van De Kuilen (2015) for a review.

are defined in the Anscombe–Aumann framework and cannot account for the observed violations of expected utility for decision under risk.

The evidence about which of these ambiguity models best describes people's preferences is surprisingly scarce. The main reason for this is that most ambiguity models use concepts that are hard to measure. The studies that compared (a subset of) the ambiguity models concluded that models that allow for both ambiguity seeking and ambiguity aversion such as α-maxmin expected utility (Hurwicz (1951), Luce and Raiffa (2012), Ghirardato, Maccheroni, and Marinacci (2004)), Choquet expected utility (Schmeidler (1989)), and prospect theory (Tversky and Kahneman (1992)), perform better than models that account only for ambiguity aversion (Baillon and Bleichrodt (2015)). Moreover, source-dependent models perform better than models that do not account for source preference (Chew, Miao, and Zhong (2017)).

Gul and Pesendorfer (2015) provide behavioral foundations for a model, that they call Hurwicz expected utility (HEU), which can account for many of the observed empirical deviations from subjective expected utility. HEU is a special case of α-maxmin where the set of priors is generated from extensions of a single prior μ defined on a σ-algebra of ideal events. HEU can address Ellsberg-style evidence as well as the joint occurrence of ambiguity aversion and ambiguity seeking. For a particular source of ambiguity, which (formally) corresponds to a probability measure defined on a specific σ-algebra of events, HEU admits a rank-dependent utility representation (Quiggin (1982)). Consequently, it can also accommodate the main deviations from expected utility that have been observed for decision under risk. Moreover, while in many applications of rank-dependent utility an ad hoc specification is used for the probability weighting function, Gul and Pesendorfer prove that the HEU probability weighting function for a given source is a convex combination of a power series and its dual. Because these power series vary across sources, HEU allows for source preference. In other words, HEU is the first axiomatic theory that is consistent with most of the available empirical evidence on decision under risk and uncertainty.

In this paper, we show that HEU can also be made tractable. The main problem in measuring HEU is that the probability weighting function has an infinite number of parameters. We restrict the power series, while still satisfying all requirements of HEU, so that the probability weighting function is characterized by two parameters, one reflecting ambiguity perception, the other ambiguity aversion. HEU can then be estimated and tested. Hence, an interesting byproduct of our paper is to propose a new functional form of the weighting function, which is axiomatically derived (to the extent that it follows from axioms of HEU). We show that this *HEU weighting function* fits our data at least as well as widely used alternatives like those of Prelec (1998) and Goldstein and Einhorn (1987) with the additional advantage that the parameters have a clear interpretation.

We test two predictions of HEU: first, that the parameter reflecting ambiguity aversion is constant across sources, and second, that ambiguity aversion and first-order risk aversion are positively correlated. First order risk aversion is a central concept in insurance theory and determines the extent to which people are willing to accept actuarially unfair insurance. It can also be used to distinguish decision theories. To the best of our knowledge, we are the first to quantify first-order risk aversion.

We found support for HEU's predictions in our main experiment where we measured HEU's parameters for three distinct sources of uncertainty (Canberra temperature, Paris temperature, and attendance of the football match FC Barcelona-Atletic Bilbao) and first- order risk aversion. We could not reject the null that the α's were the same in the aggregate analysis and for about 60% of the subjects their α's were close and HEU described their preferences well in the individual analysis. The correlation between α and our measure of first-order risk aversion was moderate and highly significant. The second experiment avoided parametric assumptions, but as a result could not fully separate ambiguity perception and ambiguity aversion. It confirmed that ambiguity attitudes were the same across sources, but the positive correlation between ambiguity aversion and first-order risk aversion was insignificant. This was due to the low number of subjects who were first order risk averse in this study and the confounding impact of ambiguity perception. It illustrates the importance of separating ambiguity perception and ambiguity aversion, which is what our HEU weighting function achieves.

2. HURWICZ EXPECTED UTILITY

Let $\mathbf{X}$ be a set of consequences and Ω an (infinite) state space, with generic elements x and ω, respectively. We refer to subsets of $E \subset \Omega$ as *events*. An *act* f is a function from Ω to $\mathbf{X}$.

A *prior* is a countably additive, complete, and nonatomic probability measure on some σ-algebra of events. Let Π denote the set of all priors, and for a given $\pi \in \Pi$, let ε_π denote the σ-algebra on which π is defined.

A prior π' constitutes an *extension* of another prior π, if $\varepsilon_{\pi'} \supset \varepsilon_\pi$ and $\pi'(E) = \pi(E)$ for all $E \in \varepsilon_\pi$. Let Π_π be the set of priors that are extensions of π. In other words, the set Π_π consists of those priors that agree with π on its domain.

We consider a decision maker who has a preference relation $\succsim$ over $\mathcal{F}$, the set of acts. In *Hurwicz Expected Utility (HEU)*, the decision-maker's preferences can be characterized by a prior $\mu \in \Pi$, an ambiguity index $\alpha \in [0, 1]$, and a (Bernoulli) utility function $v : X \to \mathbb{R}$, such that $\succsim$ admits the representation:

$$W(f) = \alpha \min_{\pi \in \Pi_\mu} \int v \circ f \, d\pi + (1 - \alpha) \max_{\pi \in \Pi_\mu} \int v \circ f \, d\pi. \tag{1}$$

Both μ and α are unique and v is unique up to unit and scale.

Gul and Pesendorfer (2015) refer to the elements of $\mathcal{E}_\mu$ as *ideal events*. The decision maker sees these events as the least uncertain. Both an ideal event and its complement satisfy Savage (1954) postulate **P2** (*the sure-thing principle*). Equation (1) shows that if an act f is ε_μ-measurable, which Gul and Pesendorfer (2015) call an *ideal act*, then HEU reduces to subjective expected utility and the ranking between any pair of such acts does not depend on α. Thus the utility function v embodies the decision-maker's risk attitudes with respect to ideal acts.

For any act f that is not ideal, decision makers apply a Hurwicz criterion to the set of extensions of μ. In other words, they act as if they are completely ignorant about probabilities beyond those implied by μ. The parameter α reflects the decision-maker's attitude toward ambiguity with higher values of α indicating more aversion to ambiguity (Gul and Pesendorfer (2015, Proposition 3)). If $\alpha = 1$, then the decision maker is always ambiguity averse. If $0 < \alpha < 1$, then the decision maker exhibits both ambiguity seeking and ambiguity aversion. If $\alpha = 0$, then the decision maker is always ambiguity seeking. HEU is a special case of α-maxmin for the given set of priors Π_μ.

The events in $\mathcal{E}_\mu$ are not necessarily the ones for which probabilities are known (decision under risk). In HEU, sources are subjective and the decision maker may not see a process generating known probabilities as the least uncertain. Moreover, different processes generating known probabilities may be seen as different sources. This is consistent with the evidence in Armantier and Treich (2016), that risk cannot (always) be treated as a single source of uncertainty.

A *bet* $x_A y$ is a binary act yielding outcome x if event A obtains and y otherwise. Throughout, we assume that writing $x_A y$ implies that $x \succsim y$. Let

$$\mu_*(B) = \sup_{E \in \mathcal{E}_\mu, E \subset B} \mu(E) \tag{2}$$

be the inner probability of an event B. For a bet $x_A y$, the least favorable extension of μ assigns $\mu_*(A)$ to A while the most favorable extension assigns $1 - \mu_*(A^c)$ to A.

Putting equation (2) into the HEU formula, equation (1) gives

$$\begin{aligned} W(x_A y) = {} & \alpha\big[\mu_*(A)v(x) + \big(1 - \mu_*(A)\big)v(y)\big] \\ & + (1-\alpha)\big[\big(1 - \mu_*\big(A^c\big)\big)v(x) + \mu_*\big(A^c\big)v(y)\big]. \end{aligned} \tag{3}$$

As equation (3) shows, the ranking of bets generally depends on the decision-maker's ambiguity aversion, reflected by α. For example, in the Ellsberg two-color example, an ambiguity averse decision maker prefers drawing a red ball from the known urn to drawing a red ball from the unknown urn whereas an ambiguity seeking decision maker has the reverse preference. Two events come from the same source if such preference reversals cannot occur: the ranking of the bets depends only on the decision-maker's prior and not on their ambiguity attitude. Formally, a prior $\pi \in \Pi$ (and its associated σ-algebra ε_π) is a *source* if for every (α, ν), $A, B \in \varepsilon_\pi$, $x_A y \succsim x_B y$ if and only if $\pi(A) \geq \pi(B)$.[2] So, in Ellsberg's 2-color problem drawing a ball from the known urn and drawing a ball from the unknown urn are different sources, while we expect that decision makers are indifferent between betting on red and betting on black when a ball is drawn from the unknown urn regardless of their ambiguity attitude.

The definition of a source π implies that for all $A, B \in \varepsilon_\pi$, $x_A y \sim x_B y$ if and only if $\pi(A) = \pi(B)$. If this condition holds, the decision maker considers A and B equally likely. Gul and Pesendorfer (2015) define sources in terms of bets for convenience, but the definition can be extended to any simple act. If $(E_1, \dots, E_n)$ is a finite partition of the state space with all $E_i \in \varepsilon_\pi$, then two events E_i and E_j are equally likely if exchanging the outcomes under events E_i and E_j does not change the indifference class. Or, more formally, if $x_{E_i} y_{E_j} f$ denotes the change of the act $f \in \mathcal{F}$ in which the outcome for any state in the event E_i is replaced by x and the outcome for any state in the event E_j by y, then E_i and E_j are viewed as equally likely if $x_{E_i} y_{E_j} f \sim y_{E_i} x_{E_j} f$ for all pairs of outcomes x and y. This corresponds to the definition of exchangeability proposed by Chew and Sagi (2008) and tested (and supported) by Abdellaoui et al. (2011).

Proposition 2 in Gul and Pesendorfer (2015) shows that every source is characterized by a power series. That is, for each source π there exists a sequence $(a_1, a_2, \dots, a_n, \dots)$ with $a_n \in [0, 1]$, $\sum_{n=1}^{\infty} a_n = 1$, such that for all $A \in \varepsilon_\pi$,

$$\mu_*(A) = \gamma\big(\pi(A)\big) = \sum_{n=1}^{\infty} a_n \big(\pi(A)\big)^n. \tag{4}$$

Let γ_π be the power series of source π. Substituting equation (4) into equation (3) gives

$$\begin{aligned} W(x_A y) = {} & \alpha\big[\gamma_\pi(\pi\big((A)\big)v(x) + (1 - \gamma_\pi\big(\pi\big((A)\big)\big)v(y)\big] \\ & + (1-\alpha)\big[(\gamma_\pi\big(1 - \pi\big((A)\big)\big)v(x) + (1 - \gamma_\pi\big(1 - \pi\big((A)\big)\big)v(y)\big]. \end{aligned} \tag{5}$$

[2] In addition, Gul and Pesendorfer (2015) impose that if $A_n \in \varepsilon_\pi$ and $x_{A_n} y$ converges pointwise to $x_A y$ then $W(x_A y)$ is the limit of the sequence $W(x_{A_n} y)$.

