Ambiguity Aversion and Epistemic Uncertainty*

Craig R. Fox
Anderson School of Management, University of California Los Angeles

Michael Goedde-Menke
School of Business and Economics, University of Münster

David Tannenbaum
Eccles School of Business, University of Utah

September 8, 2021

Abstract: We propose that ambiguity aversion reflects distaste for betting on one's relative ignorance, but only to the extent that uncertainty is seen as inherently knowable (epistemic). In contrast, uncertainty viewed as random (aleatory) can provide an attractive hedge against betting on one's ignorance. We refer to this account as the *epistemic uncertainty aversion* hypothesis, which allows for the simultaneous observation of ambiguity aversion and preference for compound lotteries involving both chance and ambiguity. In preregistered experiments involving Ellsberg urns and naturalistic events we show: (1) under conditions of ignorance, ambiguity averse decision makers prefer betting on a greater balance of aleatory to epistemic uncertainty; (2) this preference for an aleatory hedge increases as subjective knowledge decreases, and can lead decision makers to choose stochastically dominated alternatives; (3) an uncertain prospect can be framed as more epistemic or aleatory to influence its overall attractiveness. These findings collectively violate several prominent models of ambiguity aversion, but can be accomodated by generic source models. *JEL Codes: C91, D81, D91.*

Keywords: Ellsberg paradox; risk; ambiguity; epistemic uncertainty; aleatory uncertainty

---

*Authors contributed equally and authorship order was determined alphabetically. Fox: craig.fox@anderson.ucla.edu; Goedde-Menke: michael.goedde-menke@uni-muenster.de; Tannenbaum: david.tannenbaum@utah.edu. The positioning of this work was influenced by discussions with Alex Imas, and we also benefitted from feedback from seminar participants at the University of Chicago, HEC Paris, IDC Herzliya, New York University, MIT, University of California at San Diego, University of Pennsylvania, Yale University, Behavioral Decision Research in Management (2018), Society for Judgment and Decision Making (2018), Winter Judgment and Decision Making Symposium (2019) the Society for Personality and Social Psychology (2019) and the D-TEA (2020) conferences. We are grateful for helpful comments and discussions with Mohammed Abdellaoui, Dan Benjamin, Han Bleichrodt, David Budescu, Alain Cohn, Enrico Decidue, Alex Imas, Thomas Langer, Mark Machina, Kristof Madarasz, and Rakesh Sarin. Fox acknowledges support from the National Science Foundation to C.R. Fox and G. Ülkümen [SES-1427469]. Goedde-Menke thanks the Fritz Thyssen Stiftung for covering part of his travel expenses during his stay as visiting scholar at the UCLA Anderson School of Management.

1 Introduction

The earliest treatments of decision under uncertainty construe acts as chance lotteries with known probability distributions over outcomes (e.g., Pascal, 1852; Bernoulli, 1954; von Neumann and Morgenstern, 1944). When objective probabilities are not available, decision makers may assess probabilities subjectively, with some degree of vagueness or doubt. For roughly the last century decision theorists have debated how to properly characterize decisions that are made under these conditions. Knight (1921) famously distinguished decision under *risk*, in which probabilities can be measured, from decision under *uncertainty*, in which they cannot. He argued that entrepreneurs earn superior profits due in part to their willingness to bear unmeasurable uncertainty. Contemporaneously, Keynes (1921) distinguished probability, which he characterized as the balance of evidence supporting a focal proposition, from the total weight of evidence supporting its balance. He argued that one should favor actions with equal probabilities that are supported by more evidence.

Subjectivists dismissed the relevance of probability vagueness. Notably, Ramsey (1931) proposed that subjective probability should be measured by preferences between bets so that vagueness does not influence choice independently of its impact on overall belief strength. Likewise, although Savage (1954) acknowledged that subjective probabilities may be vague, he dismissed the relevance of second-order uncertainty in his development of subjective expected utility theory.

The debate concerning second-order uncertainty gained renewed attention when Ellsberg (1961) described decision problems in which the common preference to bet on known rather than unknown probabilities violate the *sure-thing principle*, a key axiom of Savage's model. Ellsberg's simplest demonstration, the "two-color problem," involves two urns, each containing red and black balls. The first urn contains 50 red balls and 50 black balls, while the second urn contains 100 red and black balls of unknown proportion. Most people report a preference to bet on a red ball being drawn at random from the 50-50 urn than from the unknown probability urn, and also prefer a bet on a black ball being drawn at random from the 50-50 urn than from the unknown probability urn. This violates the additivity assumption of subjective expected utility theory, as subjective probabilities cannot sum to 1 for both urns. Ellsberg interpreted this pattern as reflecting aversion to ambiguity, which he characterized as "a quality depending on the amount, type, reliability and 'unanimity' of the information, and giving rise to one's degree of 'confidence' in an estimate of relative likelihoods" (Ellsberg, 1961, p. 657).[1]

[1]While Ellsberg's examples were hypothetical, ambiguity aversion has been empirically validated in numerous studies using variations of Ellsberg's paradigm (for excellent reviews see Camerer and Weber, 1992; Machina and Siniscalchi, 2014; Trautmann and Van De Kuilen, 2015). Moreover, ambiguity aversion has been

The present paper tests a novel behavioral interpretation of ambiguity aversion. We provide evidence from six experimental studies that ambiguity aversion is driven by distaste for betting when one feels relatively ignorant or uninformed — but only to the extent that the relevant uncertainty is seen as inherently knowable or *epistemic* in nature. Thus, the addition of randomness or *aleatory* uncertainty provides a hedge against betting on one's relative ignorance and can make a prospect more attractive. Note that in the present account the distinction between epistemic and aleatory uncertainty is subjective and can be influenced by how options are described, as we will show in our final two studies. Our set of results cannot be accommodated by prominent models of decision under uncertainty, and contradict the common notion that ambiguity aversion is a manifestation of aversion to compound lotteries. However, generic source models provide sufficient flexibility to accommodate our findings of aversion to epistemic uncertainty.

1.1 Ambiguity and Compound Lottery Aversion

Early models in the economics literature captured ambiguity aversion by allowing decision makers to express pessimism to subjective probabilities, either by applying the worst from a range of possible subjective priors (Maxmin Expected Utility; Gilboa and Schmeidler, 1989), or by underweighting subjective probabilities (Choquet Expected Utility; Gilboa, 1987; Schmeidler, 1989). Alternative explanations construe bets on Ellsberg's ambiguous urn as a two-stage, compound lottery (Segal, 1987). The first stage can be thought of as a lottery over possible compositions of the unknown urn, for which the decision maker has a set of possible priors, while the second stage reflects a random draw of a ball from the urn that obtains in the first stage. According to this view, a systematic preference to bet on the known urn over the unknown urn represents a failure to reduce the two-stage lottery to its 50-50 equivalent, and ambiguity aversion can be viewed as aversion to compound lotteries. In particular, Halevy (2007) provided evidence that ambiguity neutrality in the two-color Ellsberg problem is strongly correlated with a tendency to price a simple 50-50 lottery equal to compound lotteries with objective probabilities that reduce to 50-50. He concludes that "...failure to reduce compound (objective) lotteries is the underlying factor of the Ellsberg paradox" (for similar results see also Abdellaoui et al., 2015; Dean and Ortoleva, 2019, but see Bernasconi and Loomes, 1992).

Four main interpretations have been proposed for why the simple risky lottery is preferred

leveraged to explain a wide range of phenomena such as portfolio choices and asset prices (Dow and da Costa Werlang, 1992; Gollier, 2011), insurance transactions (Hogarth and Kunreuther, 1985; Alary et al., 2013), incomplete contracts (Mukerji, 1998), financial markets (Mukerji and Tallon, 2001), brand choice (Muthukrishnan et al., 2009), vaccination decisions (Ritov and Baron, 1990), and strategic choice in games (Pulford and Colman, 2007; Vives and Feldman Hall, 2018).

to the compound, ambiguous lottery. First, individuals may fail to reduce compound lotteries even when all stages entail objective probabilities (Halevy, 2007). Second, individuals may accept second-stage (objective) lotteries as given, but be pessimistic with respect to their first-stage subjective probabilities, and underweight such first-stage beliefs in the valuation process (Recursive Rank-Dependent Utility; Segal, 1987, 1990). Third, individuals may be more risk averse for first-stage (subjective) lotteries than second-stage (objective) lotteries (Recursive Expected Utility; Klibanoff et al., 2005). Fourth, subjective probabilities may be more strongly underweighted than corresponding objective probabilities (source models with second-order probabilistic sophistication; Ergin and Gul, 2009).

1.2 Ambiguity and Epistemic Uncertainty Aversion

A separate stream of research in psychology breaks with the tradition of modeling ambiguity in terms of second-order probability distributions, multiple priors, or compound lotteries, and instead construes ambiguity aversion as distaste for acting in situations where a decision maker feels relatively ignorant, unskilled, or uninformed (Frisch and Baron, 1988; Heath and Tversky, 1991; Fox and Tversky, 1995; Fox and Weber, 2002; Hadar et al., 2013). In support of this hypothesis, Heath and Tversky (1991) demonstrated that although decision makers prefer betting on chance events to uncertain events of matched probability in domains where they lack expertise, they often prefer betting on uncertain events to chance events in domains where they feel particularly knowledgeable or competent. For example, in one study students who rated themselves as knowledgeable about football and ignorant about politics preferred to bet on their predicted winner of a football game that they judged to have, say, a 70% chance of winning rather than a chance gamble involving the draw of a winning poker chip from an urn containing 70 out of 100 winning chips. These same subjects, however, preferred betting on chance to their prediction of which presidential candidate would win various states in the 1988 election, consistent with ambiguity aversion. Meanwhile, students who rated themselves as knowledgeable about politics and ignorant about football exhibited the opposite pattern, preferring to bet on politics to chance and chance to football.

To motivate the link between subjective knowledge and ambiguity aversion, Heath and Tversky (1991) asserted that consequences of bets include not only their monetary outcomes but also the "psychic payoffs of satisfaction or embarrassment [that] can result from self-evaluation or evaluation by others" where "the credit and the blame associated with an outcome depend . . . on the attributions for success or failure" (p. 7). Losing a bet because of one's ignorance is more embarrassing than losing because of chance; winning a bet because of one's knowledge is more gratifying than winning because of chance. Indeed, laboratory experiments support the notion that evaluation from (real or imagined) others contributes to the ambiguity aversion phenomenon. Curley et al. (1986) found that ambiguity aversion is exacerbated in the presence of observers. Trautmann et al. (2008) replicated this effect and found that ambiguity aversion diminishes or disappears in situations where observers do not know the decision maker's preferences over outcomes and therefore cannot assign credit or blame for choices. They showed further that ambiguity aversion is more pronounced among individuals who score higher on a scale that measures fear of negative evaluation by others.

It stands to reason that decision makers may feel especially blameworthy for betting on their ignorance when they are reminded that they are more knowledgable about other bets or that other individuals are better informed. Indeed, ambiguity aversion appears to be driven by such contrasting states of knowledge (Fox and Tversky, 1995; Fox and Weber, 2002; Hadar et al., 2013). For instance, Fox and Tversky (1995) reported that ambiguity aversion is pronounced in comparative contexts where decision makers evaluate both risky and ambiguous bets simultaneously, but diminishes or disappears when separate groups of decision makers evaluate these bets in isolation so that there is no explicit contrast between urns to highlight a decision maker's comparative ignorance. Moreover, Chow and Sarin (2002) reported that people find betting on their own comparative ignorance less aversive when relevant information is available to no one. For instance, people are willing to pay more to bet on which of two apples has a greater number of seeds *before* the apples have been sliced open than *after* they have been sliced open and the seeds have been counted by someone else. Taken together, this literature suggests that ambiguity aversion reflects reluctance to bet in situations where the decision maker is relatively ignorant or uninformed concerning target outcomes, *but only to the extent that outcomes are seen as inherently predictable or knowable* at the time the decision is made.

The foregoing distinction between predictable and unpredictable outcomes recalls a longstanding philosophical distinction between *epistemic* and *aleatory* uncertainty (Hacking, 2006). Aleatory uncertainty is attributed to randomness or stochastic processes; decision under risk therefore involves pure aleatory uncertainty. Epistemic uncertainty, by contrast, is attributed to deficiencies in one's knowledge, expertise, or information; Knightian uncertainty is therefore at least partly epistemic. Laboratory studies have found that people intuitively differentiate between these two dimensions of uncertainty, and that this distinction has meaningful consequences for various domains of judgment and choice (Fox and Ülkümen, 2011; Ülkümen et al., 2016; Tannenbaum et al., 2017; Walters et al., 2020; Fox et al., 2020; Krijnen et al., 2021). Experiments have shown that outcomes of predictions made under greater epistemic uncertainty (e.g., whether or not one has correctly answered a trivia question) are associated with stronger attributions of credit for predicting correctly and blame for predicting incorrectly, whereas outcomes of predictions made under greater aleatory uncertainty (e.g., whether or not one has correctly predicted the outcome of a fair coin toss) are associated with stronger attributions of good or bad luck (Fox et al., 2020). This result suggests that ambiguity aversion will be stronger in situations where uncertainty is seen as more epistemic in nature, especially when decision makers consider themselves to be relatively ignorant, incompetent, or uninformed (and therefore more exposed to blame for making choices that result in inferior outcomes). Ambiguity aversion should be less pronounced in situations where outcomes are determined at least partly by chance (where luck plays a role in outcomes) or when decision makers consider themselves to be relatively knowledgeable, competent, or well-informed (so that there is the potential to claim some credit for making choices that result in superior outcomes).

All of these qualitative observations can be accommodated by the flexibility of generic source models such as prospect theory (Tversky and Kahneman, 1992) that weight events differently for different sources of uncertainty (Tversky and Fox, 1995; Abdellaoui et al., 2011).[2] In particular, events that the decision maker views as more epistemic and feels relatively ignorant about are assigned lower decision weights than corresponding events that the decision maker views as more aleatory or feels relatively competent about.

