

# Ambiguous Correlation\*

Larry G. Epstein      Yoram Halevy

December 14, 2017

## Abstract

Many decisions are made in environments where outcomes are determined by the realization of multiple random events. A decision maker may be uncertain how these events are related. We identify and experimentally substantiate behavior that intuitively reflects a lack of confidence in their joint distribution. Our findings suggest a dimension of ambiguity which is different from that in the classical distinction between risk and “Knightian uncertainty.”

---

\*Boston University, lepstein@bu.edu and University of Toronto, yoram.halevy@utoronto.ca. We are grateful for discussions with Kyoungwon Seo and Chew Soo Hong, for detailed comments from Aurelien Baillon and especially Peter Wakker, for insightful and thoughtful suggestions by two referees and the editor - Dimitri Vayanos, and for comments from audiences in several conferences and workshops. We are grateful also for the financial support of the National Science Foundation award SES-1216339 (Epstein) and the Social Science and Humanities Research Council of Canada grant 435-2012-0577 (Halevy). Anujit Chakraborty, David Freeman, Terri Kneeland and students from UBC’s ECON490 classes of 2013, 2014 and 2015 (Experimental Economics section) provided excellent research assistance. Earlier versions of this paper were titled “No Two Experiments Are Identical.” First version: February 2014.

# 1 Introduction

Individuals often make decisions when there is limited information about the stochastic environment and hence where there may be incomplete confidence that any single given probability law accurately describes it. Ellsberg (1961) identifies behavior that intuitively reflects such a lack of confidence, and shows thereby that the distinction between risk (where information is perfect and confidence is complete) and ambiguity is empirically meaningful. This paper considers a setting where uncertainty is generated by the realization of multiple random events and we study ambiguity about how these random events differ or are related to one another. Paralleling Ellsberg, we identify behavior that intuitively reflects a lack of confidence in (or uncertainty about) their joint distribution. Then we conduct a controlled incentivized laboratory experiment the results of which support the empirical significance of this new dimension of ambiguity.

In Ellsberg's (1961) classic two-urn thought experiment, there are two urns, each containing 100 balls that are either red or black. You are told the exact color composition for one (50-50), the risky urn, and nothing at all about the composition of the other one, the ambiguous urn. Intuition suggests, and many subsequent laboratory experiments confirm, that many people prefer to bet on drawing red (black, respectively) from the risky urn as opposed to from the ambiguous urn, thereby demonstrating a behavioral distinction between risk (known composition) and Knightian uncertainty (unknown bias or composition).

We consider a setting where there are two urns and you are told the same about the composition of each, so that you have no reason to distinguish between them. However, you are told very little; for example, only that each urn contains two balls each of which is either red or black. Accordingly, you are not given any reason to be certain that the compositions are identical, nor are you given any reason for being confident that the urns' compositions are unrelated or related in any particular way. We study choices between bets on the colors of two balls, where simultaneously one ball is drawn from each urn, and we identify behavior that intuitively reveals a lack of confidence concerning the relation between the urns. The key idea is that ambiguity about how the compositions of the two urns might differ or be related is not relevant to bets on a single urn. This leads us to focus on the choice between bets on the color of the ball drawn from one urn versus bets on the colors of the balls drawn from both urns (specifically, on whether the

two balls have the same color, or whether they have different colors). The behavior we identify is then tested in a controlled laboratory experiment which provides evidence of sensitivity to the lack of information concerning the relation between the urns' compositions. By considering also bets on an urn known to contain an equal number of red and black balls, we find that this new form of ambiguity aversion is only partially related to Knightian ambiguity aversion as measured in Ellsberg's experiment, and we study at a behavioral level the *three-fold distinction* between risk, uncertain bias (or composition) and the uncertain relation between biases.

The experimental design developed to study these preferences identifies strict preference using pairwise choices between bets. To check the robustness of the results, we also employ standard choice lists that elicit an approximation to the certainty equivalent of each bet.