If we replace in equation (5) event A by an equally likely source event B, that is, $\pi(A) = \pi(B)$, then $W(x_A y) = W(x_B y)$. The value of the bets depends only on the likelihoods of the events and not on ambiguity attitudes. Hence, probabilistic sophistication holds within sources. Because the a_i in a source's power series are nonnegative (and at least one is strictly positive), equation (5) is strictly increasing in $\pi(A)$ and all decision makers prefer to bet on more likely events.

Proposition 4 in Gul and Pesendorfer (2015) shows that the power series characterizes the ambiguity perception of decision makers. It states that one source π is more uncertain than another source π^0, if γ_π, the power series associated with π is dominated by γ_{π^0}, the power series associated with π^0. That is, $\gamma_\pi(p) \le \gamma_{\pi^0}(p)$ for all p.

Because decision makers behave probabilistically sophisticated (Machina and Schmeidler (1992)) for sources, they use the prior to reduce acts to "objectively" risky lotteries. So, for a given source π, any act f for which $f^{-1}(x) \in \varepsilon_\pi$ for all outcomes x, is evaluated as the lottery $P(x) = \pi(f^{-1}(x))$. Under HEU (Gul and Pesendorfer (2015), Proposition 6), these lotteries are evaluated by rank-dependent utility:

$$V(P) = \sum_{z \in \mathbb{R}} \left(w\left(\sum_{x: v(x) \ge z} P(x) \right) - w\left(\sum_{v(x) > z} P(x) \right) \right) \cdot v(z), \tag{6}$$

where v is the (Bernoulli) utility index from equation (1) and w is the probability weighting function given by

$$w(p) = \alpha\gamma(p) + (1-\alpha)\hat{\gamma}(p). \tag{7}$$

In equation (7), α is the decision maker's ambiguity aversion parameter from equation (1), γ is the power series associated with the source π, and $\hat{\gamma}$ is the dual of γ: $\hat{\gamma}(p) = 1 - \gamma(1-p)$. The right-hand side of equation (6) reduces to expected utility when $\gamma(p) = p$, that is, when $w(p)$ is the identity function.

Equations (6) and (7) show that utility and ambiguity aversion are the same for all sources. However, the probability weighting function is source-dependent, since it depends on the decision-maker's perception of ambiguity, γ, which can vary across sources.

As the power series γ is convex, its dual $\hat{\gamma}$ is concave, and hence, each probability weighting function is a weighted average of a convex and a concave function. Therefore, it can allow for the inverse S-shape that is usually observed in empirical studies of probability weighting (Gonzalez and Wu (1999), Abdellaoui (2000), Bleichrodt and Pinto (2000), Fox and Poldrack (2014)).

Segal and Spivak (1990) showed that probability weighting is related to first order risk aversion, the tendency to buy full insurance at actuarially unfair prices. More convex probability weighting (which follows from a higher value of α in HEU) leads to a more pronounced kink toward the origin in decision-makers' indifference curves and, consequently, to more first order risk aversion. In particular, for any pair of HEU decision makers with the same ambiguity perception (for a given source) and the same utility function, the more ambiguity averse one (the one with higher α) will exhibit more first order risk aversion. In other words, HEU predicts a positive correlation between α and first order risk aversion, ceteris paribus.

The HEU Weighting Function. We test these two predictions of HEU: that the ambiguity aversion parameter α is source-independent and that α is positively correlated with first order risk aversion. To do so, we need to measure HEU. The crucial issue to do so is the shape of the power series in equation (4). Since equation (4) can have an infinite number of positive terms, but experimental data are finite, we have to restrict the number of terms. We used the following parametric function to describe the power series:

$$\gamma_\beta(p) = \frac{(1-\beta)p}{(1+\beta)(1-p)+(1-\beta)p}. \tag{8}$$

The following lemma show that γ_β in equation (8) is indeed a power series.

LEMMA 1: *The coefficients of the Taylor series of* $\gamma_\beta(x) = \frac{(1-\beta)x}{(1+\beta)(1-x)+(1-\beta)x}$ *are always positive and sum to* 1.

The proof of Lemma 1 is in the Appendix. Lemma 1 shows that equation (8) is indeed a suitable power series for HEU. The function γ_β is decreasing in β, so larger values of β imply that a source is perceived as more ambiguous. Equations (7) and (8) jointly determine the following two-parameter probability weighting function:

$$w(p) = \frac{\alpha(1-\beta)p}{(1+\beta)(1-p)+(1-\beta)p} + \frac{(1-\alpha)(1+\beta)p}{(1-\beta)(1-p)+(1+\beta)p}. \tag{9}$$

In equation (9), the parameter α reflects ambiguity aversion and the parameter β ambiguity perception. Hence, this two-parameter weighting function has an intuitive interpretation. This is an advantage over other weighting functions used in the literature. For example, in Prelec's (1998) two-parameter weighting function, the interpretation of the two parameters is less clear and they interact. In what follows, we will refer to equation (9) as the *HEU weighting function*.

3. EXPERIMENTS

In two experiments, we tested whether the ambiguity aversion parameter α is source-independent and whether it is positively correlated with first order risk aversion. In the main experiment, the *parametric* experiment, we used a parametric method to estimate α. In the second experiment, the *nonparametric* experiment, we tested the robustness of the findings from the first experiment by measuring α by a nonparametric method, which however could not fully separate ambiguity aversion and ambiguity perception.

Prior to the experiments, we performed two pilot studies to fine-tune our experimental designs. First, we measured HEU using the data of Abdellaoui et al. (2011), which included as sources risk, Ellsberg urns, home and foreign temperature, and the stock exchange. This showed that our method to measure the parameters α and β indeed worked and that the HEU weighting function fitted well. The data in Abdellaoui et al. (2011) did not allow us to measure first order risk aversion. We, therefore, performed a second pilot study to explore how we could best measure first order risk aversion. Here, we used the tasks of Loomes and Segal (1994) and derived a measure of first order risk aversion from these. We describe these pilot studies and their results in the Online Supplementary Material (Bleichrodt, Grant, and Yang (2023)) in the Appendix.

The experiments combined the insights from these pilot studies. The first experiment measured α using the HEU weighting function and tested whether it was equal across the sources and positively correlated with first order risk aversion as measured by the test that performed best in our pilot study of first order risk aversion. In the second experiment, we did not use the (parametric) HEU weighting function, but measured α using a nonparametric method, and again tested whether it was constant across sources and correlated with first order risk aversion. The instructions of both experiments are in the Online Appendix.

3.1. *The Parametric Experiment*

3.1.1. *Subjects and Incentives*

We ran the experiment between the 16th and the 24th of February 2022. Subjects were 48 students from the subject pool of the Australian National University (ANU) Behavioral Econ Lab coming from different academic backgrounds. Due to Covid restrictions, we ran the experiment through Zoom with a maximum of 4 subjects per online session. Subjects were asked to put their camera on so we could monitor them during the experiment. They could invite the interviewer to a break out room or send private messages to the interviewer if they had questions. Subjects received a showup fee of 20 AUD (approximately 14 USD). In addition, we randomly selected one subject who played out one of their randomly selected choices for real. Subjects took on average 40 minutes to complete the experiment.

3.1.2. *Procedure*

Measuring Beliefs. The experiment consisted of 3 parts to measure HEU for 3 different sources of uncertainty and a fourth part in which we measured subjects' degree of first order risk aversion. For each part, we started with an explanation and two comprehension questions. After they had answered these correctly, subjects moved on to a practice question. After this practice question, the actual experiment started.

The sources of uncertainty were the maximum temperature in Canberra on the 5th of May 2022, the maximum temperature in Paris on the 5th of May 2022 and the attendance of the football match FC Barcelona against Athletic de Bilbao on the 27th of February 2022. We selected this match because Barcelona has the largest stadium in Europe (capacity 99,354), Europe was in the middle of the omicron crisis and it was unclear how this might affect stadium capacity, Bilbao is a good team but not top, and in previous seasons the match was never sold out, and Barcelona was going through a rough spell after Lionel Messi, one of the best players in history and seven times winner of the golden ball for best football player in the world (the last time in 2021) had left and their coach, Ronald Koeman, had just been fired.[3]

We did not include risk as a source. According to HEU, if a DM perceives a source as "objectively" risky, with known probabilities, then they should evaluate it by *SEU* and α is not identifiable. If the source risk is not perceived as "objectively" risky, then there is no particular reason to consider it separately. We performed an overall test of whether α is constant across sources, on both our data and the data of the experiments of Abdellaoui et al. (2011), which include risk (see the Online Appendix). The results were consistent with those presented below.

For each source, we elicited events with probabilities $\frac{1}{8}$, $\frac{1}{4}$, $\frac{1}{2}$, $\frac{3}{4}$, and $\frac{7}{8}$ using the exchangeability method of Abdellaoui et al. (2011). We explain the procedure for the attendance of the football match. The procedures for Canberra and Paris temperature were similar. We started by splitting the state space into two complementary events with probability $\frac{1}{2}$. Figure B.1 in the Appendix shows how we asked the questions. Subjects were

[3]We told subjects the attendance of Barcelona–Bilbao in the previous 10 seasons.

presented with two binary acts. In Option A, they win AUD1000 if attendance is at most x (and nothing otherwise), in Option B they win AUD1000 if attendance is more than x (and nothing otherwise). We varied x until subjects were indifferent between A and B. Let $a_{\frac{1}{2}}$ denote the attendance for which this indifference holds and let $E_{\frac{1}{2}} = [0, a_{\frac{1}{2}}]$ denote the event that the attendance is at most $a_{\frac{1}{2}}$. Assuming that match attendance is a source π_f (we check this later), we get by equation (5) that $\pi_f(E_{\frac{1}{2}}) = \pi_f(E^c_{\frac{1}{2}}) = \frac{1}{2}$.