1.3 The Extended Ellsberg Paradigm as a Critical Test

Consider again the Ellsberg (1961) two-color problem. Betting on a color drawn from the known probability urn is a decision under purely aleatory uncertainty, and in this case selecting the "incorrect" color can be attributed entirely to bad luck. In contrast, betting on a color drawn from the unknown probability urn represents a decision under uncertainty that is *partly* epistemic (because the composition of the urn is, in principle, knowable). Since the decision maker is more ignorant about the composition of the unknown than known probability urn, betting on the incorrect color exposes the decision maker to potential blame or self-recrimination whereas betting on the correct color confers little potential credit or self-congratulation.

The foregoing analysis suggests that a variation of Ellsberg's paradigm could yield a bet that is even less attractive than a single draw from the unknown probability urn. Betting on whether the contents of the unknown probability urn are mostly red or black is a bet under *purely* epistemic uncertainty and should therefore be less attractive for ambiguity averse decision makers than a bet on a single draw from the unknown probability urn. Importantly, a single draw from the unknown probability urn includes an *aleatory hedge*: even if the

[2] We refer to these models as generic source models to distinguish them from source models with second-order probabilistic sophistication (Ergin and Gul, 2009).

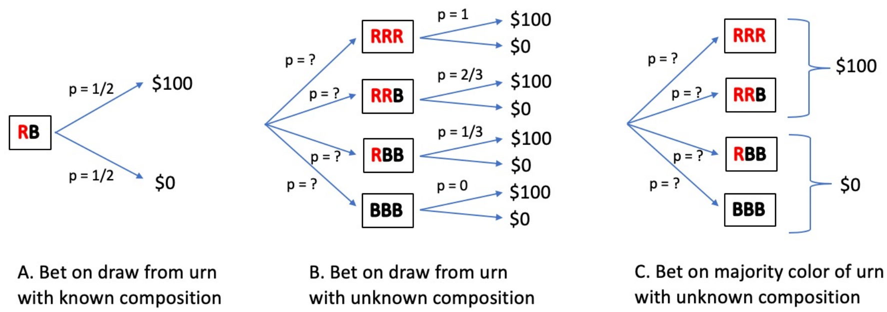

Figure 1: Extended Ellsberg two-color problem. This Figure displays a simplified extension of the Ellsberg two-color problem construed as one- and two-stage lotteries. The paradigm includes a bet that offers \$100 if a red ball is drawn from an urn containing exactly one red ball and one black ball (Panel A), \$100 if a red ball is drawn from an urn containing three balls each of which may be red or black (Panel B), and \$100 if the majority of balls are red in an urn containing three balls each of which may be red or black (Panel C). The bet in Panel A can be construed as a single-stage lottery with a 50% chance of paying \$100 (and nothing otherwise); the bet in Panel B can be construed as a two-stage lottery with an unknown probability of various possible compositions of the urn followed by a draw from the urn that obtains in the first-stage, yielding a probability of anywhere from zero to one of paying \$100 (and nothing otherwise); the bet in Panel C can be construed as a single-stage lottery in which there is an unknown probability anywhere from zero to one that the urn will contain a majority of red balls and therefore pay \$100 (and nothing otherwise).

decision maker chooses to bet on what turns out to be the minority color, she can still get lucky and win the prize, provided there is at least one ball of that color in the urn.

The extended Ellsberg paradigm outlined above can provide a critical test of conventional economic models of decision under uncertainty against the present account that ambiguity aversion is driven by a distaste for betting on epistemic (but not aleatory) uncertainty. To see how, consider a simplified extension of Ellsberg's two-color problem, illustrated in Figure 1, for a decision maker who bets on red. The first two panels display the conventional two-color problem for a bet on drawing a red ball from an urn containing one red and one black ball (Panel A) and an urn containing three balls each of which is either red or black but whose composition is unknown (Panel B). The third panel displays a bet that the majority color in the unknown composition urn is red (Panel C). Note that an even number of balls are required for the first bet in order to have a 50-50 composition, and an odd number of balls are required for the other two bets in order to guarantee a majority color.

As we show in Section 2, if we assume symmetric priors over ball color[3], Subjective

[3]Assuming symmetric priors over ball colors is common in Ellsberg-related studies on ambiguity aversion (e.g., Halevy, 2007; Chew et al., 2017) and consistent with empirical evidence (Abdellaoui et al., 2011). We rely on this assumption only for the first two out of six studies in this paper. Note that an SEU decision maker with asymmetric priors over ball colors and a choice of which color to bet on will always prefer Bet B and Bet C to Bet A and thus be ambiguity seeking. In this case, the relative attractiveness of Bet B and Bet C is determined by the probabilities assigned to the "RRB" and "RBB" states of the world. When the decision maker assigns equal probabilities to both states, she will be indifferent between Bet B and Bet C. When the "RRB" state is believed to be more likely than the "RBB" state, the decision maker will prefer

Expected Utility (SEU; Savage, 1954) predicts that decision makers will be indifferent between these three bets. Moreover, most of the prominent economic models designed to accommodate ambiguity aversion cannot also accommodate the simultaneous preference of Bet A to Bet B (the conventional Ellsberg pattern) and the preference of Bet B to Bet C (preferring to bet on a single draw from an unknown probability urn to betting on the contents of an unknown probability urn).

In sum, we propose that ambiguity aversion is driven by distaste for betting on one's relative ignorance under conditions of epistemic uncertainty. There are three main implications of the *epistemic uncertainty aversion hypothesis*: (1) under conditions of ignorance, ambiguity averse decision makers prefer betting on a greater balance of aleatory to epistemic uncertainty (holding judged probability constant); (2) this preference for an aleatory hedge diminishes with increasing subjective knowledge; and (3) viewing an uncertain prospect as epistemic or aleatory is subjective and can be influenced by the framing of decision options. In this paper we present six experimental studies that collectively test these three implications of the epistemic uncertainty aversion hypothesis, and contrast this account with prior accounts of decision under ambiguity.

2 Experimental Paradigm and Theoretical Predictions

2.1 Ellsberg Paradigm

The starting point for our experimental studies is Ellsberg's (1961) two-color problem as described above. We add a third lottery to the choice set for which, like the second lottery, the composition of the urn is unknown to the decision maker and therefore ambiguous. However, instead of betting on a single random draw from the urn, the decision maker bets on whether the majority of the urn's balls match the predicted winning color. Subjects are therefore confronted with pairwise choices between the following three options: predicting a random draw from a 50-50 urn, predicting a random draw from an unknown probability urn, and predicting the majority color of an unknown probability urn. The unknown probability urns contain an odd number of balls to guarantee a majority color.

We label these three lotteries according to their corresponding nature of uncertainty. A decision maker who chooses to draw from the 50-50 urn faces only aleatory uncertainty as the composition of the urn is known and the outcome is completely random; we therefore refer to this single-stage, risky lottery as the *A-lottery* (L_A) in the analysis that follows. The classic Ellsberg ambiguous lottery exposes the decision maker to both epistemic uncertainty (arising

Bet C to Bet B (and vice versa).

from the unknown urn composition) and aleatory uncertainty (due to the random draw from the urn given its composition); we therefore refer to this compound, mixed-uncertainty lottery as the *EA-lottery* (L_{EA}). The lottery involving the third urn asks subjects to bet on its majority color, and because no random draw is made the bet entails only epistemic uncertainty; we therefore refer to it as the *E-lottery* (L_E). This final, single-stage lottery is crucial to our analysis as it allows us to unconfound number of stages (single-stage vs. compound) from source of uncertainty (epistemic vs. aleatory). Thus, we can test whether ambiguity aversion is more closely associated with aversion to epistemic uncertainty or aversion to compound lotteries.

2.2 Theoretical Predictions

We next turn to an analysis of predictions made by various theoretic models concerning preferences across our three lotteries. The three-urn setup underlying our general paradigm can be characterized as follows: Each lottery L_i, where $i \in \{A, EA, E\}$, pays x if the decision maker correctly guesses the color — red (R) or black (B) — of the drawn ball (for the *A-lottery* or the *EA-lottery*) or correctly predicts the color of the majority of balls in the urn (for the *E-lottery*), and otherwise pays zero. To facilitate the derivation of predictions of different utility models, we assume that the decision maker is indifferent between betting on red or black, i.e., she acts as if both colors are equally likely. This symmetry assumption has been empirically validated (Abdellaoui et al., 2011) and is common in Ellsberg-related studies on ambiguity aversion (e.g., Chew et al., 2017). Due to symmetry, the utilities from betting on either color are the same. We therefore only use the utility of correctly betting on a particular color for a given urn in the following derivations. Readers wishing to forgo a technical treatment of prominent models can skip to Section 2.2.7.

2.2.1 Subjective Expected Utility

In SEU (Savage, 1954), the decision maker assigns a subjective probability p to each state of the world associated with outcomes of a lottery. When objective probabilities are provided (as for the *A-lottery*), it is assumed that subjective and objective probabilities coincide. In addition, SEU requires that decision makers adhere to probabilistic sophistication and correctly reduce compound lotteries, making the reduction of compound lottery axiom (RCLA) an integral element of SEU (e.g., Anscombe and Aumann, 1963; Segal, 1990).

The symmetry assumption, in combination with RCLA, implies that subjective probabilities for both events R and B are 0.5 for the two urns involving ambiguity (*EA-lottery* and *E-lottery*). This also holds for the *A-lottery* as the objective probability is 0.5 for drawing either color. Given a utility function u, a decision maker's subjective expected utility of lottery L_i is given by:

$$U_{SEU}(L_i) = 0.5u(x),$$

where we normalize $u(0) = 0$. SEU thus models all three lotteries as having the same utility, so that an SEU decision maker is indifferent between them:

$$SEU \Rightarrow L_A \sim L_{EA} \sim L_E.$$

Hence, SEU is unable to accommodate observations of ambiguity aversion ($L_A \succ L_{EA}$ and $L_A \succ L_E$). Moreover, SEU cannot accommodate a strict preference between the purely epistemic and mixed uncertainty lotteries, as predicted by compound lottery aversion ($L_E \succ L_{EA}$) or epistemic uncertainty aversion ($L_{EA} \succ L_E$).

2.2.2 Choquet Expected Utility

Choquet Expected Utility (CEU) is able to accommodate ambiguity aversion as it allows the decision maker to express pessimism with respect to the subjective probabilities inherent in ambiguous lotteries (Gilboa, 1987; Schmeidler, 1989).[4]

In CEU, a capacity function w is employed. The capacity function is a non-additive function that maps events onto the unit interval and is monotonic in terms of inclusion (Gilboa, 1987; Schmeidler, 1989). For a probabilistically sophisticated decision maker (Machina and Schmeidler, 1992), the capacity function can therefore be thought of as a probability weighting function that transforms subjective probabilities (Tversky and Fox, 1995; Fox and Tversky, 1998; Wakker, 2004). We follow the CEU axiomatization in Schmeidler (1989) and assume that expected utility is applied to risk and that capacities are either strictly convex or strictly concave. The shape of the capacity function determines whether a decision maker exhibits global ambiguity aversion (convex w) or ambiguity seeking (concave w). CEU with a linear capacity function (the identity function) coincides with SEU. Assuming symmetry, the utility of an ambigous lottery (L_{EA}, L_E) under CEU is given by:

$$U_{CEU}(L_i) = w(0.5)u(x).$$

For objective lotteries (L_A), CEU coincides with SEU. Under the assumption of a strictly

[4]An alternative approach to modelling pessimism regarding ambiguous prospects is employed in Maxmin Expected Utility (MEU; Gilboa and Schmeidler 1989). In MEU, the decision maker applies the worst possible prior from a convex set of priors in his evaluation of ambiguous lotteries. Because in our setup MEU can be accommodated by CEU when the capacity function is convex (which is needed to capture ambiguity aversion), we only discuss the latter model.

convex capacity ($w(0.5) < 0.5$), CEU predicts that the *A-lottery* is preferred to the *E-lottery*.[5] With respect to the *EA-lottery*, a CEU prediction can only be derived when it adopts the Anscombe-Aumann framework (Schmeidler, 1989) and incorporates RCLA. In that case the *EA-lottery* is treated the same as the *E-lottery*, and the two yield the same Choquet expected utility. If CEU is axiomatized in a Savagean domain (Gilboa, 1987; Wakker, 1987), then it is unclear how compound lotteries are evaluated (Chew et al., 2017). The following preference pattern therefore emerges for CEU incorporating convex capacities and RCLA:

$$CEU \Rightarrow L_A \succ L_{EA} \sim L_E.$$

Thus, CEU is able to accommodate ambiguity averse behavior ($L_A \succ L_{EA}$ and $L_A \succ L_E$), but cannot accommodate a strict preference between the purely epistemic and mixed uncertainty lotteries, as predicted by compound lottery aversion ($L_E \succ L_{EA}$) or epistemic uncertainty aversion ($L_{EA} \succ L_E$).

2.2.3 Recursive Rank-Dependent Utility

The foregoing discussion suggests that a strict preference for the *EA-lottery* over the *E-lottery* (or vice versa) cannot be captured when RCLA applies. Recursive rank-dependent utility (RRDU) is a valuation approach that relaxes RCLA (Segal, 1987, 1990). It builds on rank-dependent utility (Quiggin, 1982) but allows for an explicit valuation of the first and second stages within compound lotteries. Relaxing RCLA also receives empirical support, as studies find that reduction of compound risk typically fails (e.g., Abdellaoui et al., 2015).