The preceding summarizes the main contributions of the paper: highlighting ambiguity about correlation and providing supporting experimental evidence. The potential relevance for economic applications is addressed next in this introduction (and also in the concluding section). Our thought experiment is described in Section 2, and the following two sections describe the experimental implementation and results. Section 5 turns to the question of how to model the observed behavior. Finally, a concluding section elaborates on the economic relevance of ambiguity about correlation and discusses some related literature.

**Economic significance:** Betting on the draws from several urns is intended as a canonical example of choice problems where payoffs to an action depend on the (simultaneous) realization of multiple random events. A textbook example is the choice between bets on multiple tosses of a coin of unknown bias. Optimal portfolio choice is another important example. Here we indicate that the behavior we identify is relevant (albeit indirectly) to the potential gains from portfolio diversification, one of the central principles of financial economics, to the "limited stock market participation" puzzle, and to the pricing of idiosyncratic uncertainty in a cross-sectional setting. The concluding section describes other motivating examples.

It is well-known (Dow and Werlang 1992) that given a safe asset and a single uncertain stock, then nonparticipation in stocks is a knife-edge property under subjective expected utility maximization, but that it can be robustly optimal, (that is, optimal for a range of expected excess returns), if there is ambiguity aversion. This begs the question whether in the more realistic

situation where there are many uncertain assets, if diversification can diminish the effect of ambiguity on participation and possibly even restore the expected utility result asymptotically when the number of stocks is large. As a concrete example, suppose that there are  $I$  securities available and that the investor’s model of returns is a linear factor model as in arbitrage pricing theory. That is, the  $i$ th return  $s_i$  takes the form

$$s_i = \beta_i \cdot X + \epsilon_i, \quad i = 1, 2, \dots, I;$$

the vector  $X$  gives factor returns and  $\beta_i$  gives the betas or factor loadings of security  $i$ . Typically, strong assumptions on the idiosyncratic terms  $\epsilon_i$  are adopted, say that they are identically and independently distributed (i.i.d.). However, suppose that the investor is not confident that  $X$  captures all relevant factors. She may not be able to identify missing factors—if she could, then  $X$  could be expanded to include them—but she may nevertheless wish to take into account their possible existence when choosing a portfolio. Then there is no basis for taking a stand on how the  $\epsilon_i$ s may differ or are related across securities—if rates of return are influenced also by omitted factors, then the distribution of the residuals  $\epsilon_i$  depends on the nature of the omitted factors, which, by assumption, is poorly understood. In particular, if the investor is sensitive to the resulting ambiguity about the differences and relation between returns, she may perceive only limited gains from diversification which may limit her degree of stock market participation. Indeed, Epstein and Seo (2015) show, using a formal model of preference that can accommodate the new ambiguity-sensitive behavior that is our focus here, that the optimality of nonparticipation is robust to the presence of many stocks even if  $I$  grows without bound. In the literature on optimal portfolio choice, Fouque, Pun and Wong (2016), in a continuous-time model, and Liu and Zeng (2016) and Huang, Zhang and Zhu (2017), in a static mean-variance framework, show that limited diversification is optimal given suitable ambiguity about the correlation of returns (see also Jiang and Tian (2016)). Readers are referred to these papers and the references therein for arguments that estimation of the correlation of returns can be difficult, more so than for moments associated with any single asset, thus supporting the assumption of correlation ambiguity as opposed to correlation risk.<sup>1</sup>

---

<sup>1</sup>One reason cited for the difficulty given high-frequency data, is that because different assets are traded at different times it is necessary to synchronize asset returns and hence remove some data points (Ait-Sahalia et al 2010).

There are effects also on the pricing side—ambiguity about how returns are related implies that idiosyncratic uncertainty can have a positive price in equilibrium even in the limit as  $I$  goes to infinity (Epstein and Schneider 2010). Though this may strike some as contrary to the intuition based on the classic Law of Large Numbers, the point is that the latter does not apply when there is ambiguity about security returns (that is, of the  $\epsilon_i$ s) of the sort studied here (Maccheroni and Marinacci 2005).