We then split the event $E_{\frac{1}{2}}$ into two equally likely events by finding the attendance $a_{\frac{1}{4}}$ such that subjects were indifferent between winning AUD1000 if the attendance was between 0 and $a_{\frac{1}{4}}$ and winning AUD1000 if the attendance was between $a_{\frac{1}{4}}$ and $a_{\frac{1}{2}}$. It follows that the events $E_{\frac{1}{4}} = [0, a_{\frac{1}{4}}]$ and $(a_{\frac{1}{4}}, a_{\frac{1}{2}}]$ have subjective probability $\frac{1}{4}$.

To find the event $E_{\frac{3}{4}}$ with probability $\frac{3}{4}$, we split the event $E^c_{\frac{1}{2}}$ into two equally likely events by finding the attendance $a_{\frac{3}{4}}$ for which subjects were indifferent between betting on $(a_{\frac{1}{2}}, a_{\frac{3}{4}}]$ and betting on $(a_{\frac{3}{4}}, 99{,}354]$. As both of these events have probability $\frac{1}{4}$, the event $E_{\frac{3}{4}} = [0, a_{\frac{3}{4}}]$ has probability $\frac{3}{4}$. Finally, to determine the events $E_{\frac{1}{8}}$ and $E_{\frac{7}{8}}$ with probabilities $\frac{1}{8}$ and $\frac{7}{8}$, we split the events $E_{\frac{1}{4}}$ and $E^c_{\frac{3}{4}}$ into equally likely events.

The above procedure implies that if we split the event $[a_{\frac{1}{4}}, a_{\frac{3}{4}}]$ into two equally likely events $[a_{\frac{1}{4}}, a'_{\frac{1}{2}}]$ and $(a'_{\frac{1}{2}}, a_{\frac{3}{4}}]$, then we should find that $a'_{\frac{1}{2}} = a_{\frac{1}{2}}$ (except for random error). This was our test that exchangeability holds and that match attendance is indeed a source. We used the same test to check whether Canberra temperature and Paris temperature were sources. Subjects first faced the questions for Paris temperature, then for match attendance, and finally for Canberra temperature. Within the sources, the order of elicitation was $a_{\frac{1}{2}}, a_{\frac{1}{4}}, a_{\frac{3}{4}}, a_{\frac{1}{8}}, a_{\frac{7}{8}}, a'_{\frac{1}{2}}$.

All indifferences were elicited by a bisection process. The possible values for temperature in Paris and Canberra varied between -10 and 40 degrees Celsius. Possible match attendance varied between 0 and 99,354. The procedures for the starting values and the step size in the bisection process were the same as those in Abdellaoui et al. (2011). For any interval (a,b], we always started with a choice in which the interval was split in winning events $(a, a/3 + 2b/3]$ and $(2b/3, b]$ and then one in which it was split in winning events $(a, 2a/3 + b/3]$ and $(2a/3 + b/3, b]$. After these choices, we used the usual bisection procedure where in each choice the step size was halved. So, for example, to determine $E_{\frac{1}{2}}$ for match attendance, in the first choice subjects chose between betting on $[0, 67k]$ and betting on $(67k, 100k]$ and in the second choice between betting on $[0, 33k]$ and betting on $(33k, 100k]$. Depending on their choices, the bisection process then started.[4]

Measuring Utility and Probability Weighting. Let $x_p y$ denote the *risky* act that gives x with (known) probability p and y with probability $1 - p$. The *certainty equivalent* of an act $x_E y$ or $x_p y$ is the amount c for sure that is equivalent to that act: $c \sim x_E y$ or $c \sim x_p y$. To determine the utility function, we measured the certainty equivalents of the five risky acts $500_{\frac{1}{2}}0$, $1000_{\frac{1}{2}}500$, $500_{\frac{1}{2}}250$, $700_{\frac{1}{2}}500$, and $1000_{\frac{1}{2}}750$. HEU assumes that utility is source-independent, and thus, the same for risk and uncertainty. Abdellaoui et al. (2011) found support for this assumption (see also Abdellaoui, Bleichrodt, l'Haridon, and Van Dolder

[4]So, if the subject had chosen to bet twice on Option B (the one with the higher attendance) then in the next choice they chose between betting on $[0, 82K]$ and betting on $(82K, 100K]$. If they now chose Option A, the next choice was between betting on $[0, 74K]$ and betting on $(74K, 100K]$. The process ended when the highest attendance (temperature) for which A was chosen and the lowest attendance (temperature) for which B was chosen was less than 2000 (1 degree). We then took the midpoint of these values as the indifference value.

(2016)). We estimated utility by a power function $U(x) = x^{\rho}$, which is widely used in empirical research and typically gives a good fit (e.g., Wakker (2008)).

To measure the probability weights, we elicited for each source the certainty equivalents of five acts 1000_E0 with the events E equal to $E_{\frac{1}{2}}, E_{\frac{1}{4}}, E_{\frac{3}{4}}, E_{\frac{1}{8}}, E_{\frac{7}{8}}$. Using equation (6) and the estimated power utility function, we obtain the probability weights, which we can then use to estimate α and β in the HEU weighting function equation (9).[5]

We determined the certainty equivalents by a bisection procedure where the starting value was the expected value of the binary act. The process ended when the difference between the lowest amount for which the sure amount was chosen and the highest amount for which the act was chosen was less than AUD 10. We then took the midpoint of these values as the indifference value. Subjects first faced the five certainty equivalents questions for Paris temperature then for match attendance, then for risk, and finally for Canberra temperature.

Measuring First Order Risk Aversion. To measure first order risk aversion, we used the two-color question from Loomes and Segal (1994).[6] Unlike Loomes and Segal (1994), we also needed to quantify first-order risk aversion to compute the correlation between ambiguity aversion, as measured by α, and first-order risk aversion. Subjects faced an urn with blue and orange chips in known proportions $0.5 + \varepsilon$ and $0.5 - \varepsilon$ with $\varepsilon \in [0, 0.5]$. They had to allocate 17.30 AUD between the events "draw a blue chip" and "draw an orange chip." The amount 17.30 AUD was the same as in Loomes and Segal (1994) and served to counteract responding in round numbers.

We elicited the maximal value of ε ($\varepsilon_{\max}$), for which subjects chose an equal division between blue and orange (8.65 AUD on both). The value of $\varepsilon_{\max}$ was taken as a measure of first-order risk aversion. The larger is $\varepsilon_{\max}$, the more first-order risk averse subjects are. Loomes and Segal (1994) show that an expected utility maximizer with differentiable utility will choose $\varepsilon_{\max}$ equal to zero, whereas a rank-dependent utility maximizer (as in HEU) will have $\varepsilon_{\max} > 0$.

We elicited $\varepsilon_{\max}$ through bisection. In the first allocation, ε was set equal to $p/2 = 0.25$. Subjects saw a colored bar on their screen of which $(0.5 + \varepsilon)\%$ was blue and $(0.5 - \varepsilon)\%$ was orange. The numbers $0.5 + \varepsilon$ and $0.5 - \varepsilon$ were also displayed to leave no doubt what the respective probabilities were.

Subjects first saw a screen (see Figure B.2) where the amount they could put on blue[7] increased in steps of 2 AUD from 6.65 AUD to 16.65 AUD (and, consequently, the amount they could put on orange decreased from 10.65 AUD to 0.65 AUD). The allocation with 6.65 AUD on blue tested for comprehension as this allocation was first order stochastically dominated by the allocation in which 10.65 AUD was put on blue.

The bisection process varied ε until we found the maximum value ε^* for which subjects still preferred an equal allocation between blue and orange and the minimum value ε_* for which they preferred an unequal allocation between blue and orange. The difference between ε^* and ε_* was at most 2%.

Subjects then saw a second screen (see Figure B.3) where the amount they could put on blue increased in steps of 0.40 AUD from 8.25 AUD (again a first order stochastically dominated option) to 9.85 AUD. The bisection process again zoomed in at the maximum

[5]We estimated α and β by nonlinear least squares.

[6]Loomes and Segal (1994) also used two three color questions. We tested these in the pilot, but they led to less consistent results.

[7]In the printing version, dark gray refers to blue chips and light gray refers to orange chips.

value ε^* for which subjects still preferred an equal allocation between blue and orange and the minimum value ε_* for which they preferred an unequal allocation between blue and orange. The difference between these two values was at most 1% and we took their midpoint as the value of $\varepsilon_{\max}$.

3.1.3. *Results*

Tests of Exchangeability. Figure 1 shows the elicited values of $a_{\frac{1}{2}}$ and $a'_{\frac{1}{2}}$. They were close for most subjects (the dotted line shows equality), suggesting that exchangeability held. The exceptions are subjects 12 and 16 (and subjects 5 and 9 for Paris temperature). Excluding these two subjects did not alter our conclusions. Their estimated α's varied widely, violating HEU, and so keeping them in puts the bar somewhat higher for HEU.