The general mechanism of RRDU is best illustrated by an example on how the (ambiguous) two-stage, mixed-uncertainty lottery, L_{EA}, is evaluated. Suppose a decision maker constructs her own second-order (first-stage) subjective belief for an urn that contains three balls in total, as in Figure 1. She considers all possible urn compositions, including both extreme and intermediate possibilities: all red balls, two red balls (and one black ball), one red ball (and two black balls), or no red balls. The subjective probability that the urn contains balls of only one color is α for each of the two single-color scenarios (with $\alpha \leq 0.5$), and the residual

[5] Since the goal of this paper is to explore ambiguity aversion, assuming a strictly convex capacity in CEU for deriving our predictions is in order because otherwise traditional applications of CEU with no inflection point in the capacity function would not be able to accommodate the standard Ellsberg preference pattern ($L_A \succ L_{EA}$). Moreover, even when assuming a concave capacity the result that a CEU decision maker is indifferent between the *EA-lottery* and the *E-lottery* would still hold. However, in this case both urns would be preferred to the *A-lottery*, implying ambiguity seeking behavior. Finally, even if we assume a weighting function that is inverse-S shaped as in prospect theory (Tversky and Kahneman, 1992), the principle of subcertainty (weights of complementary events sum to less than 1) suggests that the assumption of convex capacities is adequate since both imply that the capacity of a 0.5 probability is less than 0.5; i.e., $w(0.5) < 0.5$.

probability $(1-2\alpha)$ is divided equally between the remaining two urn compositions (yielding $0.5-\alpha$).

Assume the decision maker bets on red, and in the first step evaluates second-stage lotteries. Depending on the urn distribution, her payoff profiles of the second-stage lotteries are as follows: $(\$x, 1; \$0, 0)$ if all balls are red, $(\$x, \frac{2}{3}; \$0, \frac{1}{3})$ if the urn contains two red balls, $(\$x, \frac{1}{3}; \$0, \frac{2}{3})$ if the urn contains one red ball, and $(\$x, 0; \$0, 1)$ if no balls are red. The decision maker's rank-dependent utility arising from such second-stage lotteries can be calculated as

$$U_{RDU}(x, p; 0, 1-p) = w(p)u(x),$$

where u is a common utility function (normalized to $u(0) = 0$) applied to both stages, and w is a probability weighting function. The weighting function must be strictly convex with nondecreasing elasticity to accommodate uniform ambiguity aversion (Segal, 1987). Such a convex weighting function implies that probabilities are always underweighted (i.e., $w(p) < p$).

In the next step, the decision maker transforms the utilities from second-stage lotteries into certainty equivalents

$$CE_{RDU}(x, p; 0, 1-p) = u^{-1}(w(p)u(x))$$

and evaluates the *EA-lottery* by employing the obtained certainty equivalents as prizes in the RDU formula. Thus, in this simplified example, the RRDU for betting on the two-stage, mixed-uncertainty lottery (L_{EA}) with prize x is

$$\begin{aligned} U_{RRDU}(L_{EA}) = u\Big(u^{-1}\big(w(1)u(x)\big)\Big) \cdot \Big(w(\alpha) - w(0)\Big) + \\ u\Big(u^{-1}\big(w(\frac{2}{3})u(x)\big)\Big) \cdot \Big(w(0.5) - w(\alpha)\Big) + \\ u\Big(u^{-1}\big(w(\frac{1}{3})u(x)\big)\Big) \cdot \Big(w(1-\alpha) - w(0.5)\Big) + \\ u\Big(u^{-1}\big(w(0)u(x)\big)\Big) \cdot \Big(w(1) - w(1-\alpha)\Big). \end{aligned}$$

Noting that $w(0) = 0$, and $w(1) = 1$, we get:

$$\begin{aligned} U_{RRDU}(L_{EA}) = u(x)\Big[w(\alpha) + \\ w(\frac{2}{3}) \cdot \Big(w(0.5) - w(\alpha)\Big) + \\ w(\frac{1}{3}) \cdot \Big(w(1-\alpha) - w(0.5)\Big)\Big]. \end{aligned}$$

Following Segal (1987, 1990), we furthermore assume that an RRDU decision maker does not distinguish between subjective and objective first-stage priors, yielding indifference between a risky single-stage lottery involving objective probabilities (*A-lottery*) and an ambiguous single-stage lottery involving subjective probabilities (*E-lottery*):

$$U_{RRDU}(L_A) = U_{RRDU}(L_E) = w(0.5)u(x).$$

It is easy to see that an RRDU decision maker displays ambiguity aversion ($L_A \succ L_{EA}$) if her priors are not extreme (i.e., she allows for the possibility that urns contain at least one red and one black ball; $\alpha < 0.5$).[6] Such a decision maker is indifferent between the *A-lottery* and the *E-lottery*, but prefers both to the *EA-lottery*:

$$RRDU \Rightarrow L_A \sim L_E \succ L_{EA}.$$

RRDU can therefore accommodate classic ambiguity averse behavior for the standard Ellsberg lotteries ($L_A \succ L_{EA}$), but not for the purely aleatory and epistemic lotteries ($L_A \succ L_E$). RRDU can also accomodate compound lottery aversion ($L_A \succ L_{EA}$ and $L_E \succ L_{EA}$), but not both choice patterns predicted by epistemic uncertainty aversion ($L_A \succ L_{EA}$ and $L_{EA} \succ L_E$).

2.2.4 Recursive Expected Utility

RRDU assumes that the decision maker applies the same utility function to first- and second-stage lotteries. This assumption is relaxed in the Recursive Expected Utility (REU) model advanced by Klibanoff et al. (2005). In this model, the decision maker applies a utility function u_s to (subjective) first-stage lotteries and a utility function u_o to (objective) second-stage lotteries, and ambiguity attitudes are determined by the relative concavities of the two utility functions (which are again normalized to $u_s(0) = 0$ and $u_o(0) = 0$). Aversion to ambiguous prospects holds when u_s is more concave than u_o, implying that the decision maker demonstrates a higher degree of risk aversion when facing lotteries involving subjective probabilities as compared to objective probabilities. REU coincides with SEU when u_s and u_o are identical.

We begin with an illustration of how an REU decision maker evaluates the *EA-lottery* — an ambiguous, two-stage lottery with subjective probabilities in the first stage and objective probabilities in the second stage. REU is formulated in three main steps. First, a decision maker forms a subjective belief over possible urn distributions and derives the corresponding second-stage lotteries based on this belief. Second, the decision maker constructs certainty

[6]In the case of extreme priors ($\alpha = 0.5$) an RRDU decision maker would be indifferent between all three lotteries.

equivalents for all second-stage lotteries based on their respective expected utilities according to u_o. Third, these certainty equivalents are then employed as input in a subjective expected utility evaluation using u_s. Consider the same decision scenario as in the previous section. The unknown urns contain three balls and all potential urn compositions are considered possible (all red, two red, one red, no red). Again, the decision maker assigns a probability α to each of the two single-color scenarios (with $\alpha \leq 0.5$), while the residual probability $(1 - 2\alpha)$ is divided equally between the remaining two urn compositions (yielding $0.5 - \alpha$). Due to symmetry, the decision maker's REU arising from the *EA-lottery* for either color is

$$\begin{aligned}
U_{REU}(L_{EA}) = & \Big[\alpha u_s\Big(u_o^{-1}\big(u_o(x)\big)\Big)\Big] + \\
& (0.5-\alpha)\Big[u_s\Big(u_o^{-1}\big(\frac{2}{3}\cdot u_o(x)\big)\Big)\Big] + \\
& (0.5-\alpha)\Big[u_s\Big(u_o^{-1}\big(\frac{1}{3}\cdot u_o(x)\big)\Big)\Big] + \\
& \Big[\alpha u_s\Big(u_o^{-1}\big(u_o(0)\big)\Big)\Big]
\end{aligned}$$

which reduces to

$$U_{REU}(L_{EA}) = \alpha u_s(x) + (0.5-\alpha)\Big[u_s\Big(u_o^{-1}\big(\frac{2}{3}\cdot u_o(x)\big)\Big) + u_s\Big(u_o^{-1}\big(\frac{1}{3}\cdot u_o(x)\big)\Big)\Big].$$

Because prior applications of REU do not distinguish risk attitudes when evaluating objective lotteries in the first versus second stage (Halevy, 2007), the *A-lottery* is evaluated as follows (applying the required transformation function $u_s \circ u_o^{-1}$):

$$U_{REU}(L_A) = u_s\Big(u_o^{-1}\big(0.5u_o(x) + 0.5u_o(0)\big)\Big) = u_s\Big(u_o^{-1}\big(0.5u_o(x)\big)\Big).$$

Applying the same argument to single-stage, ambiguous lotteries involving subjective probabilities (*E-lottery*) yields

$$\begin{aligned}
U_{REU}(L_E) &= \alpha u_s(x) + \alpha u_s(0) + (0.5-\alpha)u_s(x) + (0.5-\alpha)u_s(0) \\
&= \alpha u_s(x) + 0.5u_s(x) - \alpha u_s(x) \\
&= 0.5u_s(x).
\end{aligned}$$

Because ambiguity aversion implies that u_s is more concave than u_o (Klibanoff et al., 2005), an REU decision maker with extreme priors (i.e., $\alpha = 0.5$) exhibits the following preference pattern:

$$REU \Rightarrow L_A \succ L_{EA} \sim L_E.$$

If an REU decision maker's priors are not extreme (i.e., $\alpha < 0.5$), the preference ordering is as follows:

$$REU \Rightarrow L_A \succ L_{EA} \succ L_E.$$

REU therefore accommodates ambiguity aversion ($L_A \succ L_{EA}$ and $L_A \succ L_E$), but cannot accommodate the full preference pattern predicted by compound lottery aversion ($L_A \succ L_{EA}$ and $L_E \succ L_{EA}$). However, REU can accommodate the full preference pattern predicted by epistemic uncertainty aversion ($L_A \succ L_{EA}$ and $L_{EA} \succ L_E$), but only if priors are not extreme (i.e., decision makers allow for the possibility that urns contain at least one red and one black ball).

2.2.5 Source Models with Second-order Probabilistic Sophistication

A model that can capture preferences for betting on different sources of uncertainty while assuming second-order probabilistic sophistication (SPS) has been suggested by Ergin and Gul (2009). Their model allows for violations of RCLA and can accommodate distinct non-expected utility preferences across lottery stages.

In the following, we will assume an SPS decision maker who maximizes RRDU but employs different probability weighting functions in the two stages. Instead of assuming different utility functions for objective and subjective lotteries as in REU, the decision maker expresses her lack of confidence in subjective probabilities by underweighting them more strongly than objective probabilities. For empirical evidence suggesting that individuals tend to process subjective probabilities more pessimistically than objective probabilities, see Tversky and Fox (1995), Abdellaoui et al. (2011), and Baillon et al. (2018).

To illustrate, assume the same decision scenario as in the previous two sections when evaluating the *EA-lottery* (L_{EA}). The only difference is now that instead of applying the same probability weighting function in both stages, we distinguish a weighting function that is applied to first-stage (subjective) probabilities, w_s, from a probability weighting function that is applied to second-stage (objective) probabilities, w_o. Both functions are assumed to be strictly convex with nondecreasing elasticity (Segal, 1987). To capture the increased pessimism towards subjective probabilities, w_s must be more convex than w_o. The decision maker's RDU arising from objective, second-stage lotteries therefore remains

$$U_{RDU}(x, p; 0, 1-p) = w_o(p)u(x),$$

where u is a common utility function (normalized to $u(0) = 0$). The transformation of objective, second-stage utilities into certainty equivalents is also unchanged,

$$CE_{RDU}(x, p; 0, 1-p) = u^{-1}\big(w_o(p)u(x)\big).$$

However, the next step in the valuation process differs. When evaluating the (ambiguous) two-stage, mixed-uncertainty lottery (L_{EA}), the decision maker still employs the obtained certainty equivalents as prizes in the RDU formula. As before, the subjective probability of the urn containing balls of only one color is α for each of the two single-color scenarios, and the residual probability $(1-2\alpha)$ is divided equally between the remaining two urn compositions (two red or one red), yielding $0.5-\alpha$. Now instead of using the objective weighting function w_o, she employs the more convex subjective probability weighting function w_s to weight subjective priors, resulting in

$$\begin{aligned} U_{SPS}(L_{EA}) = u(x)\Big[&w_s(\alpha)+ \\ &w_o(\frac{2}{3})\cdot\Big(w_s(0.5)-w_s(\alpha)\Big)+ \\ &w_o(\frac{1}{3})\cdot\Big(w_s(1-\alpha)-w_s(0.5)\Big)\Big]. \end{aligned}$$

As the SPS representation allows for distinguishing between objective and subjective stage one priors, a corresponding decision maker is no longer indifferent between single-stage objective (L_A) and subjective lotteries (L_E), but assigns distinct values to each:

$$\begin{aligned} U_{SPS}(L_A) &= w_o(0.5)u(x), \\ U_{SPS}(L_E) &= w_s(0.5)u(x). \end{aligned}$$

Due to the stronger convexity of w_s than w_o, an SPS decision maker therefore prefers the *A-lottery* to the *E-lottery* and prefers both to the *EA-lottery*:

$$SPS \Rightarrow L_A \succ L_E \succ L_{EA}.$$

Hence, SPS accommodates ambiguity aversion ($L_A \succ L_{EA}$ and $L_A \succ L_E$) and compound lottery aversion ($L_A \succ L_{EA}$ and $L_E \succ L_{EA}$), but not epistemic uncertainty aversion ($L_A \succ L_{EA}$ and $L_{EA} \succ L_E$).

2.2.6 Generic Source Models

Although source models with second-order probabilistic sophistication cannot accommodate the simultaneous presence of ambiguity aversion and epistemic uncertainty aversion, source models that relax the assumption of probabilistic sophistication between sources are sufficiently flexible to accommodate such preference patterns (see Abdellaoui et al., 2011, for a model based on subjective probabilities from revealed preferences; see Tversky and Fox, 1995; Fox and Tversky, 1998, for a model based on judged probabilities from introspective judgment). We refer to these accounts as Generic Source Models (GSM).