## 2 A Thought Experiment

Two urns, numbered 1 and 2, each contain two balls, each of which is either red or black; no additional information about the urns' compositions is given.<sup>2</sup> One ball is to be drawn from each urn simultaneously. Thus the set of possible outcomes is  $\{R_1B_2, B_1R_2, R_1R_2, B_1B_2\}$ , where the characters correspond to the color of the ball (red or black) and the subscript to the urn (1 or 2). Its subsets are called events. Before the balls are drawn, an individual is asked to choose between specified bets on the colors of the two balls. Denote by  $R_1B_2$  both the obvious event and the corresponding bet that yields the (positive) prize  $x$  if that event is realized and the prize 0 otherwise; similarly for other events and bets. Prizes are denominated in dollars. Both the events  $\{R_1B_2, R_1R_2\}$  and  $\{B_1R_2, B_1B_2\}$  and the corresponding bets on the color of the ball drawn from urn 1 are sometimes denoted simply  $R_1$  and  $B_1$  respectively. Let  $\succeq$  denote a preference relation on the set of bets.

As outlined in the introduction, our thought experiment, which we term *One vs Two*, offers the individual the choice between betting on the color drawn from one urn as opposed to betting on the colors drawn from both urns. More precisely, consider the following rankings between the bets *Same* =  $\{R_1R_2, B_1B_2\}$ , *Diff* =  $\{R_1B_2, B_1R_2\}$  and the bets  $R_1$  and  $B_1$  on the draw from urn 1:

$$R_1 \succ \textit{Same} \quad \text{and} \quad B_1 \succ \textit{Diff}. \quad (2.1)$$

Why would an individual exhibit these rankings? The intuition is that only the bets on the draws from both urns are subject to ambiguity about how urns differ or are related, which may, depending on the degree of aversion to

---

<sup>2</sup>In Ellsberg's two-urn experiment, the individual is given different information about the two urns. In particular, since one of the urns in that experiment has known composition, there is no scope for ambiguity about how the two compositions may be related which is our focus here.

such ambiguity, lead to the preference for the bets  $R_1$  and  $B_1$ . To elaborate, consider the choice between  $R_1$  and *Same*. Betting on the same color being drawn from both urns is attractive if it is believed that the compositions of the two urns are similar, which would make “positive correlation” between the colors drawn likely. Since you are not told anything to the contrary, this belief is plausible but no more so than the belief that the two compositions are different—one urn is biased towards red and the other towards black—which would make “negative correlation” between the colors drawn more likely and render *Same* an unattractive bet. Given a conservative attitude, this uncertainty would act against choosing *Same*. Of course, there is also reason for a conservative individual to discount the bet  $R_1$  because the composition of each urn is ambiguous. Conclude that a preference for  $R_1$  may arise if there is *greater aversion* to ambiguity about the relation between urns than to ambiguity about the bias of (the first urn and hence, presumably) any single urn.

A simpler rationale for a strict preference for  $R_1$  is that the individual believes that the first urn has many more red than black balls, which rationale does not rely on sensitivity to ambiguity of any sort. To rule out this alternative rationale, we follow Ellsberg and the ensuing literature in considering also another choice—in our case, that between  $B_1$  and *Diff*. A strict preference for  $B_1$  can be understood as above as a reflection of an aversion to ambiguity about the relation between urns, (here the unfavorable scenario for *Diff* is that the two urns might both be biased towards the same color), which thus “explains” both rankings indicated in (2.1). In contrast, the pair of rankings is inconsistent with beliefs that can be represented by any probability measure: there does not exist a measure  $P$  on the four possible pairs of colors satisfying

$$\begin{aligned} P(R_1) &> P(\textit{Same}), \quad P(B_1) > P(\textit{Diff}) \\ P(R_1) + P(B_1) &= 1 = P(\textit{Same}) + P(\textit{Diff}). \end{aligned}$$

More formally, the rankings contradict probabilistic sophistication as defined by Machina and Schmeidler (1992).<sup>3</sup> Note that such a contradiction exists also if (2.1) is weakened so that at most one of the rankings is weak. Such behavior is abbreviated below as *One*  $\succ$  *Two*.