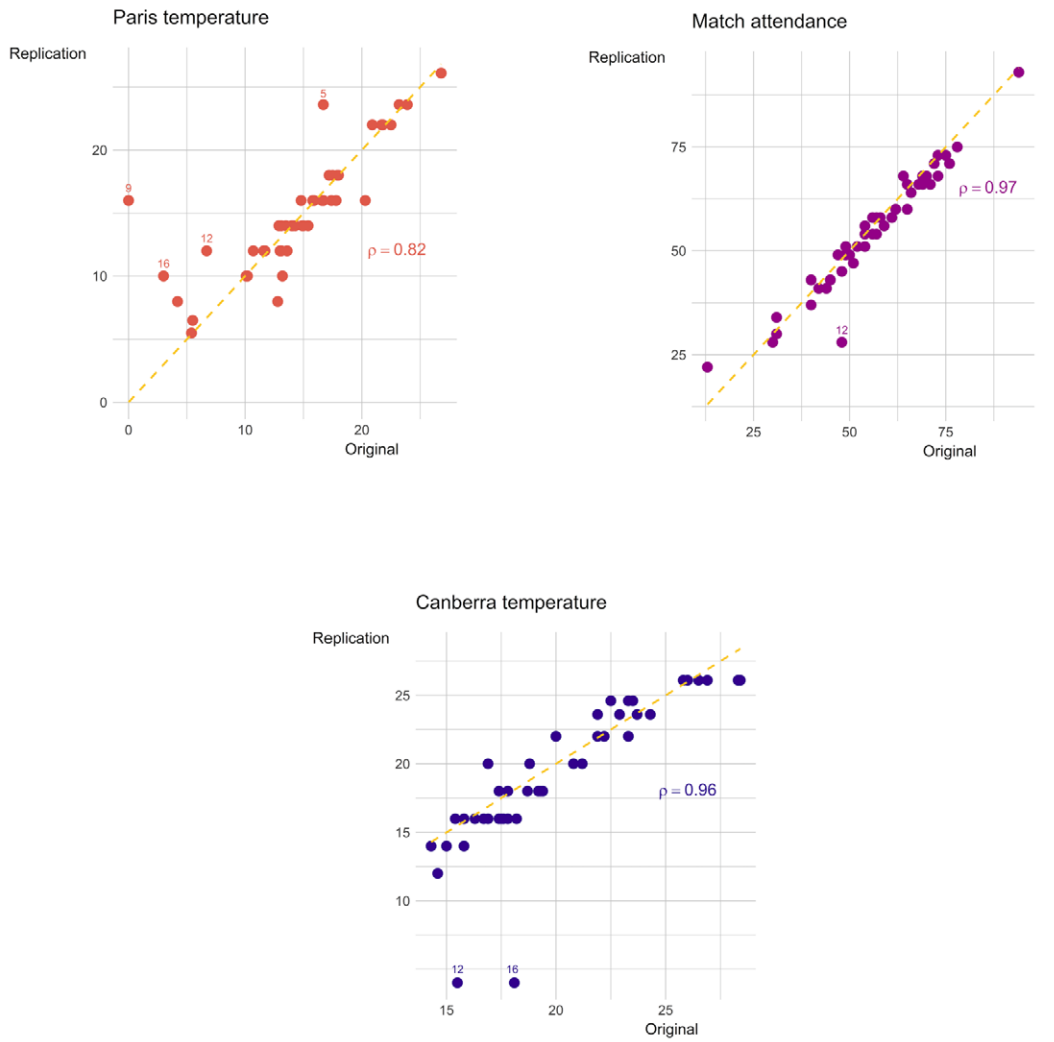

FIGURE 1.—Tests whether exchangeability holds and $a_{\frac{1}{2}}$ and $a'_{\frac{1}{2}}$ are equal.

Statistical tests confirmed that exchangeability held and Paris temperature, Canberra temperature, and match attendance can be taken as sources. We could not reject that $a_{\frac{1}{2}}$ and $a'_{\frac{1}{2}}$ were equal (ANOVA, $p = 0.79$).[8] Moreover, as shown in Figure 1, the correlations between $a_{\frac{1}{2}}$ and $a'_{\frac{1}{2}}$ were almost perfect: they exceeded 0.82 for all sources.[9] This is comparable to the correlations in Abdellaoui et al. (2011).

In the measurement of first order risk aversion, three subjects violated first order stochastic dominance (FOSD), two of which violated FOSD in every choice. Excluding these three subjects did not affect our results. Their implied value of $\varepsilon_{\max}$ was 0, that is, consistent with expected utility. As we will explain below, HEU makes no prediction about the correlation between α and first-order risk aversion for expected utility maximizers and so expected utility maximizers had to be excluded from the correlation analysis anyhow.

The median power coefficient of the utility function was 0.97 showing that aggregate utility was approximately linear (t-test, $p = 0.72$). Abdellaoui et al. (2011) also could not reject the null that utility was linear. Apparently, utility does not contribute much to explaining risk aversion once source preference, ambiguity aversion, and probability weighting are taken into account.

The HEU Weighting Function. Figure 2 shows that the HEU function fitted the data well for each of the three sources. This also held for most of our subjects individually. These estimations are in the Online Appendix.

Table I compares the fit of the HEU weighting function with that of popular alternative parametric forms: the weighting functions of Goldstein and Einhorn (1987), Tversky and Kahneman (1992), Prelec (1998), and the neo-additive weighting function of Chateauneuf, Eichberger, and Grant (2007).[10] We used the same symbols for the parameters as in the HEU weighting function to highlight similarities even though, apart from the neo-additive weighting function, the other functional forms do not give the same clean separation between ambiguity perception and ambiguity aversion as the HEU weighting function. We included Tversky Kahneman's weighting function because of its widespread use, even though we expected it to fit worse as it has only one parameter.

The table shows that HEU, Prelec, and GE fit about equally well as measured by their sums of squared residuals (SSR). Indeed, a two-way ANOVA with functional form and source as factors found no significant differences between these three parametric forms ($p = 0.91$). Neo-additive fitted somewhat worse, but adding the neo-additive function to the ANOVA still gave no significant difference ($p = 0.44$). However, when adding the Tversky Kahneman weighting function, Tukey's HSD test indicated that it fitted clearly worse ($p < 0.001$).[11]

We also determined for each subject separately which function fitted best (in terms of SSR). Here, HEU, Prelec, and neo-additive performed best. The main thing to note from Table I is that the HEU function fits as well as other widely used forms, which gives support for our parametric analysis presented next.

---

[8] We checked the robustness of all of our findings using Bayesian statistics. This led to the same conclusions.

[9] In describing the strength of correlations, we use the scheme of Landis, Richard, and Koch (1977).

[10] We follow the convention in the literature to restrict the parameters of the Prelec and Goldstein Einhorn (GE) weighting functions to satisfy subproportionality (Kahneman and Tversky (1979)). The analysis with the parameters unrestricted led to the same conclusions.

[11] Using the data of Abdellaoui et al. (2011) led to the same conclusions: the fit of HEU, Prelec, and GE is similar and better than neo-additive and Tversky Kahneman. See the Online Appendix for details.

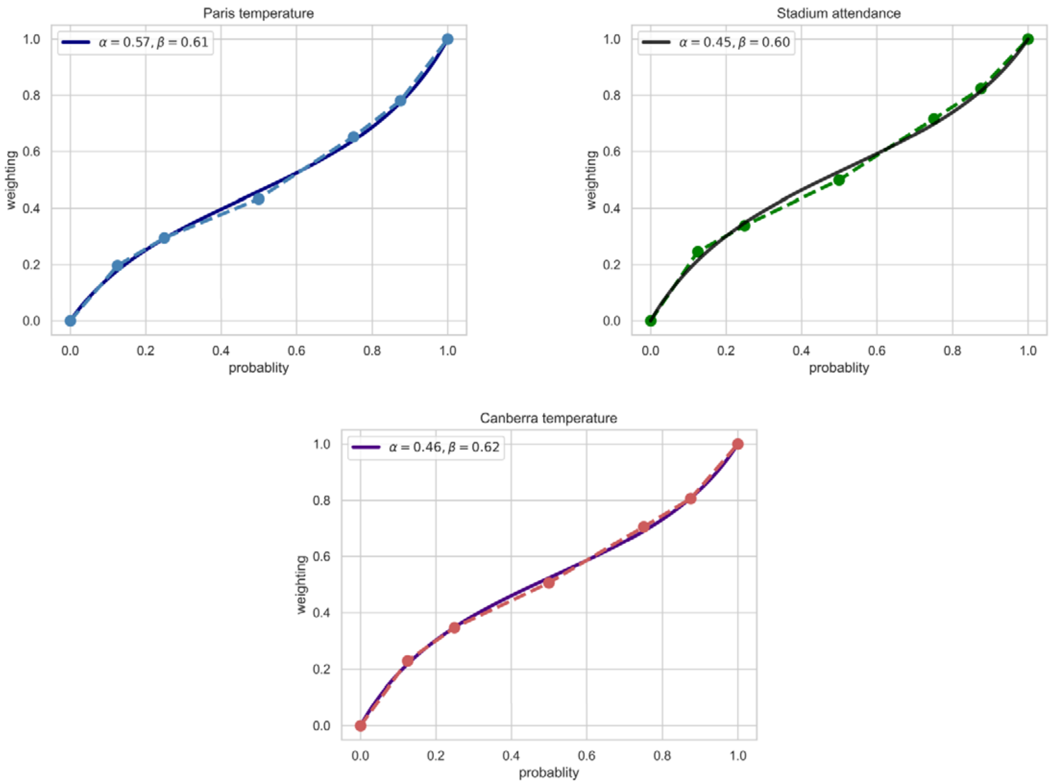

FIGURE 2.—HEU weighting functions for the three sources of uncertainty by mean data.

First Test: Are the α's Source-Independent? Our data support the first prediction of HEU, that the α's are equal across the three sources of uncertainty. Figure 2 shows that the overall values of α were close. Figure 3 shows that the distributions of the elicited

TABLE I

THE FIT OF THE HEU WEIGHTING FUNCTION IN COMPARISON WITH WIDELY USED ALTERNATIVES.

		Sources					
		Paris		Match		Canberra	
Functional Form	Expression	Sum of SSRs	No. Subjects	Sum of SSRs	No. Subjects	Sum of SSRs	No. Subjects
HEU	Equation (9)	1.55	21	2.06	10	0.95	12
Goldstein Einhorn	$\frac{\alpha p^\beta}{\alpha p^\beta+(1-p)^\beta}$	1.52	5	2.01	6	0.95	7
Tversky Kahneman	$\frac{p^\beta}{(p^\beta+(1-p)^\beta)^{\frac{1}{\beta}}}$	6.67	0	8.41	0	7.52	0
Prelec	$\exp(-\alpha(-\ln p)^\beta)$	1.42	14	1.87	20	0.91	21
Neo-additive	$\begin{cases} 0, & p=0, \\ \beta(1-\alpha)+(1-\beta)p, & p\in(0,1), \\ 1, & p=1 \end{cases}$	2.06	13	2.43	13	1.35	9

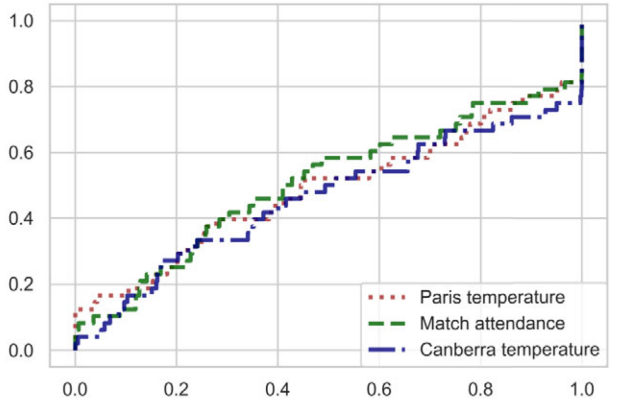

FIGURE 3.—Cumulative distribution functions of α for the three sources of uncertainty.