Consider a decision maker who maximizes RDU, with probability weighting functions w that vary by source of uncertainty i:

$$U_{GSM}(x, p; 0, 1-p) = w_i(p)u(x).$$

Assume, due to symmetry, no preference for betting on red versus black so that the subjective probability[7] of winning any lottery is 0.5. Also assume that each lottery, L_i, entails a distinct source of uncertainty, $i \in \{A, EA, E\}$. Under GSM, the utility of each lottery is given by:

$$U_{GSM}(L_i) = w_i(0.5)u(x),$$

where w_i is the weighting function associated with source i, also known as the "source function." The epistemic uncertainty aversion hypothesis states that ambiguity aversion is driven by distaste for betting in situations where the decision maker feels comparatively ignorant or unformed, especially when the balance of epistemic to aleatory uncertainty is high. This can be accommodated within the GSM framework simply by assuming that the elevation of the source function w_i decreases (that is, the "degree of pessimism" in Abdellaoui et al., 2011, increases) with the interaction of a decision maker's comparative ignorance and perceived balance of epistemicness to aleatoriness. Assuming decision makers consider themselves less knowledgeable concerning unknown probabilities than known probabilities for Ellsberg urns (Fox and Tversky, 1995), and assuming the majority color in an urn is seen as more purely epistemic in nature than the outcome of a single draw from such an urn, then

$$w_A(0.5) > w_{EA}(0.5) > w_E(0.5),$$

so that

$$GSM \Rightarrow L_A \succ L_{EA} \succ L_E.$$

Hence, generic source models are, in principle, able to accommodate the simultaneous presence

[7]For simplicity in this example, we assume due to symmetry that the subjective probability p is 0.5. More generally, in a prospect theory framework the source-dependent decision weight can be expressed for a prospect that offers \$x if event E obtains and nothing otherwise as $w_i[P(E)]$, where P is the subjective probability (or, alternatively, judged probability) of event E (see Tversky and Fox, 1995; Fox and Tversky, 1998; Wakker, 2004).

of ambiguity aversion and epistemic uncertainty aversion. This, however, requires decision makers to be more pessimistic about subjective probabilities when lotteries entail greater epistemic uncertainty (i.e., more ignorance $\times$ epistemicness).

2.2.7 Theoretical Predictions: Summary

Taking stock of six prominent models of decision under uncertainty (SEU, CEU, RRDU, REU, SPS, GSM), all models other than SEU can accommodate standard ambiguity averse preferences ($L_A \succ L_{EA}$), and all models other than SEU and RRDU can also accommodate strict ambiguity averse preferences in our extended Ellsberg paradigm ($L_A \succ L_E$). The critical test concerns preferences between the *E-lottery* and the *EA-lottery* among ambiguity averse individuals. While RRDU and SPS predict strict compound lottery aversion ($L_E \succ L_{EA}$) among ambiguity averse individuals, only REU can, under the condition that priors are not extreme, accommodate strict compound lottery seeking between these two lotteries ($L_{EA} \succ L_E$) as predicted by the epistemic uncertainty aversion hypothesis. Additionally, the GSM framework provides sufficient flexibility to allow for modelling preference patterns that are consistent with our account.[8]

3 Experimental Evidence

We next turn to an experimental investigation of ambiguity and epistemic uncertainty aversion. Studies 1 and 2 employ our extended Ellsberg paradigm to test whether ambiguity aversion is more closely associated with epistemic uncertainty aversion or compound lottery aversion. Study 3 extends this test to more naturalistic bets involving soccer matches. Study 4 examines whether epistemic uncertainty aversion can lead to violations of stochastic dominance among less knowledgeable individuals who seek an aleatory hedge. Studies 5 and 6 examine whether the attractiveness of an aleatory hedge diminishes when reframed as more epistemic in nature. As we will show, the results of Studies 1-4 contradict all aforementioned models other than REU and GSM, and the results of Studies 5 and 6 contradict all models, including REU, but can be accomodated by the flexibility of GSM.

For all studies we determined sample sizes in advance of data collection. We preregistered hypotheses and analysis plans for all studies except Study 1. Materials, data, and code for all studies can be found at `https://researchbox.org/128&PEER_REVIEW_passcode=DFIFVM`.

[8] Of course, flexibility of generic source models comes at the expense of specificity, a point on which we will elaborate in the discussion.

3.1 Study 1: Extended Ellsberg Paradigm

We recruited 200 subjects from an online labor market (`www.prolific.ac`) to participate in a brief study in exchange for a £0.40 payment. Following the extended Ellsberg paradigm outlined in Section 2.1, subjects made a series of pairwise choices between three lotteries described in Table 1. The *A-lottery* involved a standard risky urn containing 50 red and 50 black poker chips, in which subjects would pick a color and then win a prize if a single randomly drawn chip matched that color; as such, it can be represented as a single-stage lottery involving purely aleatory uncertainty. The *EA-lottery* involved a standard Ellsberg urn containing 101 red and black poker chips[9] of unknown proportion, in which subjects would pick a color and then win a prize if a single randomly drawn chip matched that color; as such, it can be represented as a compound lottery involving a mixture of epistemic and aleatory uncertainty. The *E-lottery* was identical to the *EA-lottery* except that instead of drawing a single chip from the urn, the subject would draw all chips from the urn and win a prize if he or she had correctly predicted the majority color; as such, it can be represented as a single-stage lottery involving purely epistemic uncertainty.[10]

We presented lottery pairs in an order that was randomized for each subject, and for each pair asked subjects to indicate the lottery they prefer. The large majority of participants (91%) exhibited transitive preferences among their three choices.[11] Looking first at choices between the *A-lottery* and *EA-lottery*, we replicate the standard Ellsberg effect with 71% of subjects preferring L_A over L_{EA} ($p < 0.01$ by a binomial test). Likewise, 70% of subjects chose the risky single-stage *A-lottery* over the ambiguous single-stage *E-lottery* ($p < 0.01$), another manifestation of ambiguity aversion. Most subjects (56%) demonstrate consistent ambiguity aversion by preferring the *A-lottery* to both the *EA-lottery* and the *E-lottery*.

Our critical test involves the choice between the *EA-lottery* and *E-lottery*, in which a distaste for compound lotteries predicts a preference for the single-stage lottery ($L_E \succ L_{EA}$) among consistently ambiguity averse subjects (i.e., those who exhibit both $L_A \succ L_{EA}$ and $L_A \succ L_E$; $N = 113$). In contrast, a distaste for epistemic uncertainty predicts a preference for the mixed-uncertainty lottery ($L_{EA} \succ L_E$) among ambiguity averse subjects. In accord with the epistemic uncertainty aversion hypothesis, the majority of consistently ambiguity averse subjects prefer the *EA-lottery* to the *E-lottery* (65%; $p < 0.01$). Hence, individuals who exhibited consistent ambiguity aversion were nearly twice as likely to choose the compound, mixed-uncertainty *EA-lottery* than the single-stage (and purely epistemic) *E-lottery*. We

[9]We increase the traditional number of balls from 100 to 101 for the *EA-lottery* in order to guarantee a single majority color for the *E-lottery* and keep these two lotteries comparable.

[10]In this experiment we used generic labels "Lottery A," "Lottery B," and "Lottery C."

[11]Note that if subjects responded randomly we would expect 75% transitive orderings.

	A-lottery	EA-lottery	E-lottery
Description:	A bag is filled with exactly 50 red and 50 black poker chips. First, choose a color to bet on (Red or Black). Next, draw a single chip without looking from the bag. If the chip you pulled out is the color you predicted then you win \$100, otherwise you win nothing.	A bag is filled with 101 poker chips that are red and black, but you do not know their relative proportion. First, choose a color to bet on (Red or Black). Next, draw a single chip without looking from the bag. If the chip you pulled out is the color you predicted then you win \$100, otherwise you win nothing.	A bag is filled with 101 poker chips that are red and black, but you do not know their relative proportion. First, choose a color to bet on (Red or Black). Next, empty out the entire bag. If the majority of chips in the bag are the color you predicted then you win \$100, otherwise you win nothing.
Dimension of uncertainty:	Aleatory	Epistemic and Aleatory	Epistemic
Knightian uncertainty:	Risk	Ambiguity	Ambiguity
Stage type:	Single-stage	Compound	Single-stage

Table 1: Menu of alternatives employed in Study 1. The three lotteries employed in Study 1 as described to subjects, who made choices between each pair. The table also displays key characteristics of each lottery.

also note that the preference ordering predicted by a distaste for epistemic uncertainty ($L_A \succ L_{EA} \succ L_E$) is the modal preference ordering in the data, with 37% of all subjects displaying this pattern of preferences (see columns 6 and 7 of Table 2).[12]

3.2 Study 2: Extended Ellsberg Paradigm – Replication and Extension

In Study 2 we replicate the results of Study 1 while modifying the experimental design in a few important respects. First, we recruited a sample of German university students who made their decisions in a classroom setting rather than online. Second, choices were incentive-compatible: subjects were informed that some respondents would be selected at random to have one of their choices (also selected at random) to be played for real money. In particular, subjects would win a €50 gift card if the color associated with their lottery choice was correct (and €0 otherwise). Finally, we increased the sample size in Study 2 ($N = 567$) to provide greater statistical power to analyze behavior not only among consistently ambiguity averse subjects, but also among consistently ambiguity *seeking* individuals. Doing so allows us to further distinguish preferences over compound lotteries versus epistemic uncertainty. If ambiguity preferences are more generally associated with preferences for compound lotteries, then rare instances of ambiguity *seeking* ($L_{EA} \succ L_A$ and $L_E \succ L_A$) should also be associated with

[12] We can also examine the inverse analysis involving the proportion of subjects exhibiting classic ambiguity aversion ($L_A \succ L_{EA}$) conditional on a consistent preference against the purely epistemic lottery. Among subjects with a consistent distaste for the purely epistemic lottery ($L_A \succ L_E$ and $L_{EA} \succ L_E$; $N = 90$), a significant majority exhibited classic ambiguity aversion (82% showed $L_A \succ L_{EA}$; $p < 0.01$).

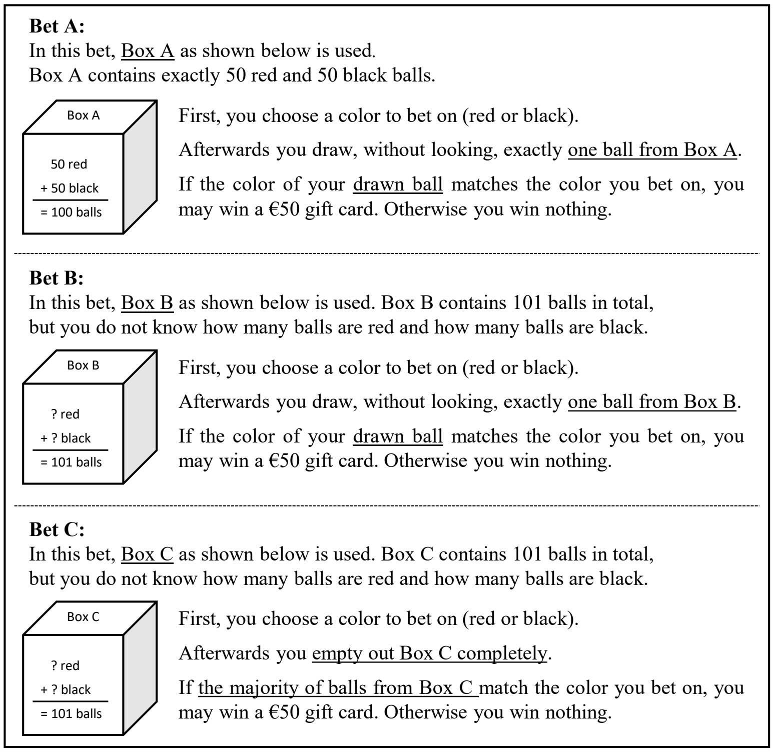

Figure 2: Menu of alternatives employed in Study 2. The three lotteries employed in Study 2 as described to subjects, who made choices between each pair. Bet A can be interpreted as a single-stage *A-lottery* which entails aleatory uncertainty only. Bet B can be interpreted as a two-stage, mixed-uncertainty *EA-lottery*. Bet C can be interpreted as a single-stage *E-lottery* which only entails epistemic uncertainty.

compound lottery seeking ($L_{EA} \succ L_E$). If ambiguity preferences are instead associated with preferences for epistemic uncertainty, then ambiguity seeking behavior should be associated with epistemic uncertainty seeking ($L_E \succ L_{EA}$). Note that none of the prominent models considered in Section 2, other than REU without extreme priors or GSM, can accommodate the simultaneous observation of ambiguity seeking and a strict preference for the *E-lottery* over the *EA-lottery* as predicted by the epistemic uncertainty hypothesis.

The structure of Study 2 was otherwise similar to that of Study 1. We presented first-year undergraduate business administration students with three pairs of lotteries using our extended Ellsberg paradigm, which they completed in a classroom setting. We presented choices on separate questionnaire pages and in an order that was randomized for each subject. For each choice, subjects indicated their preferred lottery and chose a color to bet on (red or black). Figure 2 provides a description of lotteries used in Study 2 (translated from German). The proportion of subjects indicating each preference ordering is displayed in the last two columns of Table 2.