---

<sup>3</sup>Familiarity with their formal definition is not needed in the sequel. The reader can take probabilistic sophistication to mean simply that there exists a probability measure such that, when choosing between bets, the individual always prefers to bet on the event having higher probability.

A contradiction to probabilistic sophistication exists also if both rankings in (2.1) are reversed, with at least one of them strict, which we abbreviate by  $Two \succ One$ . In the Ellsberg experiment, where the composition of one urn is known, a preference to bet on the unknown urn is naturally understood as ambiguity seeking. Here, however, because both urns have unknown compositions and all four bets considered above are ambiguous, the behavior we study does not justify taking a stand on whether the individual likes or dislikes ambiguity about correlation in an absolute rather than relative sense. As a result, just as we interpreted (2.1) above in terms of *relative* ambiguity aversion, we interpret the pair of reverse rankings as indicating a *lesser aversion* to (or a greater affinity for) ambiguity about the relation between urns than to ambiguity about the bias of any single urn.

Besides those discussed thus far, all other choices in One vs Two are consistent with probabilistic sophistication. For example, the rankings

$$Same \succeq R_1 \text{ and } B_1 \succeq Diff \tag{2.2}$$

can be rationalized by any probability measure satisfying  $P(B_1 B_2) \geq P(R_1 B_2)$ . A Bayesian with an i.i.d. prior uniform within each urn would be indifferent between all four bets indicated.

We turn now to describing our experimental investigation of One vs Two.

### 3 Experimental Design

We conducted three experiments (in 2013, 2014 and 2015) that study the behavior in One vs Two. We report the two earlier studies in Appendix A.1.<sup>4</sup> We concentrate here on the final experiment which included two experimental designs.

Subjects were presented with two (ambiguous) urns and were told that each contained two balls, each ball being either red or black, and also a third (risky) urn that contained one red and one black ball.<sup>5</sup>

---

<sup>4</sup>In 2014 we included important control treatments that tested alternative explanations for failure of probabilistic sophistication in One vs Two.

<sup>5</sup>The language in the experiment used jars and marbles that were blue or green.

### 3.1 Pairwise choices

In the main experimental design, which was also employed in the earlier experiments, subjects were presented with versions of the following six choice problems. Note that in the first four choice problems, subjects were asked to consider only the two ambiguous urns.

***Red vs Different (1)***: Choose between a bet that pays if the color of the ball drawn from a single urn is red and a bet that pays if the balls drawn from the two urns have different colors.

***Black vs Different (2)***: Choose between a bet that pays if the color of the ball drawn from a single urn is black and a bet that pays if the balls drawn from the two urns have different colors.

***Black vs Same (3)***: Choose between a bet that pays if the color of the ball drawn from a single urn is black and a bet that pays if the balls drawn from the two urns have the same color.

***Red vs Same (4)***: Choose between a bet that pays if the color of the ball drawn from a single urn is red and a bet that pays if the balls drawn from the two urns have the same color.

***Standard Ellsberg Red (5)***: Choose between a bet that pays if the color of the ball drawn from one of the two ambiguous urns is red and a bet that the color of a ball drawn from a third urn containing one red ball and one black ball (a risky urn as in Ellsberg) is red.

***Standard Ellsberg Black (6)***: Choose between a bet that pays if the color of the ball drawn from one of the two ambiguous urns is black and a bet that the color of a ball drawn from a third urn containing one red ball and one black ball (a risky urn as in Ellsberg) is black.

Note that the first four choice problems included two variations of (2.1), where each pair of problems (1/3 and 2/4) allowed the experimenter to detect violation of probabilistic sophistication.