α's were also close and an ANOVA test could not reject the null that they are equal ($p = 0.57$). We have the same result using the data of Abdellaoui et al. (2011), where they have four sources including risk (see the Online Appendix for details). Figure 4 shows the correlations between the α's for the three sources. They are moderate to substantial and vary from 0.56 (Paris and Canberra temperature) to 0.68 (Paris temperature and match attendance) to 0.76 (Canberra temperature and match attendance). Correlation tests showed that they are all significantly different from zero ($p < 0.01$ in all cases). The figure shows that for most subjects their α's were indeed approximately equal.

To get a clearer picture of how many subjects behaved in line with HEU and had source-independent ambiguity attitude, we computed for each subject the maximum difference between their α's. Figure 5 shows this distribution. For around 60% of the subjects, the maximum difference was less than 0.20, which is about half the standard deviation of the distribution of α's that we observed both in our experiment and when using the data of Abdellaoui et al. (2011).[12] There are clearly some subjects for whom the α's differed substantially, but for a majority HEU's assumption of source-independent ambiguity attitudes describes their preferences well.

Figure 6 shows that this good individual fit did not occur because subjects perceived the sources as equally ambiguous. The figure shows that there was no relation between the difference in α and the difference in β. The correlation was 0.02 and was not significantly different from zero ($p = 0.91$). This analysis also shows that the parameters α and β in the HEU weighting function are independent and measure different concepts. We see this as an important advantage of the HEU weighting function.

Second Test: Are Ambiguity Aversion and First Order Risk Aversion Correlated? The mean of the index of first order risk aversion was 3.11, which by a t-test was significantly different from 0 ($p < 0.01$). Thus, on average subjects did not behave according to expected utility. Seventeen subjects (35.4%) had their index equal to zero. This was close to the 37.2% who behaved in line with expected utility in the pilot.[13]

[12]The proportion of subjects who behaved according to HEU was similar in the analysis of the data in Abdellaoui et al. (2011). See the Online Appendix for details.

[13]In Loomes and Segal (1994), 24.3% of the subjects behaved in line with expected utility. The difference might be explained by the larger number (16) of tests in their study.

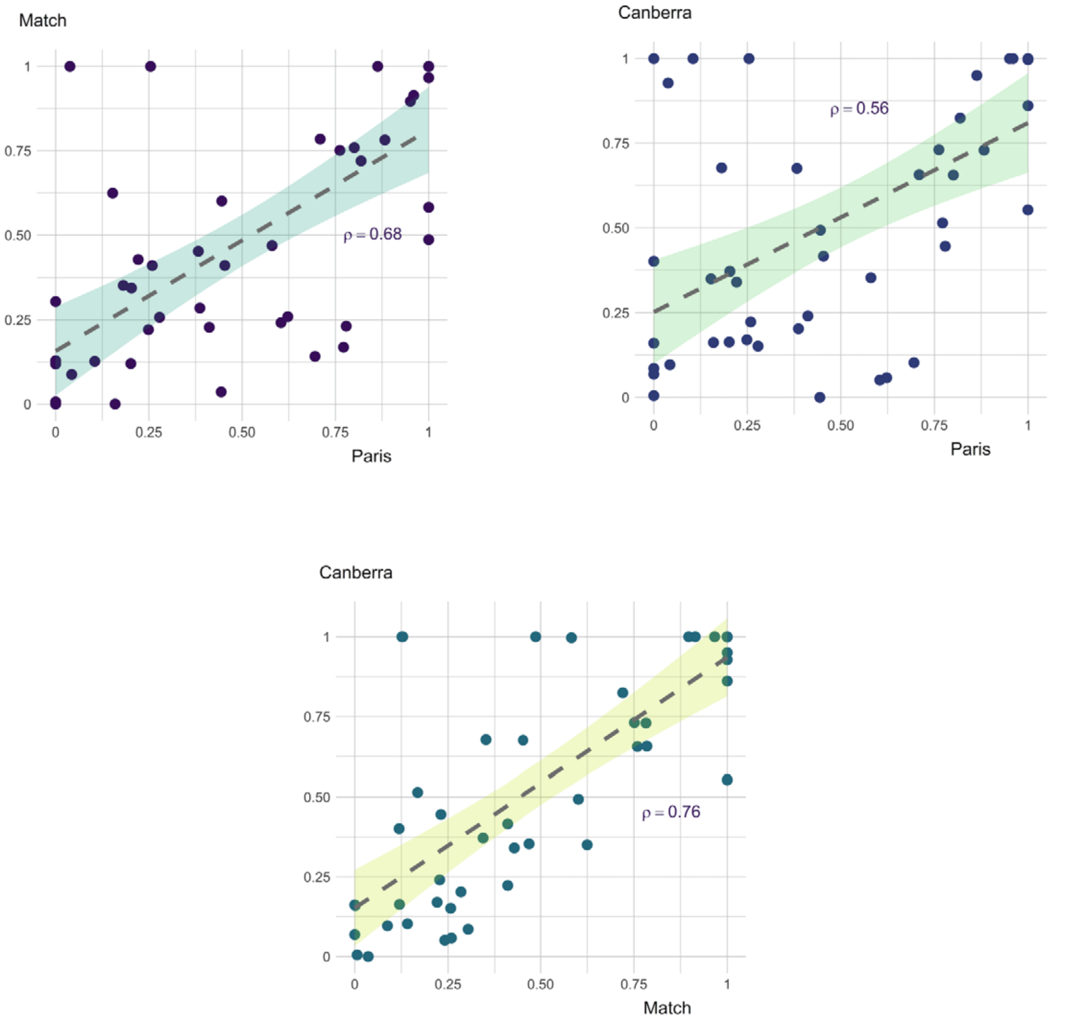

FIGURE 4.—Correlations between the measurements of the α's in the first experiment.

Under HEU, subjects who are first order risk neutral for a source, view all events in that source as "ideal" and are expected utility maximizers for acts measurable for that source. That is, $\gamma(p) = p$ and it follows from equation (7) that regardless of α, $w(p) = p$. Hence, for expected utility maximizers HEU predicts no relation between ambiguity aversion and first-order risk aversion and we, therefore, removed these subjects from the analysis about the correlation between ambiguity aversion and first order risk aversion.

Figure 7 shows the correlation between ambiguity aversion (measured by the mean value of α across the three sources) and our measure of first order risk aversion for the subjects who were first order risk averse. In line with the predictions of HEU, there is a moderate positive correlation ($\rho = 0.51$, correlation test $p < 0.01$).[14]

[14]In the Online Appendix, we show the correlations between α and first-order risk aversion for each source separately.

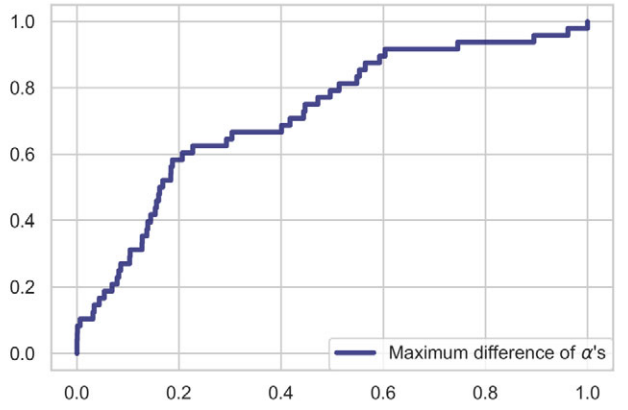

FIGURE 5.—Cumulative distribution function of the maximum difference between α's for the three sources of uncertainty.

3.1.4. *Conclusion First Experiment and Prologue to the Second*

Our first experiment confirms both predictions of HEU. We found no significant differences between the α's that we elicited for the three different sources and we found a positive correlation between α and first order risk aversion. We also found that the HEU probability weighting function fitted at least as well as widely-used alternatives in the literature (while the parameters have a clear interpretation).

The first experiment made parametric assumptions about probability weighting and the utility function. We wanted to explore in a second experiment whether we could replicate our results without these assumptions. We designed a new measurement of α that did not make any parametric assumption and tested again whether the α's were source-independent and positively correlated with first order risk aversion.

The *matching probability* $m(E)$ of an event E is defined as the probability p such that $x_p y \sim x_E y$. Dimmock, Kouwenberg, and Wakker (2016, Theorem 3.1) showed how to measure ambiguity attitudes using matching probabilities without the need to know utility. Their result is also valid under HEU. In the second experiment, we measure am-

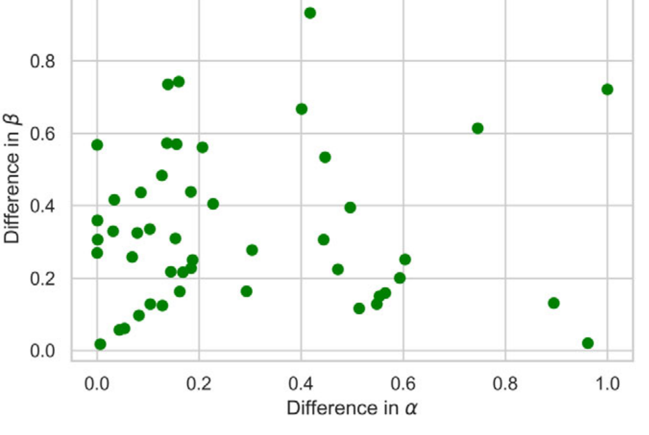

FIGURE 6.—The relation between the difference in α and the difference in β.

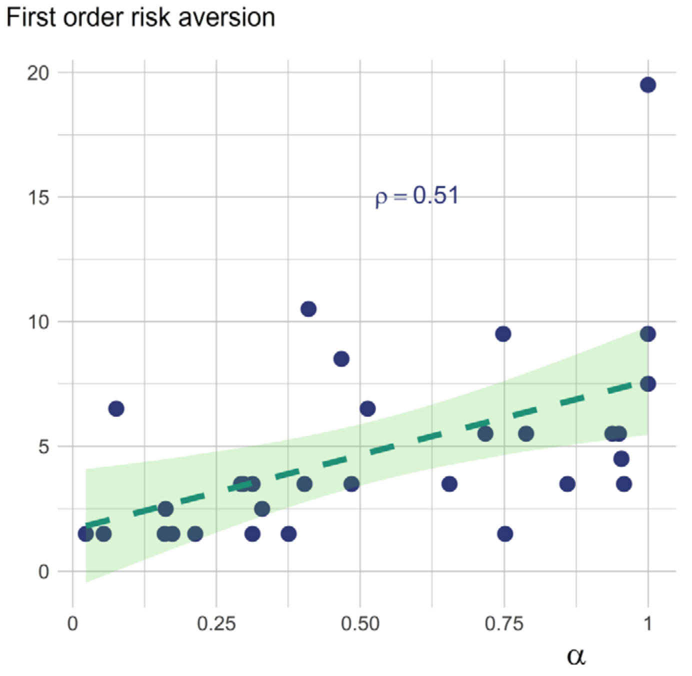

FIGURE 7.—Positive correlation between ambiguity aversion (α) and first order risk aversion..

biguity aversion as the difference between $\frac{1}{2}$ and the matching probability of an event with subjective probability $\frac{1}{2}$. Our method shows how ambiguity attitudes can be measured using matching probabilities while controlling for beliefs. Baillon, Huang, Selim, and Wakker (2018) showed how this can be done under the neo-additive weighting function. Our method extends this to general weighting functions. It measures the matching probabilities of the event and its complement for which subjects have belief $\frac{1}{2}$ and by induction, it can be used for other values of beliefs.