Similar to Study 1, more than 90% of respondents exhibited transitive preference orderings. Also similar to Study 1, most subjects displayed ambiguity aversion, preferring the *A-lottery*

Choice patterns					Study 1		Study 2	
$L_A \succ L_{EA}$	$L_A \succ L_E$	$L_{EA} \succ L_E$	Preference ordering	Interpretation	N	%	N	%
1	1	1	$L_A \succ L_{EA} \succ L_E$	consistent ambiguity aversion epistemic uncertainty aversion	74	37%	175	31%
1	1	0	$L_A \succ L_E \succ L_{EA}$	consistent ambiguity aversion compound lottery aversion	39	19%	140	25%
0	0	0	$L_E \succ L_{EA} \succ L_A$	consistent ambiguity seeking epistemic uncertainty seeking	16	8%	69	12%
0	0	1	$L_{EA} \succ L_E \succ L_A$	consistent ambiguity seeking compound lottery seeking	17	9%	39	7%
0	1	1	$L_{EA} \succ L_A \succ L_E$	no consistent ambiguity attitude compound lottery seeking	16	8%	42	7%
1	0	0	$L_E \succ L_A \succ L_{EA}$	no consistent ambiguity attitude compound lottery aversion	19	10%	47	8%
0	1	0	intransitive		10	5%	30	5%
1	0	1	intransitive		9	4%	25	4%

Table 2: Preference orderings among lotteries in Studies 1 and 2. The first column indicates whether or not subjects in the corresponding row preferred the purely aleatory lottery to the mixed epistemic-aleatory lottery (i.e., 1 indicates $L_A \succ L_{EA}$; 0 indicates $L_{EA} \succ L_A$); the second column indicates whether or not subjects preferred the purely aleatory lottery to the purely epistemic lottery; the third column indicates whether or not subjects preferred the mixed epistemic-aleatory lottery to the purely epistemic lottery. The fourth column presents the resulting preference ordering among lotteries. The fifth column indicates the interpretations most consistent with these preference orderings. The last four columns provide the absolute and relative frequencies of corresponding choice patterns among subjects in Studies 1 and 2.

over the *EA-lottery* (68%; $p < 0.01$) and the *A-lottery* over the *E-lottery* (68%; $p < 0.01$). Again, most subjects (56%) exhibited consistent ambiguity aversion by preferring the *A-lottery* to both the *EA-lottery* and the *E-lottery*.

Replicating our key result from Study 1, the majority of consistently ambiguity averse subjects (who indicated both $L_A \succ L_{EA}$ and $L_A \succ L_E$; $N = 315$) also exhibited epistemic uncertainty aversion rather than compound lottery aversion (56% of these subjects indicated $L_{EA} \succ L_E$; $p < 0.05$). As in Study 1, this was the modal preference ordering (31% of all subjects). Meanwhile, the majority of consistently ambiguity seeking subjects ($L_{EA} \succ L_A$ and $L_E \succ L_A$; $N = 108$) also exhibited epistemic uncertainty seeking behavior rather than compound lottery seeking (64% of these subjects indicated $L_E \succ L_{EA}$; $p < 0.01$). Finally, we note that the small proportion of transitive subjects without consistent ambiguity attitudes ($L_A \succ L_{EA}$ and $L_E \succ L_A$, or $L_{EA} \succ L_A$ and $L_A \succ L_E$; $N = 89$) exhibit no significant preference between the purely epistemic lottery and the mixed-uncertainty lottery (47% of these subjects indicated $L_{EA} \succ L_E$; $p = 0.40$).[13]

[13] As in Study 1, we also conducted the inverse analysis involving the proportion of subjects exhibiting classic ambiguity aversion ($L_A \succ L_{EA}$) conditional on a consistent preference for or against the purely epistemic lottery. Among subjects with a consistent distaste for the purely epistemic lottery ($L_A \succ L_E$ and

For Studies 1 and 2, we also examined the proportion of preference orderings, among those with consistent ambiguity attitudes, that were uniquely consistent with our account concerning sensitivity to epistemic uncertainty (i.e., $L_A \succ L_{EA} \succ L_E$ or $L_E \succ L_{EA} \succ L_A$) to preference orderings uniquely consistent with the notion of sensitivity to compound lotteries (i.e., $L_A \succ L_E \succ L_{EA}$ or $L_{EA} \succ L_E \succ L_A$). That is, we compared the choice proportions in rows 1 and 3 to choice proportions in rows 2 and 4 of Table 2. For both studies a significantly larger percentage of choices are consistent with sensitivity to epistemic uncertainty than sensitivity to compound lotteries (Study 1: 45% vs 28%, $p < 0.01$; Study 2: 43% vs 32%, $p < 0.01$).

3.3 Study 3: Naturalistic Bets

Studies 1 and 2 provided a critical test of the present account against prominent economic models that cannot accommodate the simultaneous observation of ambiguity aversion and compound lottery seeking, or ambiguity seeking and compound lottery aversion (which are predicted by none of the above models, except for REU and GSM). In Study 3 we develop a research design using analogous bets on an upcoming soccer match, and in which we also ask participants to rate their level of knowledge concerning the match. The experimental design provides three enhancements over the previous studies. First, we move from the highly stylized domain of balls and urns to naturalistic events on which it is common to bet outside the laboratory. Second, our analysis no longer relies on the assumption of symmetric priors, and instead examines preferences among complementary bets. Finally, we exploit natural variation in self-rated expertise to examine the hypothesis that preference for an aleatory hedge is stronger among subjects who feel relatively ignorant concerning the events in question.

A bet on which team will win an upcoming soccer match can be interpreted as a two-stage lottery much like a random draw from the unknown urn in the two-color Ellsberg paradigm. In the first stage, decision makers identify the team favored to win the game by assessing the prior probability of each team winning, analogous to choosing a color on which to bet by assessing one's prior over the composition of the unknown urn. The second stage represents the realization of a particular game outcome between the two teams, analogous to a random draw from the unknown urn. Because the first-stage exclusively depends on knowledge about the relative strength or skill of the two teams and how they match up (similar to determining

$L_{EA} \succ L_E$; $N = 217$), a significant majority exhibited classic ambiguity aversion (81% showed $L_A \succ L_{EA}$; $p < 0.01$). Among subjects with a consistent preference for the purely epistemic lottery ($L_E \succ L_A$ and $L_E \succ L_{EA}$, $N = 116$), a significant majority exhibited classic ambiguity seeking (59% showed $L_{EA} \succ L_A$; $p < 0.01$). These patterns again accord with the hypothesis that classic ambiguity aversion is associated with aversion toward epistemic uncertainty rather than compound lotteries.

the proportion of red to black balls in an urn), selecting which of two teams is currently favored to win by bookmakers entails purely epistemic uncertainty. In contrast, the second stage in which outcomes are realized is largely aleatory in nature because different outcomes can occur by chance — sometimes the weaker team prevails.[14]

We exploit these differences in uncertainty across stages to design lotteries that entail purely epistemic uncertainty or a mixture of epistemic and aleatory uncertainty. Prior to an upcoming soccer match in Germany between Cologne and Augsburg, we asked Münster residents to choose between one of the following two bets:

(f) €100 if Cologne is currently favored by bookmakers to win the upcoming match with Augsburg.

(g) €100 if Cologne wins the upcoming match with Augsburg.

We also asked subjects to choose between two additional bets involving events that are complementary to the first pair:

(f') €100 if Augsburg or neither team is currently favored by bookmakers to win the upcoming match with Cologne.

(g') €100 if Augsburg wins or ties the match with Cologne.

Let p be the subjective probability that bookmakers currently favor Cologne, and q be the subjective probability that Cologne wins the match. It follows under SEU that $1-p$ is the subjective probability that bookmakers currently favor Augsburg or neither team, and $1-q$ is the subjective probability that Augsburg wins or ties the match. It is clear that under SEU, $f \succ g$ iff $g' \succ f'$ and thus, that SEU can accommodate only two of the four possible preference patterns in which subjects bet on one team to be favored and the other team to win (f and g', or g and f').

Our first prediction is that subjects who rate themselves more knowledgeable about the upcoming game are more likely to bet consistently with SEU. For instance, if a high-knowledge subject strongly believes that Cologne is currently favored by bookmakers over Augsburg, then she should bet on Cologne being favored (f) rather than Cologne winning the match (g), since soccer game outcomes are partly random and Augsburg could potentially win or tie the match. By the same logic, this individual should also bet on Augsburg winning or tying the match (g') rather than Augsburg or neither team being currently favored by bookmakers (f'), since she strongly believes the latter to be false.

[14]Of course, bookmaker odds do not necessarily capture the objective prior odds over the outcome of the soccer match, but they are the best available and most canonical proxy that is seen as a knowable (epistemic) and not particularly random (aleatory) event.

Our more central prediction concerns the betting behavior of less knowledgeable subjects. Decision makers can violate SEU in two distinct ways. They can choose to bet twice on the team currently favored by bookmakers (f and f'), and thereby exhibit a preference for the single-stage, purely epistemic lotteries. Alternatively, subjects can bet twice on the outcome of the upcoming game (g and g') and thereby exhibit a preference for two-stage lotteries that offer an aleatory hedge. While the former combination ($f \succ g$ and $f' \succ g'$) is consistent with compound lottery aversion, the latter combination ($g \succ f$ and $g' \succ f'$) is consistent with aversion to epistemic uncertainty. Thus, we predict that most subjects who violate SEU will bet twice on the game rather than twice on which team is currently favored.

We conducted a paper-and-pencil based survey at a local government-run citizen center in Münster, Germany ($N = 721$).[15] Subjects were offered a chocolate bar as compensation for completing the survey and were also informed that some respondents would be selected at random to have one of their (randomly-selected) bets played for real money. We distributed gift cards worth €600 in total among subjects who were selected and won their bet. Specifically, subjects chose between (f) "Win €100 gift card if the current betting odds on bet365.com say that [Cologne] is favored to win the game against [Augsburg] on November 26," or (g) "Win €100 gift card if [Cologne] wins the game against [Augsburg] on November 26." For the second bet, subjects chose between (f') "Win €100 gift card if the current betting odds on bet365.com say that [Augsburg] is favored to win or draw the game against [Cologne] on November 26,"), or (g') "Win €100 gift card if [Augsburg] wins or draws the game against [Cologne] on November 26."[16] Half of subjects completed the survey in the order described above, while the other half of subjects completed the survey in the reverse order (in both cases, bets were shown on separate pages).[17] After choosing between bets, subjects rated how

[15]The citizen center is an attractive location to recruit survey subjects for two reasons. First, all residents must come to the citizen center to get their IDs, passports, or file their changes of address, enabling us to sample a broad pool of the general public. Second, visitors of the citizen center usually face considerable waiting times which made our study a welcome distraction, ensuring high participation rates.

[16]The quotes above are translated from the original German. We note that the two teams (Cologne and Augsburg) were not popular among our sample (which was drawn from residents of Münster). This was by design to minimize the likelihood that betting behavior would be driven primarily by fan loyalty. In fact, none of the subjects in our sample indicated that they rooted for either Cologne or Augsburg. Second, it was unlikely that subjects would know definitively which team was favored, as the betting odds at that time were fairly even. While Cologne was favored to win the game (payoff multiplier in case of winning the bet was 1.72), a tie was also considered quite likely (3.50), and even an Augsburg win was still imaginable among bookmakers (5.00). The game was to be played about two weeks after we conducted the survey.

[17]For half of subjects, we included the following statement after describing the betting context and before presenting the actual bets: "We chose these bets because the vast majority of citizens in Münster report that they are familiar with these two teams." Emphasizing the existence of a relevant peer group that is well-informed regarding the decision context is an effective tool to induce a comparative ignorance effect (Fox and Tversky, 1995). Our results indicate that including the comparison with a well-informed peer group did indeed have the desired effect: subjects who were exposed to the comparison reported lower subjective knowledge than their unexposed counterparts (means were 2.31 vs. 2.68 on a 7-point scale;

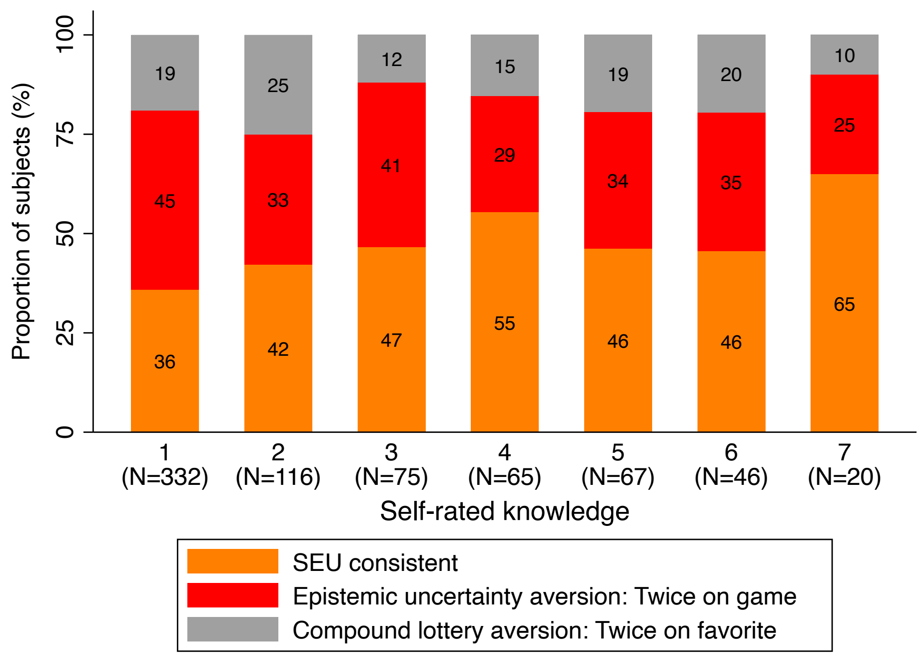

Figure 3: Impact of self-rated knowledge on subjects' betting behavior in Study 3. The figure displays betting behavior across self-rated knowledge levels. For each knowledge level, it provides the percentage of subjects who bet consistently with SEU preferences (by betting once on the game, and once on the favorite), bet inconsistently and in line with epistemic uncertainty aversion (by betting twice on the outcome of the game), or bet inconsistently in line with compound lottery aversion (by betting twice on the favorite).

knowledgeable they felt about the upcoming soccer match on a scale ranging from 1 ("not at all") to 7 ("very much"). The mean knowledge rating in our sample was 2.50 with a standard deviation of 1.80. About half of the survey participants were female (52%) and their average age was 30 years old.