The choice problems above were organized in triplets, which allowed us to infer strict and weak rankings from choices by slightly varying the prizes, assuming monotone and transitive preferences. For example, problem 3' asked the subject to choose between a bet paying \$25 if the ball drawn

from urn 1 is black and a bet paying \$25 if the balls drawn from the two urns are of the same color. Choice made in this problem reveals only weak preference. Problem 3 (3'') was similar except that the winning prize on the latter (former) bet was increased to \$26. The choice to bet on the single urn in problem 3 implies strict preference when the two prizes are equal; and symmetrically if one chooses in problem 3'' to bet on the two colors being the same. Choice of the two bets (in both 3 and 3'') that pay \$25 is inconsistent with monotone and transitive preference. Choice of the two bets that pay \$26, together with either choice in 3' reveals only weak preference between the bets.<sup>6</sup> The rationale behind this design was explained to subjects before they answered any questions.

In the same way, the experimental design can reveal strict and weak ambiguity attitudes through versions of the two standard Ellsberg choice problems. Note that a subject who weakly prefers bets on the color of the ball drawn from the risky urn to the corresponding bets on the color of the ball drawn from an ambiguous urn does *not* necessarily violate probabilistic sophistication. However, if at least one of the above preferences is strict, we can infer that she must be *strictly ambiguity averse* and not probabilistically sophisticated. This approach yields lower and upper bounds on the frequency of ambiguity attitude.

Finally, a similar procedure is applied to versions of the first four problems above (concerning *One vs Two*). Here also the distinction between weak and strict preference is important. For example, a subject who weakly prefers  $R_1$  to *Diff* and  $B_1$  to *Same* is only weakly more averse to ambiguity about correlation than towards bias and may be probabilistically sophisticated. However, if she *strictly prefers* either  $R_1$  or  $B_1$ , then  $\text{One} \succ \text{Two}$  and there is necessarily a violation of probabilistic sophistication. In a symmetric manner we define a subject who is strictly (weakly) more averse to ambiguity about bias than about correlation, and denote this preference by  $\text{Two} \succ (\succeq) \text{One}$ .

It may be helpful to “behaviorally” classify subjects into those who always choose the higher prize in every pair of problems in which the prizes are different, and those who sometimes choose the bet with the lower prize. The first group never exhibits strict preference, and therefore is consistent

---

<sup>6</sup>It is important to remember that the choice of the two bets paying \$26 in the problems in which prizes differ is consistent with indifference. However, it could be that lowering the higher prize to  $\$25 + \varepsilon$  (for example, \$25.1) in these problems might cause a subject to choose the same bet in all three comparisons, thus revealing strict preference when prizes are equal.

with probabilistic sophistication. Subjects in the second group may exhibit behavior that is inconsistent with probabilistic sophistication (Ellsbergian ambiguity aversion or seeking,  $\text{One} \succ \text{Two}$  or  $\text{Two} \succ \text{One}$ ), or make choices that are consistent with “non-symmetric” probabilistic beliefs (for example,  $R_1 \succ R_3$  and  $B_3 \succeq B_1$  in the Ellsberg case, and  $\text{Diff} \succ R_1$  and  $B_1 \succeq \text{Same}$  in One vs Two).

This design of choice problems, which we have not seen used previously, allows us to identify a strict ordinal ranking without using a cardinal valuation. Methods based on elicitation of cardinal valuations of bets (“matching tasks”), such as Becker-DeGroot-Marschak (1964) used in Halevy (2007) or a discrete version using a choice list used in Abdellaoui et al (2011), can obviously identify strict ranking. The problem is that often the response modes (matching vs. choice tasks) disagree. One may be worried, however, that our new design increases the complexity of the choice problems subjects are required to consider. To partially answer this concern, we included in the experiment a choice list design such as has been used in many experimental studies.