A limitation of this nonparametric test is that we cannot separate the effects of ambiguity aversion and ambiguity perception as cleanly as we did in the first experiment. However, the first experiment showed that at the aggregate level the differences in ambiguity perception between the sources were modest and we, therefore, used them again in the second experiment (except for adjustments for dates, see below).[15]

3.2. *The Nonparametric Experiment*

3.2.1. *Subjects and Incentives*

We ran the experiment between the 1st and the 7th of June 2022. Subjects were 34 students drawn from the ANU Behavioral Econ Lab subject pool. We ran the experiment in person in sessions with at most 4 subjects per session.[16] The incentives were the same

[15]Moreover, it can be shown that if subjects are expected utility for risk, then matching probabilities for an event with subjective probability $\frac{1}{2}$, as we use in our second experiment, maximally reflect changes in α.

[16]Recruiting subjects proved difficult, because even though Covid restrictions were gradually lifted at the time, we performed the experiment, ANU still did almost everything online, and most students avoided campus.

as in the first experiment: subjects received a showup fee of 20 AUD (approximately 14 USD) and, in addition, we randomly selected one subject to play out one of their randomly selected choices for real. Subjects took on average 35 minutes to complete the experiment.

3.2.2. *Procedure*

Measuring α Nonparametrically. The experiment consisted of two parts. First, we determined for the three sources of uncertainty the events with subjective probability $\frac{1}{2}$ and their corresponding matching probabilities. Second, we measured subjects' first-order risk aversion using the same method as in the first experiment. The sources were the maximum temperature in Canberra on the 4th of September 2022, the maximum temperature in Paris on the 4th of September 2022, and the attendance at the first home match of FC Barcelona in the 2022–2023 season.[17] For both parts, we started with an explanation and two comprehension questions. After they had answered these correctly, subjects moved on to a practice question. After this practice question, the actual experiment started.

To determine the events with subjective probabilities $\frac{1}{2}$, we selected an event and elicited both its matching probability and that of its complement. For example, for match attendance we started by determining the matching probabilities of the events that at most 30,000 people would attend the first home match of FC Barcelona in the 2022/2023 season and that at least 30,000 people would attend this match.[18] If the matching probabilities differed, then we adjusted the events until the matching probabilities were the same. The exact elicitation procedure is described in the Online Appendix.

Under HEU, if the elicited matching probabilities for an event E and its complement E^c are the same, they must be equally likely, and thus, $P(E) = P(E^c) = \frac{1}{2}$. However, by equation (5) the elicited matching probabilities can be different from $\frac{1}{2}$ due to ambiguity attitudes. We took the difference between $\frac{1}{2}$ and the matching probability as our measure of ambiguity aversion. As mentioned above, this difference is also affected by ambiguity perception and, consequently, is not a perfect measure, but that is the price to pay for trying to measure α without making parametric assumptions.

3.2.3. *Results*

First Test: Are the α's Source-Independent? Figure 8 shows the correlations between the measurements of α for the three sources. The correlations are lower than in the first experiment, but we expected this given the confounding impact of ambiguity perception. They vary from moderate (Paris–Canberra and Paris–Match) to substantial (Match–Canberra). An ANOVA could not reject the null that the α's were equal ($p = 0.44$). The data from the second experiment confirm the support for HEU's prediction of source-independent ambiguity aversion that we found in the first experiment.

Second Test: Are Ambiguity Aversion and First Order Risk Aversion Correlated? Three subjects violated first-order stochastic dominance. As these responses probably reflect confusion, we removed them from the remaining analyses. This left 31 subjects in the

[17] It was unknown at the time of the experiment when this match would be and who would be the opponent. We told subjects what the attendance had been during Barcelona's first home match for the past 10 years and who the opponent had been. The actual opponent turned out to be Rayo Vallecano.

[18] For Paris temperature, we started with the events "the temperature on the 4th of September 2022 is at most 20.5 degrees celsius" and "the maximum temperature on the 4th of September 2022 is at least 20.5 degrees celsius. For Canberra temperature, we used the events "the temperature on the 4th of September 2022 is at most 10.5 degrees celsius" and "the temperature on the 4th of September 2022 is at least 10.5 degrees celsius.

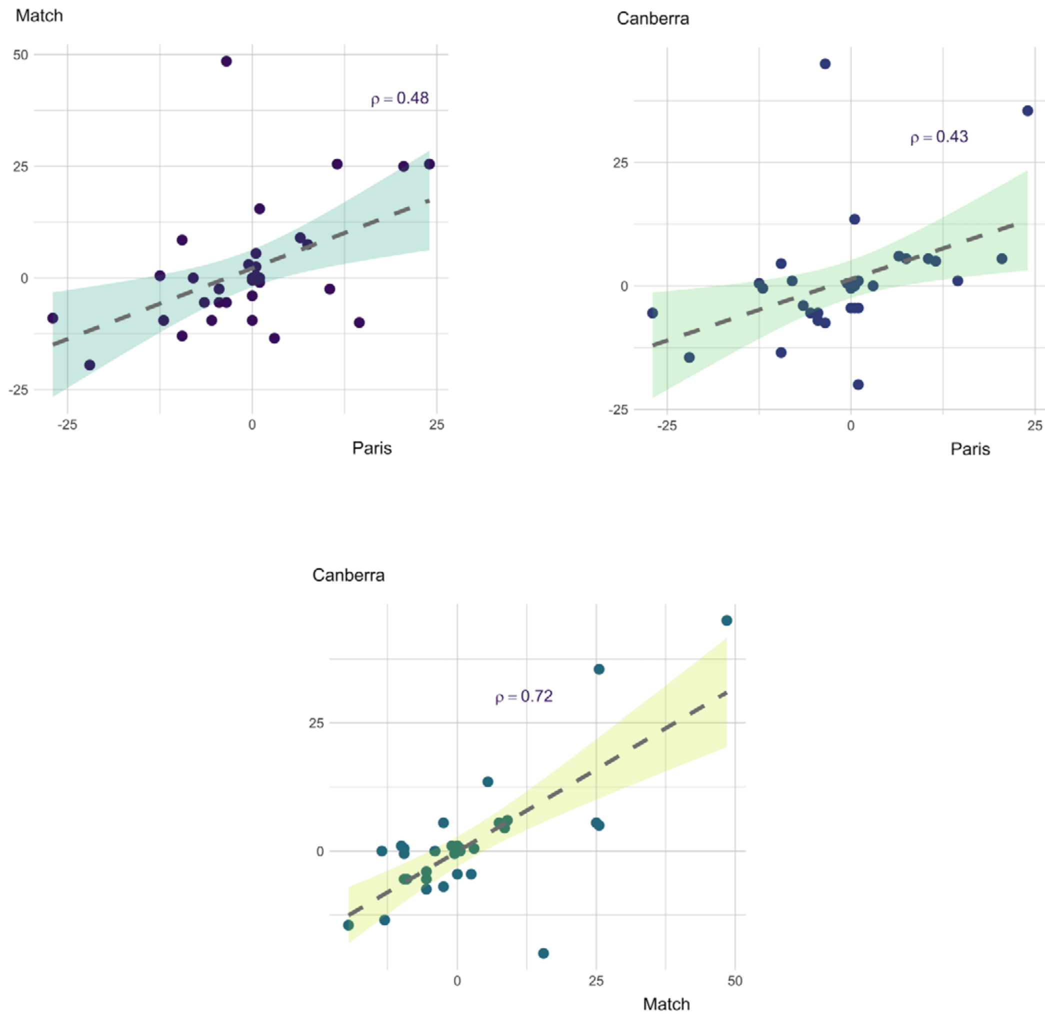

FIGURE 8.—Correlations between the measurements of the α's in the second experiment.

analyses. The mean of our measure of first order risk aversion was 2.24 and differed from 0, the prediction of expected utility (t-test, $p < 0.001$). 38.7% of the subjects had their index equal to zero, which is about the same as in the first experiment and in the pilot. As HEU makes no prediction for these subjects about their correlation between ambiguity aversion and first-order risk aversion, we removed them from the subsequent analysis.

Figure 9 shows the relation between the measures of ambiguity aversion and first order risk aversion. The overall correlation was slight and by a correlation test insignificant ($p = 0.65$). The lower correlation compared with the first experiment is not surprising given that ambiguity aversion was measured with some imprecision and there were only 19 subjects with an index of first-order risk aversion exceeding 0.[19] In the pilot on first-order risk aversion, we measured ambiguity aversion using an Ellsberg urn task. Here, ambiguity perception also interfered and we observed no significant positive correlation

[19]A Bayesian test indicated that the data were inconclusive.

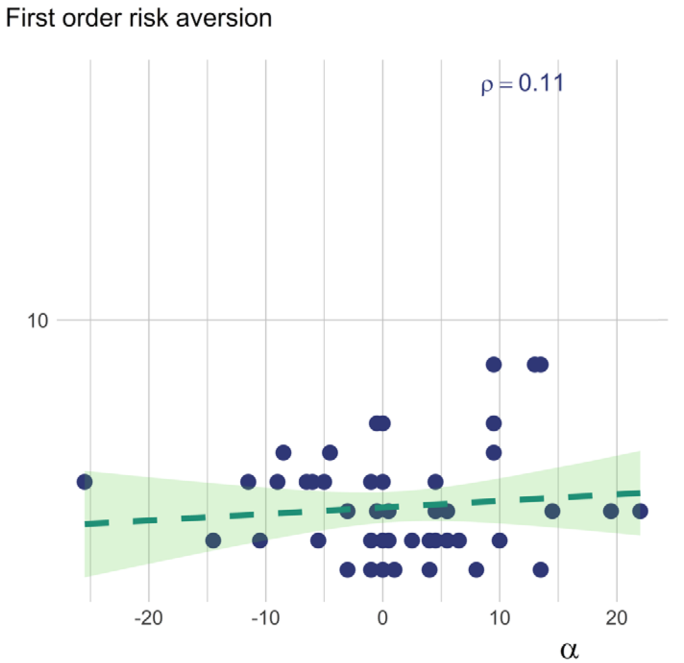

FIGURE 9.—Positive correlation between ambiguity aversion (α) and first order risk aversion.

between ambiguity aversion and first-order risk aversion. See the Online Appendix for details. The data from experiment 2 and the pilot indicate that the separation between ambiguity aversion and ambiguity perception that we realized in our main experiment is important to test HEU properly.