Figure 3 displays the percentage of response profiles consistent with SEU, epistemic uncertainty aversion, and compound lottery aversion as a function of self-rated knowledge. Overall, 42% of participants provided responses consistent with SEU (by betting once on the current favorite and once on the upcoming game) and, as expected, this tendency increased as a function of self-rated knowledge. Based on a multinomial probit model, with choice patterns regressed on subjective knowledge, the likelihood of providing SEU-consistent responses increased by an average of 3.4 points for each one-point increase in self-rated knowledge ($p < 0.01$ based on the average marginal effect). As can be seen in Figure 3, the probability of betting consistently with SEU increased by more than half when comparing the highest-knowledge group of subjects (those rating their knowledge a 7 out of 7) to the

$p < 0.01$). Because the choice patterns across conditions did not otherwise differ significantly, we do not distinguish between them in the results that follow.

lowest-knowledge group of subjects (those rating their knowledge a 1 out of 7).

Our second and more central prediction concerns whether subjects who violated SEU prefer single-stage but purely epistemic lotteries (f and f') to mixed-uncertainty but compound lotteries (g and g'). Consistent with the epistemic uncertainty aversion hypothesis (and contrary to compound lottery aversion), 68% of SEU-inconsistent subjects preferred to bet on both sides of the upcoming match (i.e., on mixed-uncertainty, compound lotteries), while only 32% of subjects preferred to bet on both sides of the team currently favored by bookmakers (i.e., on purely epistemic, single-stage lotteries; $p < 0.01$ by a binomial test). Furthermore, and consistent with the epistemic uncertainty aversion hypothesis, the tendency to bet twice on the upcoming match decreased as a function of self-rated knowledge ($p < 0.01$ based on the average marginal effect). Meanwhile, we observe no significant correlation between self-rated knowledge and the tendency to bet twice on which team was currently favored by bookmakers ($p = 0.43$ based on the average marginal effect), indicating that compound lottery averse choices were not reliably related to decision maker knowledge.

3.4 Study 4: Violations of Dominance

Our central thesis is that ambiguity aversion reflects a distaste for acting on one's relative ignorance, to the extent that uncertainty is seen as epistemic in nature. In Study 4 we exogenously vary feelings of relative ignorance and examine its effect on preference for an aleatory hedge. Study 4 also examines whether epistemic uncertainty aversion can induce decision makers to hedge against their ignorance by choosing stochastically dominated alternatives that add aleatory uncertainty.

We employ two-alternative forced choice trivia questions (i.e., questions in which one of two possible answers is correct) as target events that entail purely epistemic uncertainty. In particular, we asked subjects to choose between a lottery where outcomes depend entirely on correctly answering the trivia question versus a lottery where outcomes depend on both trivia performance and chance. In this setup, the second option is stochastically dominated by the first option, and only the second option entails a compound lottery.[18] Thus, strict preferences for the second option are inconsistent with both SEU and compound lottery aversion. Moreover, it can be shown that a strict preference for the second option cannot be accommodated by the other prominent models of decision under uncertainty discussed above (CEU, RRDU, Source models with SPS), except for REU and GSM.[19] However, to the extent

[18]To illustrate why the compound lottery is stochastically dominated by the simple lottery, consider a decision maker with a subjective probability of 0.80 that she correctly answered the trivia question. This decision maker would face a choice between a (subjective) 0.80 chance of winning €100 or a $(0.80 \times 0.90) + (0.20 \times 0.10) = 0.74$ chance of winning €100.

[19]In case of GSM, the explicit compound lottery would be treated as a distinct source of uncertainty from

	Alternative A	Alternative B
Description:	If your answer to the geography question above is right, you will win a €100 gift card. If your answer to the geography question above is wrong, you will win nothing.	If your answer to the geography question above is right, you will receive a 90% chance to win a €100 gift card. If your answer to the geography question above is wrong, you will receive a 10% chance to win a €100 gift card.
Dimension of uncertainty:	Epistemic	Epistemic and Aleatory
Stage type:	Single-stage	Compound

Table 3: Menu of alternatives employed in Study 4. After answering their respective trivia questions, subjects in the easy question and hard question treatments were asked to choose between the lotteries listed above. The table also lists the type of uncertainty and stage type for each lottery.

that individuals are reluctant to bet on their own ignorance under epistemic uncertainty, we expect that less confident subjects will find the second option more attractive.[20] In addition, we experimentally vary the difficulty of the trivia question to causally identify whether feelings of relative ignorance drive an aversion to epistemic uncertainty.

We recruited 132 advanced undergraduate students enrolled in a course on rational decision making at the University of Münster, who had received some instruction on formal approaches to decision under uncertainty prior to completing this study. At the beginning of the paper-and-pencil based survey, which was administered in class, we informed subjects that some respondents would be selected at random to receive the outcome of their preferred lottery for real money. Half the subjects were randomly assigned to the *easy question* treatment and chose which of two German states, Baden-Württemberg or North Rhine-Westphalia, was larger in square kilometers. The other half of subjects were assigned to the *difficult question* treatment and chose whether Schleswig-Holstein or Thuringia was larger in square kilometers.[21] After providing their answer to the geography question, subjects in both treatments chose between the two options displayed in Table 3.

For Alternative A, winning the €100 gift card depended entirely on answering the geography question correctly, and as such entailed purely epistemic uncertainty. For Alternative B, subjects who answered correctly would win the gift card with 90% probability and subjects

the purely epistemic lottery.

[20] Again, this prediction can be accomodated only by REU and GSM.

[21] Schleswig-Holstein and Thuringia are smaller in size, less populated, and generally less familiar than either German state in the easy question treatment. The two states in our difficult question treatment are also closer to each other in geographic size than Baden-Württemberg and North Rhine-Westphalia. We thus expected these features to make for a more difficult general knowledge question. Confirming our a priori expectation, the average judged probability that a subject answered correctly was lower for the difficult question treatment than for the easy question treatment (mean judged probabilities of answering correctly were 0.65 and 0.76, respectively; $p < 0.01$).

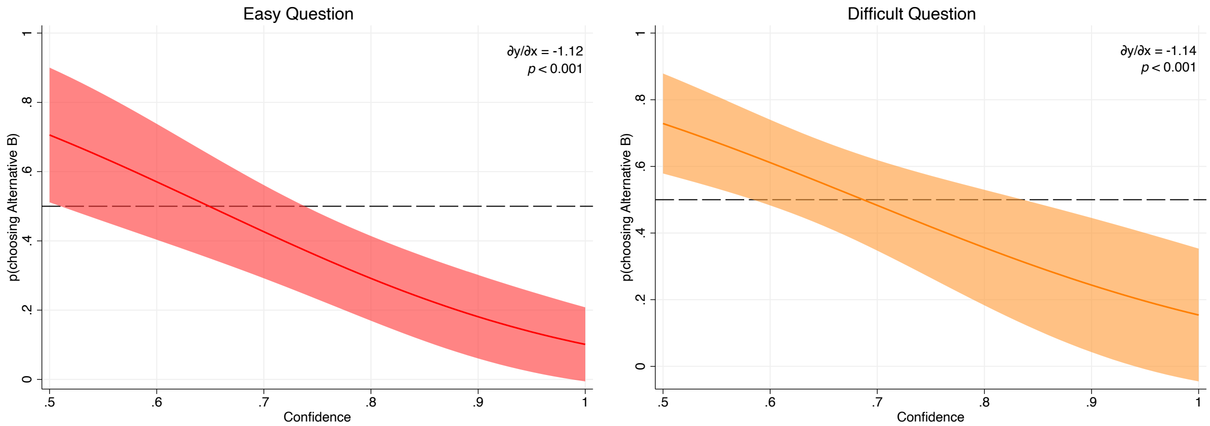

Figure 4: Violations of dominance as a function of confidence in Study 4. This figure displays the predicted probability of selecting the (dominated) compound lottery over the single-stage lottery as a function of confidence (i.e., judged probability of answering the trivia question correctly). The left panel plots this relationship for the easy question treatment, and the right panel plots this relationship for the difficult question treatment. Predicted probabilities are calculated based on the average marginal effects from a probit regression. Error bands represent 95% confidence intervals.

who answered incorrectly would win the gift card with 10% probability. Due to the randomness in the second stage, Alternative B represents a compound lottery that entails both epistemic and aleatory uncertainty. Finally, subjects provided their subjective probability of having answered the trivia question correctly (from 50% to 100%). Note that, when reduction of compound lotteries applies, opting for Alternative B strictly violates first-order stochastic dominance for any subjective probability greater than 0.50.

Results of Study 4 accord with the epistemic uncertainty aversion hypothesis. Subjects were more likely to choose the compound lottery (Alternative B) when responding to a difficult question (55%) than when responding to an easy question (36%), $z = 2.10$, $p < 0.05$. Note that a preference for Alternative B (46% of subjects across conditions) violates both SEU and most prominent models of ambiguity aversion, regardless of question difficulty. We also find, as predicted, that in both treatments the preference for Alternative B increases as confidence decreases (see Figure 4). Based on a probit model, with choices regressed on confidence, a one-point decrease in subjective probability of answering the geography question correctly resulted in an average 1.1-point increase in choosing Alternative B (for both easy and difficult questions; in both cases $p < 0.01$ based on the average marginal effects).

Because subjects were randomly assigned to easy and difficult questions, we can obtain direct causal estimates of how confidence affects a preference for the aleatory hedge (Alternative B). Using an instrumental variables probit model, with confidence instrumented on question difficulty, we find that on average a one-point decrease in confidence leads to a 1.6 percentage point increase in selecting Alternative B ($p < 0.01$ based on the average marginal effect). This suggests that increasing ignorance is associated with a stronger preference for compound lotteries that provide an aleatory hedge.[22]

3.5 Study 5: Locating the Aleatory Hedge in the Past vs. Future

Our thesis in this paper is that ambiguity aversion reflects a distaste for acting on one's relative ignorance to the extent that uncertainty is seen as epistemic in nature. In Study 4 we experimentally varied question difficulty and find that relative ignorance amplifies epistemic uncertainty aversion. In Study 5 we experimentally vary whether a compound lottery stage is treated as more epistemic (versus more aleatory) influences the attractiveness of an aleatory hedge among participants who view themselves as relatively ignorant. In particular, we examine whether attractiveness of the aleatory hedge diminishes when chance is located in the past, prior to the epistemic stage of the lottery (rather than in the future, following resolution of the epistemic stage).

Past behavioral research suggests that people prefer betting on random events that have not yet been resolved to random events whose results are already resolved but unknown. For instance, Rothbart and Snyder (1970) found that subjects are willing to bet more money on the outcome of a die that has yet to be rolled than one that has already been rolled but whose outcome is unknown (for related results see Brun and Teigen, 1990; Heath and Tversky, 1991). One interpretation of this result is that random events that have already been resolved, even if unknown, take on an appearance of epistemic uncertainty in a way that future random events do not. According to our hypothesis, then, locating the aleatory component of a mixed-uncertainty prospect in the past should make that lottery appear less attractive to decision makers. Such a timing-related dependency of the value of the aleatory hedge is consistent with theoretical predictions of Saito (2015), Ke and Zhang (2020), and GSM, however cannot be accommodated by REU.

Our experimental intervention locates the aleatory hedge in the past, which affects not only stage timing (randomization takes place in the past versus future) but also stage ordering (aleatory hedge represents the first versus second stage of the lottery). Note that for compound lotteries with independent stages, such as those we examine here, models of decision under uncertainty and ambiguity typically imply that the ordering of lottery stages is irrelevant. In particular, Anscombe and Aumann's (1963) axiomatization of SEU explicitly assumes

[22] It is worth acknowledging that the degree of stochastic dominance also increases as subjective probability of answering correctly increases (the percentage difference in subjective probability of winning the prize decreases by 0.20 for every 1 percentage point increase in confidence). Thus, the negative correlation between subjective confidence and choosing the dominated option may be driven in part by sensitivity to the degree of stochastic dominance. However, this factor cannot explain the systematic preference for dominated lotteries among participants who feel relatively ignorant (i.e., report probabilities close to 0.50).

that the order of lottery stages should not affect preferences (Assumption 2: Reversal of Compound Lotteries, p. 201). Thus, a diminished preference for the aleatory hedge when it is located in the first rather than second lottery stage would not only violate REU but also Anscombe and Aumann's (1963) second assumption.[23]

We recruited 500 subjects from an online labor market (`www.prolific.ac`) to participate in a brief study in exchange for a \$0.30 payment. Because our manipulation of timing is subtle and depends on keen attention of subjects, we preregistered recruitment of a large sample in order to allow us to only retain subjects who passed a rigorous comprehension check at the end of the experiment.[24] We randomly assigned subjects to one of two treatments. In the first treatment (*randomization after choosing*), we asked subjects to first answer which river they thought was longer: the Amazon or the Nile. They next estimated the probability of having answered correctly (from 50% to 100%) and then chose one of the following options:

> *Alternative A*
> If your answer to the trivia question above is <u>correct</u>, then you receive \$100.
> If your answer to the trivia question above is <u>incorrect</u>, then you receive nothing.
>
> *Alternative B*
> If your answer to the trivia question above is <u>correct</u> then an experimenter will roll a fair six-sided die. If the experimenter's die lands on 1, 2, 3, 4, or 5 then you receive \$100.
> If your answer to the trivia question above is <u>incorrect</u> then an experimenter will roll a fair six-sided die. If the experimenter's die lands on 6 then you receive \$100.

As in Study 4, Alternative A can be interpreted as a single-stage lottery that entails purely epistemic uncertainty. Meanwhile, Alternative B can be interpreted as a compound lottery that entails epistemic uncertainty in the first stage and aleatory uncertainty in the second stage. Note that, as was the case in Study 4, Alternative B is stochastically dominated

[23] In REU, more uncertain events are always resolved first, followed by resolving the less uncertain events (for a discussion see He, 2021). Hence, reversing the stages of a compound lottery that differ in the degree of uncertainty (aleatory vs. epistemic) does not change the fact that an REU decision maker treats such a prospect as a "horse-roulette" (i.e., epistemic-aleatory) lottery. Changing the order of lottery stages therefore cannot alter the attractiveness of the aleatory hedge in REU.