4. CONCLUSION

Since the Schmeidler (1989) classic paper, the literature has produced a wide variety of models to account for ambiguity attitudes. Gul and Pesendorfer (2015) HEU is a promising theoretical model that can accommodate many known behavioral phenomena. It accounts for the joint occurrence of ambiguity aversion and ambiguity seeking and for source-dependence of ambiguity attitudes. For decision under risk, HEU is equivalent to rank-dependent utility. Consequently, it is consistent with all the deviations of expected utility that rank-dependent utility and prospect theory for gains can explain.

In this paper, we show an additional advantage of HEU: it can be measured empirically. One challenge of ambiguity models is that they are hard to measure and test, because they involve abstract concepts like sets of priors (Gilboa and Schmeidler (1989), Ghirardato, Maccheroni, and Marinacci (2004), Maccheroni, Marinacci, and Rustichini (2006)) or (subjective) second-order distributions (Klibanoff, Marinacci, and Mukerji (2005), Nau (2006)). This challenge also makes them hard to apply to real-world problems. HEU also uses an abstract concept (the ideal sets), but we show that this does not inhibit the possibility to measure it. The crucial step in our measurement of HEU was to operationalize the power series that characterizes each source of uncertainty. We show how this can be done by adding one parameter that measures ambiguity perception to the model. This leads to a new two-parameter specification of the probability weighting function, in which the parameters have a clear interpretation: one measures ambiguity aversion, the other ambiguity perception. We showed that the HEU weighting function fitted our data at least as well as popular alternatives. Because risk can be seen as a source of uncertainty on its own, the HEU weighting function can also be applied in decision under risk and is an attractive alternative for the weighting functions that are commonly used in empirical studies.

We tested two predictions of HEU: (i) that the ambiguity aversion parameter α is constant across sources of uncertainty, and (ii) that α is positively correlated with first-order risk aversion. To test the second prediction, we had to quantify first-order risk aversion, and proposed a measure to do so.

We observed support for both predictions. Both in our main experiment and in a robustness check that used a nonparametric measure of ambiguity aversion, we could not reject the null that the α's were the same across sources of uncertainty. Around 60% of our subjects behaved in line with HEU and had their α's close. For the remaining subjects, ambiguity aversion depended on the source. To model their behavior, we might need a more general α-maxmin model with α source-dependent. This, in turn, would be a special case of the model in Ghirardato, Maccheroni, and Marinacci (2004) where α can depend on the act.

We found that α and first order risk aversion were moderately positively correlated. In the robustness check, the correlation was slight and insignificant, but this was caused by the low number of first-order risk averse subjects and the fact that we could not fully separate ambiguity aversion and ambiguity perception in the nonparametric measurement. Our data suggest that to properly measure HEU, ambiguity aversion and ambiguity perception should be separated. As we showed in our main experiment, the HEU weighting function can achieve this. Capturing ambiguity perception by a single parameter may be ambitious, and sometimes a simplification. On the other hand, it offers substantial advantages for applying HEU and we found that the HEU weighting function fitted the data of most individual subjects well.

Ambiguity is rife and many real-world decisions are made in the face of uncertainty. In spite of this, ambiguity is often ignored in such decisions. One reason is that ambiguity models are hard to apply. The main contribution of our paper is to bridge the gap between the theoretical and applied literature. We show that a promising model, HEU, that captures the main findings from the empirical literature, can be measured, and thus applied. Moreover, we find support for its main predictions. We hope that in so doing our paper contributes to helping people and governments make better decisions in the face of ambiguity.

APPENDIX A: PROOFS

PROOF OF LEMMA 1: First, $\gamma_\beta(1) = 1$. Notice that

$$\gamma_\beta(x) = \frac{1}{\dfrac{1+\beta}{1-\beta}\left(\dfrac{1}{x} - 1\right) + 1}.$$

If we set

$$g_\beta(x) := \frac{1+\beta}{1-\beta}\left(\frac{1}{x}-1\right)+1\big(=1/\gamma_\beta(x)\big),$$

then we see that $g_\beta^{(n)}(x) = \frac{1+\beta}{1-\beta}(-1)^n n! x^{-(n+1)}$ and $g_\beta^{(n)}(x) = (-1)nx^{-1}g_\beta^{(n-1)}(x)$ for all $n \geq 2$.

Moreover, since $\gamma_\beta(x)g_\beta(x) = 1$, it follows that

$$\big(\gamma_\beta(x)g_\beta(x)\big)^{(n)} = 0 \quad \text{for all } n \geq 1, \tag{10}$$

By the rule of derivatives, (10) is equivalent to: for all $n \geq 1$,

$$\sum_{i=0}^{n}\binom{n}{i}\gamma_\beta^i(x)g_\beta^{n-i}(x) = \sum_{i=0}^{n}\frac{n!}{i!(n-i)!}\gamma_\beta^i(x)g_\beta^{n-i}(x) = 0. \tag{11}$$

For $i \geq 2$, since $g_\beta^{n-i}(x) = (-1)(n-i)x^{-1}g_\beta^{n-i-1}(x)$, we have

$$\frac{n!}{i!(n-i)!}f_\beta^i(x)g_\beta^{n-i}(x) = \frac{n!}{i!(n-i)!}\gamma_\beta^i(x)(-1)(n-i)x^{-1}g_\beta^{n-i-1}(x) \tag{12}$$

$$= \frac{-n}{x}\frac{(n-1)!}{i!(n-i-1)!}\gamma_\beta^i(x)g_\beta^{n-i-1}(x). \tag{13}$$

Substituting each term when $i > 2$ in (11) by (13), we have

$$\begin{aligned}
\sum_{i=0}^{n}\frac{n!}{i!(n-i)!}\gamma_\beta^i(x)g_\beta^{n-i}(x) &= \gamma_\beta^n(x)g_\beta(x) + nf_\beta^{n-1}(x)g_\beta^{(1)}(x) \\
&\quad + \frac{-n}{x}\sum_{i=0}^{n-2}\frac{(n-1)!}{i!(n-i-1)!}\gamma_\beta^i(x)g_\beta^{n-i-1}(x) \\
&= \gamma_\beta^n(x)g_\beta(x) + n\gamma_\beta^{n-1}(x)g_\beta^{(1)}(x) + \frac{-n}{x}\big(-\gamma_\beta^{n-1}(x)g_\beta(x)\big) \\
&= \gamma_\beta^n(x)g_\beta(x) + n\gamma_\beta^{n-1}(x)g_\beta^{(1)}(x) + \frac{n}{x}\big(\gamma_\beta^{n-1}(x)g_\beta(x)\big) \\
&= \gamma_\beta^n(x)g_\beta(x) + n\gamma_\beta^{n-1}(x)\left(g_\beta^{(1)}(x) + \frac{1}{x}g_\beta(x)\right) = 0.
\end{aligned}$$

It is straightforward to check that $g_\beta^{(1)}(x) + \frac{1}{x}g_\beta(x) < 0$, and so by induction if $\gamma_\beta^{n-1}(x) > 0$, then $\gamma_\beta^n(x) > 0$, as required. *Q.E.D.*

APPENDIX B: EXPERIMENTAL DISPLAYS

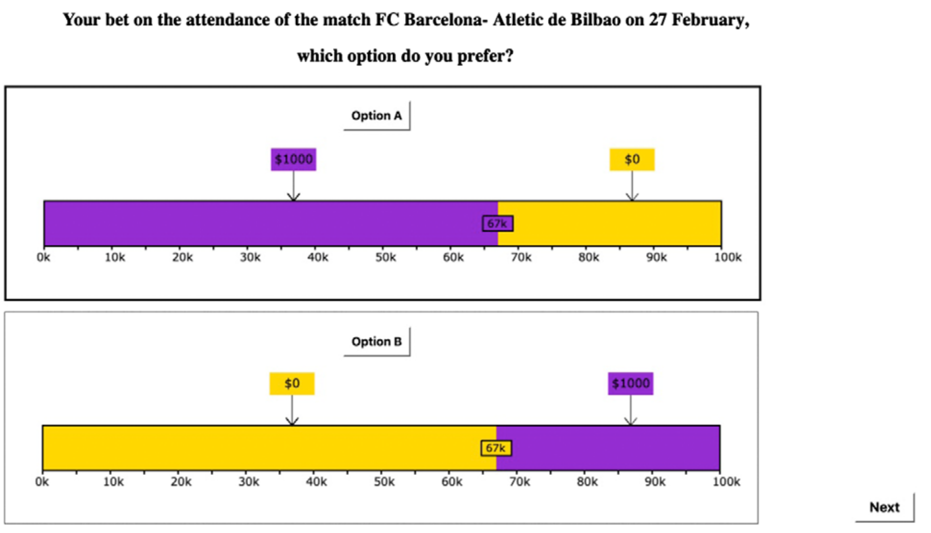

FIGURE B.1.—A screenshot of the experiment for exchangeable event of the match attendance.

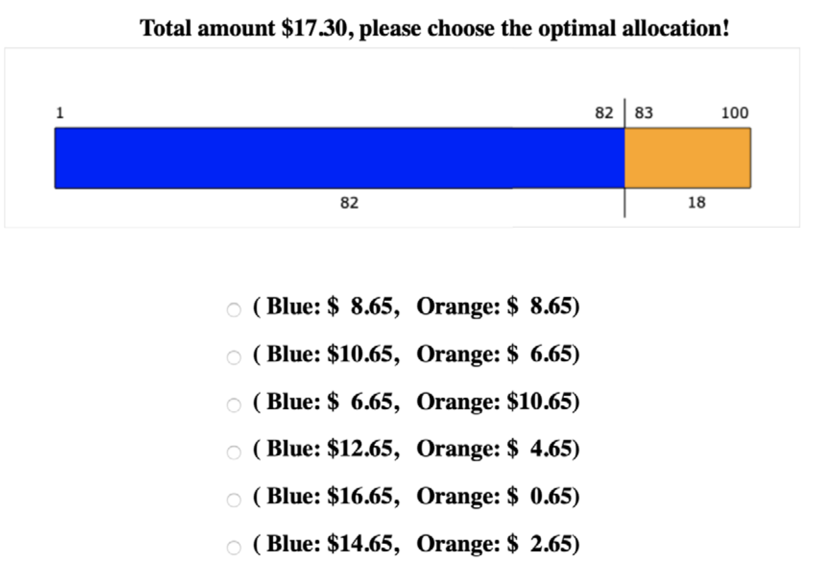

FIGURE B.2.—A screenshot of the experiment for first order risk aversion I.