[24] At the end of the study subjects responded to three comprehension questions: (1) "Did the option with the dice roll involve a die thrown by the experimenter before or after you answered the trivia question?" (2) "Did the previous task involve an option where you could win \$100 even if you answered the trivia question incorrectly?" (3) "What was the probability that participants who chose the second option would win \$100 if they answered the trivia question correctly?" We excluded subjects who answered any of these comprehension questions incorrectly. Also, since our experiment involved a trivia question that could be answered through an internet search, we implemented a second exclusion criteria that subjects not browse away at any point during the experiment (we used software to track browsing behavior; Permut et al., 2019). Our final sample based on these criteria consisted of responses from 149 subjects.

by Alternative A for subjective probabilities greater than 0.50, assuming that reduction of compound lotteries applies.

The second treatment (*randomization before choosing*) was identical to the first treatment except that we reversed the ordering of the lottery stages for Alternative B. We first informed subjects about the experimenter's dice roll as follows:

> An experimenter rolls a fair six-sided die and privately observes what number the die landed on. If the die landed on 1, 2, 3, 4, or 5 then the experimenter writes down 'Correct.' If the die landed on 6 then the experimenter writes down 'Incorrect.' You do not know what the experimenter wrote down.

After receiving this information subjects then answered the same trivia question involving the length of the Amazon and Nile and estimated the probability of having answered the trivia question correctly (from 50% to 100%). They next chose one of the following options:

> *Alternative A*
> If your answer to the trivia question above is correct, then you receive $100.
> If your answer to the trivia question above is incorrect, then you receive nothing.
>
> *Alternative B*
> If the experimenter wrote down "Correct" then you receive $100 only if you answered correctly.
> If the experimenter wrote down "Incorrect" then you receive $100 only if you answered incorrectly.

Note that the choice menu for the second treatment involves the same probability distributions over outcomes as the first treatment. Importantly, we do not observe significant differences on responses to the trivia question as a function of experimental treatment ($z = 1.12$, $p = 0.265$), nor do we see significant differences in confidence of having answered the question correctly[25] ($t_{147} = 0.15$, $p = 0.880$).

Results accord with the epistemic uncertainty aversion hypothesis. Subjects were more likely to choose Alternative B when the aleatory hedge was located as a second-stage lottery in the future (46%) than as a first-stage lottery in the past (30%), $z = 1.92$, $p = 0.055$. Note that as in Study 4 any subject who prefers Alternative B (40% of subjects across treatments) violates both SEU and compound lottery aversion. Additionally, we find that preference for Alternative B generally increases as confidence decreases (see Figure 5).

In the randomization-after-choosing treatment, which closely resembles the choice menu from Study 4, a one-point decrease in confidence resulted on average in a 0.8-point increase of

[25] We also find a null effect when examining differences in the distributions of confidence ratings across the two conditions ($p = 0.697$ by a Kolmogorov-Smirnov test).

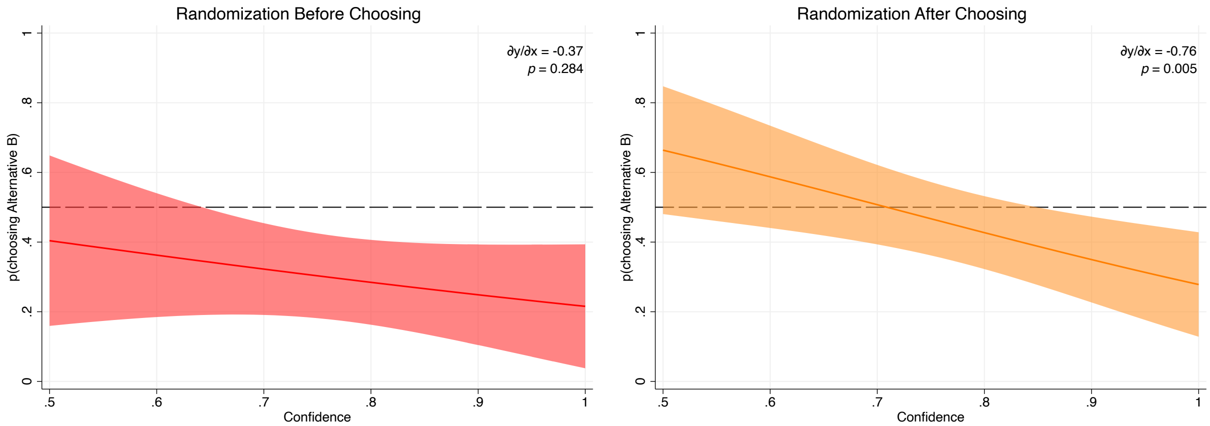

Figure 5: Violations of dominance as a function of confidence in Study 5. This figure displays the predicted probability of selecting the (dominated) compound lottery over the single-stage lottery as a function of confidence (judged probability of answering the trivia question correctly). The left panel plots this relationship in the randomization-before-choosing treatment, and the right panel plots this relationship for the randomization-after-choosing treatment. Predicted probabilities are calculated based on the average marginal effects from a probit regression. Error bands represent 95% confidence intervals.

choosing Alternative B ($p < 0.01$ based on the average marginal effect). In the randomization-before-choosing treatment, however, the effect of confidence on choice is no longer statistically significant ($p = 0.284$ based on the average marginal effect). It worth noting that the dampening effect for the second treatment compared to the first is consistent with the interpretation that random events located in the past are construed as more epistemic, thereby making Alternative B less attractive to low-confidence subjects because it is less likely to be seen as an aleatory hedge. This said, a post hoc comparison of the coefficients for the two average marginal effect are not statistically different from one another ($p = 0.379$), so this interpretation should be treated as tentative.

3.6 Study 6: Reframing the Aleatory Hedge as Epistemic

In our previous study, we attempted to influence the impact of the aleatory hedge by varying the temporal order of the stages involved in the compound lottery (Alternative B). In Study 6, we keep the temporal ordering of the compound lottery's stages fixed and instead attempt to influence the impact of the aleatory hedge by reframing the description of Alternative B to either highlight or obscure the aleatory component of the lottery. We expected that reframing the compound lottery in a way that integrates the chance component with the outcome of the answer to the trivia question should neutralize the attractiveness of the aleatory hedge. Alternative B should thus be viewed as less attractive when framed in this manner, especially among subjects who feel relatively ignorant about the trivia question. Among the models reviewed in this paper, only GSM can accomodate the difference in attractiveness of the compound lottery under different frames.

We recruited 404 subjects from an online labor market (`www.mturk.com`) to participate in a brief study in exchange for a $0.35 payment and also a chance to win a $100 prize for an unrelated task that came after the present study. We randomly assigned subjects to one of two treatments. In the first treatment (*standard frame*), we asked subjects to first answer which river they thought was longer between the Amazon and the Nile. Next, they choose between the following options:

> *Alternative A*
> If your answer to the question above is right, you win $200.
> If your answer to the question above is wrong, you win nothing.
>
> *Alternative B*
> If your answer to the question above is right, you have a 90% chance to win $200.
> If your answer to the question above is wrong, you have a 10% chance to win $200.

Finally, we asked subjects to estimate their probability of having answered the trivia question correctly (from 50% to 100%).

In the second treatment (*switch frame*), the subjects completed the same task except were offered the following two options after answering the trivia question:

> *Alternative A*
> If your answer to the question above is right, you win $200.
> If your answer to the question above is wrong, you win nothing.
>
> *Alternative B*
> There is a 10% chance that your answer to the question above will be switched, and a 90% chance your answer will remain the same. This will be your final answer.
>
> If your final answer is right, you win $200. If your final answer is wrong, you win nothing.

Note that options in the switch frame are extensionally equivalent to options in the standard frame. The only difference is that the description of Alternative B no longer explicitly segregates the aleatory hedge, but instead frames the chance element as part of the epistemic component of the lottery. Importantly, we do not observe significant differences on responses to the trivia question as a function of experimental treatment ($z = 0.20$, $p = 0.844$), nor

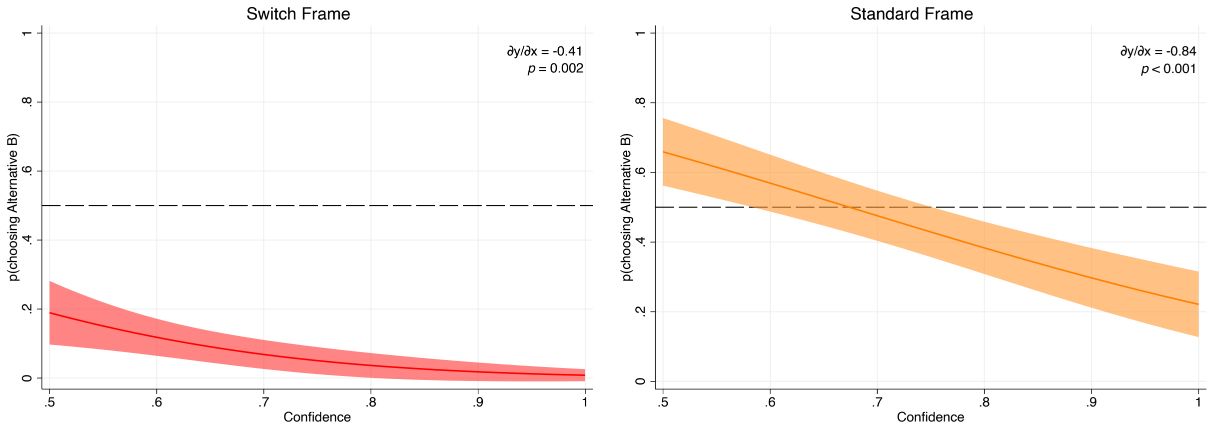

Figure 6: Violations of dominance as a function of confidence in Study 6. This figure displays the predicted probability of selecting the (dominated) compound lottery over the single-stage lottery as a function of confidence (judged probability of answering the trivia question correctly). The left panel plots this relationship for the switch frame treatment, and the right panel plots this relationship for the standard frame treatment. Predicted probabilities are calculated based on the average marginal effects from a probit regression. Error bands represent 95% confidence intervals.

do we see significant differences in confidence of having answered the question correctly[26] ($t_{402} = 1.45$, $p = 0.149$).

Results again accord with the epistemic uncertainty aversion hypothesis. Subjects were considerably more likely to choose Alternative B in the standard frame where the aleatory hedge is clearly segregated from the epistemic component of the compound lottery (47%) than in the switch frame where the chance component is integrated with the epistemic component of the compound lottery (8%), $z = 8.64$, $p < 0.01$. Note again that any subject preferring Alternative B (27% of subjects across treatments) violates both SEU and compound lottery aversion. Finally, we replicate the finding that preference for Alternative B increases as confidence decreases (see Figure 6). In the standard frame, a 1-point decrease in confidence resulted on average in a 0.8-point increase of choosing Alternative B ($p < 0.01$ based on the average marginal effect). In the switch frame, the effect of confidence on choice is dampened as in Study 5 but in this case remains statistically significant ($p < 0.01$ based on the average marginal effect). However, this time we find that the two average marginal effects are statistically different from one another ($p = 0.033$). This dampening effect for the second treatment is consistent with the interpretation that obscuring the chance component of Alternative B made the option especially less attractive to low-confidence subjects because it is less likely to be seen as an aleatory hedge.

[26] We also find a null effect when examining differences in the distributions of confidence ratings across the two conditions ($p = 0.496$ by a Kolmogorov-Smirnov test).

4 Discussion

We extend prior psychological accounts of decision under uncertainty to argue that ambiguity aversion reflects a distaste for betting on one's relative ignorance to the extent that events are seen as inherently knowable or *epistemic* in nature. In contrast, uncertainty viewed as random or *aleatory* in nature can provide an attractive hedge against betting on one's ignorance. We provide experimental evidence that: (1) under conditions of ignorance, ambiguity averse decision makers prefer betting on a greater balance of aleatory to epistemic uncertainty; (2) this preference for an aleatory hedge increases as subjective knowledge decreases, and can lead decision makers to choose stochastically dominated alternatives; and (3) an ambiguous prospect can be made more or less attractive depending on whether uncertainty is framed as more aleatory or epistemic in nature.

Our findings contradict prominent models of decision under uncertainty. However, the present results can be accommodated by generic source models with rank-dependent utility (such as prospect theory) that include probability weighting functions that vary by source of uncertainty (i.e., "source functions"), as discussed in Section 2.2.6. Of course, a more complete model would formalize the relationship between characteristics of sources of uncertainty and corresponding source functions. The present findings suggest that the elevation of such functions diminishes (i.e., pessimism increases) with the interaction of a decision maker's comparative ignorance and perceived balance of epistemicness to aleatoriness of the source. We leave a formalization of this account to future research.

In his seminal paper, Ellsberg (1961) noted that "ambiguity is a subjective variable" (p. 660). Likewise, experimental studies have shown that level of knowledge or competence that drives source preferences is subjective and context-dependent (Fox and Tversky, 1995; Fox and Weber, 2002; Hadar et al., 2013). For instance, Fox and Weber (2002) report that betting on a moderately familiar event (e.g., the winner of an upcoming election in Russia) is more attractive after decision makers are reminded of a *less* familiar event (e.g., the winner of an upcoming election in the Dominican Republic) than after they are reminded of a *more* familiar event (e.g., the winner of an upcoming election in the United States). Thus, a decision maker's subjective level of knowledge appears to be driven by a contrast with related events that the decision maker has recently been contemplating. These results suggest that decision makers' confidence in betting under ambiguity is more closely tethered to their subjective rather than objective level of knowledge (for more direct evidence see Hadar et al., 2013).[27] Furthermore, the last two studies in the current paper suggest that assessments of

[27]The interpretation that decision makers are especially wary of acting on their relative ignorance is also consistent with neuroeconomic studies that compare brain activation when participants make betting decisions under ambiguity versus risk, high versus low absolute knowledge, and high versus low knowledge

epistemicness and aleatoriness of a source of uncertainty, like assessments of competence or knowledge, are context-dependent and subject to framing effects. Whether a random event is located in the future or past, or whether it is segregated as a chance lottery or integrated with a knowable event, critically influences the attractiveness of betting on a source of uncertainty.