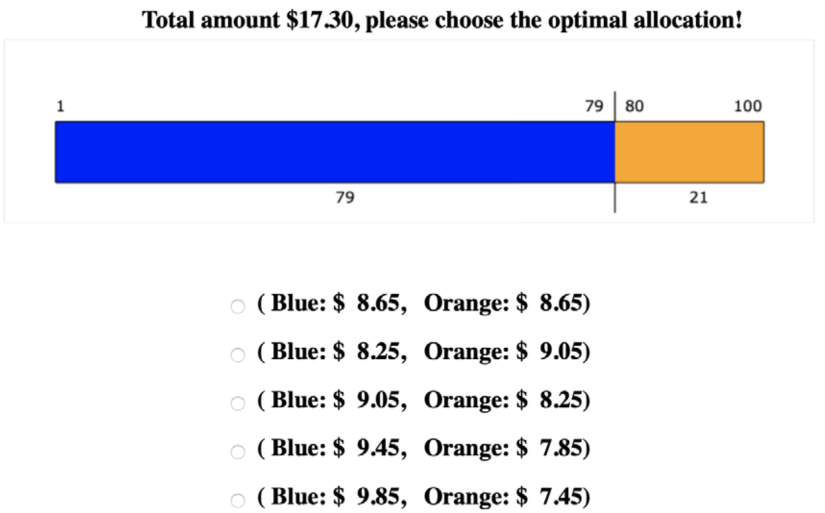

FIGURE B.3.—A screenshot of the experiment for first order risk aversion II.

REFERENCES

ABDELLAOUI, MOHAMMED (2000): "Parameter-Free Elicitation of Utility and Probability Weighting Functions," *Management Science*, 46 (11), 1497–1512. [1397]

ABDELLAOUI, MOHAMMED, AURÉLIEN BAILLON, LAETITIA PLACIDO, AND PETER P. WAKKER (2011): "The Rich Domain of Uncertainty: Source Functions and Their Experimental Implementation," *American Economic Review*, 101 (2), 695–723. [1393,1396,1398-1400,1403,1405]

ABDELLAOUI, MOHAMMED, HAN BLEICHRODT, OLIVIER L'HARIDON, AND DENNIE VAN DOLDER (2016): "Measuring Loss Aversion Under Ambiguity: A Method to Make Prospect Theory Completely Observable," *Journal of Risk and Uncertainty*, 52 (1), 1–20. [1400,1401]

ARMANTIER, OLIVIER, AND NICOLAS TREICH (2016): "The Rich Domain of Risk," *Management Science*, 62 (7), 1954–1969. [1395]

BAILLON, AURÉLIEN, AND HAN BLEICHRODT (2015): "Testing Ambiguity Models Through the Measurement of Probabilities for Gains and Losses," *American Economic Journal: Microeconomics*, 7 (2), 77–100. [1394]

BAILLON, AURÉLIEN, ZHENXING HUANG, ASLI SELIM, AND PETER WAKKER (2018): "Measuring Ambiguity Attitudes for All (Natural) Events" *Econometrica*, 86 (5), 1839–1858. [1408]

BLEICHRODT, HAN, AND JOSE LUIS PINTO (2000): "A Parameter-Free Elicitation of the Probability Weighting Function in Medical Decision Analysis," *Management Science*, 46 (11), 1485–1496. [1397]

BLEICHRODT, HAN, SIMON GRANT, AND JINGNI YANG (2023): "Supplement to 'Testing Hurwicz Expected Utility'," *Econometrica Supplemental Material*, 91, https://doi.org/10.3982/ECTA19221. [1398]

CHATEAUNEUF, ALAIN, JÜRGEN EICHBERGER, AND SIMON GRANT (2007): "Choice Under Uncertainty With the Best and Worst in Mind: Neo-Additive Capacities," *Journal of Economic Theory*, 137 (1), 538–567. [1403]

CHEW, SOO HONG, AND JACOB S. SAGI (2008): "Small Worlds: Modeling Attitudes Toward Sources of Uncertainty," *Journal of Economic Theory*, 139 (1), 1–24. [1396]

CHEW, SOO HONG, BIN MIAO, AND SONGFA ZHONG (2017): "Partial Ambiguity," *Econometrica*, 85 (4), 1239–1260. [1394]

DIMMOCK, STEPHEN G., ROY KOUWENBERG, AND PETER P. WAKKER (2016): "Ambiguity Attitudes in a Large Representative Sample," *Management Science*, 62 (5), 1363–1380. [1407]

ELLSBERG, DANIEL (1961): "Risk, Ambiguity, and the Savage Axioms," *The Quarterly Journal of Economics*, 75 (4), 643–669. [1393]

FOX, CRAIG R., AND RUSSELL A. POLDRACK (2014): "Prospect Theory and the Brain," in *Handbook of Neuroeconomics*, ed. by Glimcher Paul and Fehr Ernst, 533–567. [1397]

FRENCH, KENNETH R., AND JAMES M. POTERBA (1991): "Investor Diversification and International Equity Markets," Tech. rep., National Bureau of Economic Research. [1393]

GHIRARDATO, PAOLO, FABIO MACCHERONI, MASSIMO MARINACCI (2004): "Differentiating Ambiguity and Ambiguity Attitude," *Journal of Economic Theory*, 118 (2), 133–173. [1394,1411,1412]

GILBOA, ITZHAK, AND MASSIMO MARINACCI (2016): "Ambiguity and the Bayesian Paradigm," in *Readings in Formal Epistemology*. Springer, 385–439. [1393]
GILBOA, ITZHAK, AND DAVID SCHMEIDLER (1989): "Maxmin Expected Utility With Non-Unique Prior," *Journal of Mathematical Economics*, 18 (2), 141–153. [1411]
GOLDSTEIN, WILLIAM M., AND HILLEL J. EINHORN (1987): "Expression Theory and the Preference Reversal Phenomena," *Psychological Review*, 94 (2), 236–254. [1394,1403]
GONZALEZ, RICHARD, AND GEORGE WU (1999): "On the Shape of the Probability Weighting Function," *Cognitive Psychology*, 38 (1), 129–166. [1397]
GUL, FARUK, AND WOLFGANG PESENDORFER (2015): "Hurwicz Expected Utility and Subjective Sources," *Journal of Economic Theory*, 159, 465–488. [1393-1397,1411]
HEATH, CHIP, AND AMOS TVERSKY (1991): "Preference and Belief: Ambiguity and Competence in Choice Under Uncertainty," *Journal of Risk and Uncertainty*, 4 (1), 5–28. [1393]
HURWICZ, LEONID (1951): "Some Specification Problems and Applications to Econometric Models," *Econometrica*, 19 (3), 343–344. [1394]
KAHNEMAN, DANIEL, AND AMOS TVERSKY (1979): "Prospect Theory: An Analysis of Decision Under Risk," *Econometrica*, 47, 263–291. [1403]
KLIBANOFF, PETER, MASSIMO MARINACCI, AND SUJOY MUKERJI (2005): "A Smooth Model of Decision Making Under Ambiguity," *Econometrica*, 73 (6), 1849–1892. [1411]
KOCHER, MARTIN G., AMREI MARIE LAHNO, AND STEFAN T. TRAUTMANN (2018): "Ambiguity Aversion Is Not Universal," *European Economic Review*, 101, 268–283. [1393]
LANDIS, J. RICHARD, AND GARY G. KOCH (1977): "The Measurement of Observer Agreement for Categorical Data," *Biometrics*, 33 (1), 159–174. [1403]
LOOMES, GRAHAM, AND UZI SEGAL (1994): "Observing Different Orders of Risk Aversion," *Journal of Risk and Uncertainty*, 9 (3), 239–256. [1398,1401,1405]
LUCE, R. DUNCAN, AND HOWARD RAIFFA (2012): *Games and Decisions: Introduction and Critical Survey*. Courier Corporation. [1394]
MACCHERONI, FABIO, MASSIMO MARINACCI, AND ALDO RUSTICHINI (2006): "Ambiguity Aversion, Robustness, and the Variational Representation of Preferences," *Econometrica*, 74 (6), 1447–1498. [1411]
MACHINA, MARK J., AND DAVID SCHMEIDLER (1992): "A More Robust Definition of Subjective Probability," *Econometrica*, 60 (4), 745–780. [1397]
NAU, ROBERT F. (2006): "Uncertainty Aversion With Second-Order Utilities and Probabilities," *Management Science*, 52 (1), 136–145. [1411]
PRELEC, DRAZEN (1998): "The Probability Weighting Function," *Econometrica*, 66, 497–528. [1394,1403]
QUIGGIN, JOHN (1982): "A Theory of Anticipated Utility," *Journal of Economic Behavior & Organization*, 3 (4), 323–343. [1394]
SAVAGE, LEONARD J. (1954): *The Foundations of Statistics*. John Wiley & Sons. [1395]
SCHMEIDLER, DAVID (1989): "Subjective Probability and Expected Utility Without Additivity," *Econometrica*, 57, 571–587. [1394,1411]
SEGAL, UZI, AND AVIA SPIVAK (1990): "First Order versus Second Order Risk Aversion," *Journal of Economic Theory*, 51 (1), 111–125. [1397]
TRAUTMANN, STEFAN T., AND GIJS VAN DE KUILEN (2015): "Ambiguity Attitudes," in *The Wiley Blackwell Handbook of Judgment and Decision Making*, Vol. 1, 89–116. [1393]
TVERSKY, AMOS, AND DANIEL KAHNEMAN (1992): "Advances in Prospect Theory: Cumulative Representation of Uncertainty," *Journal of Risk and Uncertainty*, 5 (4), 297–323. [1394,1403]
WAKKER, PETER P. (2008): "Explaining the Characteristics of the Power (CRRA) Utility Family," *Health Economics*, 17 (12), 1329–1344. [1401]

Co-editor Alessandro Lizzeri handled this manuscript.

Manuscript received 6 December, 2020; final version accepted 14 November, 2022; available online 20 December, 2022.