Another critical contextual factor is whether different sources of uncertainty are evaluated in a comparative or noncomparative choice setting. Fox and Tversky (1995) find that ambiguity aversion involving Ellsberg bets is pronounced when decision makers evaluate both risky and ambiguous bets simultaneously, but diminishes or disappears when urns are evaluated separately (so that there is no explicit contrast between urns to highlight one's comparative ignorance concerning the unknown probability urn). Likewise, it is worth noting that all of the studies in the present paper entail choices involving contrasting sources of uncertainty, for example, a choice between the *EA-lottery* and *E-lottery* in Studies 1 and 2. We speculate that the attractiveness of the aleatory hedge provided by the *EA-lottery* may be amplified by such contrasts and could be diminished when sources of uncertainty are evaluated separately (as when lotteries involving purely epistemic versus mixed uncertainty are priced separately), in a similar vein to Fox and Tversky (1995).

In this paper we have examined preferences for compound lotteries involving a simple chance component and an ambiguous component. We have not examined preferences for *objective* compound lotteries involving only chance components (e.g., Halevy, 2007). Interestingly, Baars and Goedde-Menke (2021) find that probability weighting distortions are more pronounced when subjects are less familiar with the games of chance underlying risky lotteries — a phenomenon which they call the "ignorance illusion." Relatedly, Armantier and Treich (2016) provide evidence that probabilities are more distorted for objective lotteries with higher complexity than for simpler lotteries, on par with distortions observed for ambiguous lotteries. These results may explain instances of compound lottery aversion and its occasional association with ambiguity aversion (e.g., Halevy, 2007; Aydogan et al., 2020), since the more complex, two-stage design of compound (objective) lotteries may make them more difficult to understand for some decision makers. This said, it is important to note that complexity aversion cannot explain the preference for aleatory hedges in the present studies. Purely epistemic lotteries are no more complex to describe than mixed-uncertainty lotteries in the case of the Ellsberg paradigm (Studies 1-2); betting on a favored team is no more complex than betting on the outcome of a game (Study 3); and the compound, mixed-uncertainty lottery that subjects favor is, in fact, *more* complex than the simple epistemic lottery (Study 4).

relative to another participant (Hsu et al., 2005). In particular, these authors find that level of (relative) ignorance in choices correlates positively with activation in the amygdala (which is generally associated with increased vigilance), and negatively with activation in the striatum (which is generally associated with reward).

The context-dependent character of both subjective knowledge and nature of uncertainty that mutually drive ambiguity aversion poses a significant challenge to formal modeling. This would require one to mathematically represent the state of mind of a decision maker — which may be affected by fleeting ruminations and associations — at the moment when a particular decision is made. Although such a formalization of epistemic uncertainty aversion is beyond the scope of this paper, we have provided empirical evidence that prominent models of ambiguity aversion are inadequate, noted generic source models as a promising scaffolding on which to construct a viable model, and pointed to features of these functions that track epistemic uncertainty aversion and more accurately capture the ambiguity aversion phenomenon.

References

Abdellaoui, M., Baillon, A., Placido, L., and Wakker, P. P. (2011). The rich domain of uncertainty: Source functions and their experimental implementation. *American Economic Review*, 101(2):695–723.

Abdellaoui, M., Klibanoff, P., and Placido, L. (2015). Experiments on compound risk in relation to simple risk and to ambiguity. *Management Science*, 61(6):1306–1322.

Alary, D., Gollier, C., and Treich, N. (2013). The effect of ambiguity aversion on insurance and self-protection. *Economic Journal*, 123(573):1188–1202.

Anscombe, F. J. and Aumann, R. J. (1963). A definition of subjective probability. *Annals of Mathematical Statistics*, 34(1):199–205.

Armantier, O. and Treich, N. (2016). The rich domain of risk. *Management Science*, 62(7):1954–1969.

Aydogan, I., Berger, L., Bosetti, V., and Liu, N. (2020). Three layers of uncertainty and the role of model misspecification.

Baars, M. and Goedde-Menke, M. (2021). Ignorance illusion in decisions under risk: The impact of perceived expertise on probability weighting. *Journal of Risk and Insurance*, forthcoming.

Baillon, A., Huang, Z., Selim, A., and Wakker, P. P. (2018). Measuring ambiguity attitudes for all (natural) events. *Econometrica*, 86(5):1839–1858.

Bernasconi, M. and Loomes, G. (1992). Failures of the reduction principle in an ellsberg-type problem. *Theory and Decision*, 32(1):77–100.

Bernoulli, D. (1738/1954). Exposition of a new theory on the measurement. *Econometrica*, 22(1):23–36.

Brun, W. and Teigen, K. H. (1990). Prediction and postdiction preferences in guessing. *Journal of Behavioral Decision Making*, 3(1):17–28.

Camerer, C. and Weber, M. (1992). Recent developments in modeling preferences: Uncertainty and ambiguity. *Journal of Risk and Uncertainty*, 5(4).

Chew, S. H., Miao, B., and Zhong, S. (2017). Partial ambiguity. *Econometrica*, 85(4):1239–1260.

Chow, C. C. and Sarin, R. K. (2002). Known, unknown, and unknowable uncertainties. *Theory and Decision*, 52(2):127–138.

Curley, S. P., Yates, J. F., and Abrams, R. A. (1986). Psychological sources of ambiguity avoidance. *Organizational Behavior and Human Decision Processes*, 38(2):230–256.

Dean, M. and Ortoleva, P. (2019). The empirical relationship between nonstandard economic behaviors. *Proceedings of the National Academy of Sciences*, 116(33):16262–16267.

Dow, J. and da Costa Werlang, S. R. (1992). Uncertainty aversion, risk aversion, and the optimal choice of portfolio. *Econometrica*, pages 197–204.

Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms. *The Quarterly Journal of Economics*, 75(4):643–669.

Ergin, H. and Gul, F. (2009). A theory of subjective compound lotteries. *Journal of Economic Theory*, 144(3):899–929.

Fox, C., Tannenbaum, D., Ülkümen, G., Walters, D., and Erner, C. (2020). Credit, blame, luck, and the perceived nature of uncertainty. Unpublished manuscript.

Fox, C. R. and Tversky, A. (1995). Ambiguity aversion and comparative ignorance. *The Quarterly Journal of Economics*, 110(3):585–603.

Fox, C. R. and Tversky, A. (1998). A belief-based account of decision under uncertainty. *Management science*, 44(7):879–895.

Fox, C. R. and Ülkümen, G. (2011). Distinguishing two dimensions of uncertainty. In W, B., G, K., G, K., and H, M., editors, *Perspectives on Thinking, Judging, and Decision Making*, pages 21–35. Universitetsforlaget, Oslo.

Fox, C. R. and Weber, M. (2002). Ambiguity aversion, comparative ignorance, and decision context. *Organizational Behavior and Human Decision Processes*, 88(1):476–498.

Frisch, D. and Baron, J. (1988). Ambiguity and rationality. *Journal of Behavioral Decision Making*, 1(3):149–157.

Gilboa, I. (1987). Expected utility with purely subjective non-additive probabilities. *Journal of Mathematical Economics*, 16(1):65–88.

Gilboa, I. and Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153.

Gollier, C. (2011). Portfolio choices and asset prices: The comparative statics of ambiguity aversion. *Review of Economic Studies*, 78(4):1329–1344.

Hacking, I. (2006). *The Emergence of Probability: A Philosophical Study of Early Ideas about Probability, Induction and Statistical Inference*. Cambridge University Press, Cambridge, UK.

Hadar, L., Sood, S., and Fox, C. R. (2013). Subjective knowledge in consumer financial decisions. *Journal of Marketing Research*, 50(3):303–316.

Halevy, Y. (2007). Ellsberg revisited: An experimental study. *Econometrica*, 75(2):503–536.

He, Y. (2021). Revisiting ellsberg's and machina's paradoxes: A two-stage evaluation model under ambiguity. *Management Science*, forthcoming.

Heath, C. and Tversky, A. (1991). Preference and belief: Ambiguity and competence in choice under uncertainty. *Journal of Risk and Uncertainty*, 4(1):5–28.

Hogarth, R. M. and Kunreuther, H. (1985). Ambiguity and insurance decisions. *American Economic Review*, 75(2):386–390.

Hsu, M., Bhatt, M., Adolphs, R., Tranel, D., and Camerer, C. F. (2005). Neural systems responding to degrees of uncertainty in human decision-making. *Science*, 310(5754):1680–1683.

Ke, S. and Zhang, Q. (2020). Randomization and ambiguity aversion. *Econometrica*, 88(3):1159–1195.

Keynes, J. M. (1921). *A Treatise on Probability*. Macmillan, London.

Klibanoff, P., Marinacci, M., and Mukerji, S. (2005). A smooth model of decision making under ambiguity. *Econometrica*, 73(6):1849–1892.

Knight, F. H. (1921). *Risk, Uncertainty and Profit*. Cambridge University Press, Cambridge, UK.

Krijnen, J. M., Ülkümen, G., Bogard, J., and Fox, C. R. (2021). Lay theories of financial well-being predict political and policy message preferences. *Journal of Personality and Social Psychology*.

Machina, M. J. and Schmeidler, D. (1992). A more robust definition of subjective probability. *Econometrica*, 60(4):745–780.

Machina, M. J. and Siniscalchi, M. (2014). Ambiguity and ambiguity aversion. In *Handbook of the Economics of Risk and Uncertainty*, volume 1, pages 729–807. North-Holland.

Mukerji, S. (1998). Ambiguity aversion and incompleteness of contractual form. *American Economic Review*, 88(5):1207–1231.

Mukerji, S. and Tallon, J. (2001). Ambiguity aversion and incompleteness of financial markets. *Review of Economic Studies*, 68(4):883–904.

Muthukrishnan, A., Wathieu, L., and Xu, A. J. (2009). Ambiguity aversion and the preference for established brands. *Management Science*, 55(12):1933–1941.

Pascal, B. (1852). *Pensées*. Dezobry et E. Magdeleine.

Permut, S., Fisher, M., and Oppenheimer, D. M. (2019). Taskmaster: A tool for determining when subjects are on task. *Advances in Methods and Practices in Psychological Science*, 2(2):188–196.

Pulford, B. D. and Colman, A. M. (2007). Ambiguous games: Evidence for strategic ambiguity aversion. *Quarterly Journal of Experimental Psychology*, 60(8):1083–1100.

Quiggin, J. (1982). A theory of anticipated utility. *Journal of Economic Behavior and Organization*, 3(4):323–343.

Ramsey, F. P. (1931). Truth and probability. In *The Foundations of Mathematics and Other Logical Essays*, pages 156–198. Routledge and Kegan Paul.

Ritov, I. and Baron, J. (1990). Reluctance to vaccinate: Omission bias and ambiguity. *Journal of Behavioral Decision Making*, 3(4):263–277.

Rothbart, M. and Snyder, M. (1970). Confidence in the prediction and postdiction of an uncertain outcome. *Canadian Journal of Behavioural Science*, 2(1):38–43.

Saito, K. (2015). Preferences for flexibility and randomization under uncertainty. *American Economic Review*, 105(3):1246–71.

Savage, L. J. (1954). *The Foundations of Statistics*. Wiley, New York.

Schmeidler, D. (1989). Subjective probability and expected utility without additivity. *Econometrica*, 57(3):571–587.

Segal, U. (1987). The ellsberg paradox and risk aversion: An anticipated utility approach. *International Economic Review*, 28(1):175–202.

Segal, U. (1990). Two-stage lotteries without the reduction axiom. *Econometrica*, 58(2):349–377.

Tannenbaum, D., Fox, C. R., and Ülkümen, G. (2017). Judgment extremity and accuracy under epistemic versus aleatory uncertainty. *Management Science*, 63(2):497–518.

Trautmann, S. T. and Van De Kuilen, G. (2015). Ambiguity attitudes. In *The Wiley Blackwell Handbook of Judgment and Decision Making*, volume 1, pages 89–116. John Wiley & Sons, Ltd.

Trautmann, S. T., Vieider, F. M., and Wakker, P. P. (2008). Causes of ambiguity aversion: Known versus unknown preferences. *Journal of Risk and Uncertainty*, 36(3):225–243.

Tversky, A. and Fox, C. R. (1995). Weighing risk and uncertainty. *Psychological Review*, 102(2):269–283.

Tversky, A. and Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5(4):297–323.

Ülkümen, G., Fox, C. R., and Malle, B. F. (2016). Two dimensions of subjective uncertainty: Clues from natural language. *Journal of Experimental Psychology: General*, 145(10):1280–1297.

Vives, M. and Feldman Hall, O. (2018). Tolerance to ambiguous uncertainty predicts prosocial behavior. *Nature Communications*, 9(1):1–9.

von Neumann, J. and Morgenstern, O. (1944). *Theory of Games and Economic Behavior*. Princeton university press, Princeton, NJ.

Wakker, P. P. (1987). Nonadditive probabilities and derived strenghts of preferences. unpublished manuscript.

Wakker, P. P. (2004). On the composition of risk preference and belief. *Psychological Review*, 111(1):236–241.

Walters, D. J., Ülkümen, G., Erner, C., Tannenbaum, D., and Fox, C. R. (2020). Investment behavior under epistemic versus aleatory uncertainty. Unpublished manuscript.