## 3.2 Choice lists

In addition to the main pairwise choice design, in 2015 we included a standard choice list design. The latter elicits an approximation to the certainty equivalents of the six bets (each with winning and losing prizes \$25 and \$0 respectively): Red, Black, Same, Different, Risky (50%) Red, Risky (50%) Black. In each choice list the subject was asked to make 20 pairwise choices between the bet and a sure amount, which varied in increments of at least \$1 from \$0 to \$25. The first comparison in which a subject chooses the sure amount approximates the subject’s certainty equivalent, and can serve as a cardinal index for utility. Thus, the comparison made in One vs Two and in the Ellsberg standard problem may be achieved by comparing the switching points of the different bets. In order to mimic the strict preference elicitation achieved in the pairwise choice design, we elicited strict preference between two bets only if the difference between their switching points is at least two lines.<sup>7</sup>

---

<sup>7</sup>A common practice is to use a single line difference as indicating strict preference. However, it could be that a subject is indifferent between two bets and a sure amount that appears on the choice list, and when forced to choose breaks the indifference differently in the two lists. The use of two line differences makes this scenario impossible (a similar

### 3.3 Incentives and other details

Payment was determined by randomly selecting a single choice problem for payment *before* subjects made their choices, by distributing (before the actual experiment started) sealed envelopes containing a number of a randomly selected choice problem (and a randomly selected line in the case of choice lists) among the participants in the experiment. The envelopes were opened only after subjects made all their choices and the balls were drawn from the urns.

This version of the Random Incentive System (RIS) is theoretically incentive compatible in eliciting ambiguity attitudes (Baillon, Halevy and Li 2014), because the order suggests that choices are between lotteries over Savage acts wherein ambiguity cannot be hedged. We do not find evidence that subjects used the RIS to hedge the different sources of ambiguity. See Appendix A.2.4 for further discussion and examples, and for detailed evidence.

Subjects were recruited from UBC’s Vancouver School of Economics subject pool using ORSEE (Greiner 2015) to an experiment that promised participants a chance to earn up to \$36 during a one hour experiment in decision making (including a show-up fee of \$10). After consent forms were signed, the instructions were read aloud.

In order to eliminate a potential suspicion that the experimenter could manipulate the composition of the urns, each subject was asked at the beginning of the experiment to choose the urn (1 or 2) on which to bet in all those bets that involve a single urn. To simplify exposition of the results, we proceed as though all individuals chose urn 1. Thus, for example, the bet on red from the single urn is represented by  $R_1$ . Throughout the experiment the language used in the description of the different bets was completely symmetric, enumerating the states in which each bet pays. Complete instructions may be found in the Online Appendix.

## 4 Experimental Results

This section describes the main results of the experiments described above (and conducted in 2015). The results of previous experiments (conducted in 2014 and 2013), which had higher stakes and included additional treatments that support our interpretation of the behavior, are included in Appendix

---

approach was taken in Halevy (2015)).

**A.2.** The latter also contains some additional information regarding the 2015 results.

Out of 153 subjects who took part in the pairwise choice experiment, 19 made choices inconsistent with monotone and transitive preferences in at least one (out of six) triplet of choices. An additional 5 subjects made intransitive choices between triplets of choices problems.<sup>8</sup> The results reported below are restricted to the remaining 129, although adding (to the extent possible<sup>9</sup>) some of these subjects does not substantially change any of the findings. The average payment was \$17.01 with a standard deviation of \$12.79.

The choice list treatment included 77 subjects, 3 of whom had multiple switching lines or did not respond in some of the decision problems. The reported results are restricted to the remaining 74 subjects. The average payment was \$19.62 (including a show-up payment of \$5) with a standard deviation of \$9.20.

There was no overlap between the subjects who completed the pairwise choice experiment and those who completed the choice list experiment.

Note that there are many ways in which “random” choices in the pairwise choice design would lead to cyclical rankings (since subjects made choices in two different variations of (2.1)). We find no empirical evidence of such behavior.<sup>10</sup> Similarly, the low frequency of multiple switching in the choice list design (only 3 subjects), and the symmetry in the valuations of different colors (only 10 subjects had color valuations that vary in more than \$1) lead us to conclude that choices were deliberate.

---

<sup>8</sup>For example, a subject may exhibit  $R_1 \succ Diff$ ,  $Diff \succ B_1$ ,  $B_1 \succ Same$ , and  $Same \succ R_1$ .

<sup>9</sup>13 subjects made choices inconsistent with monotone or transitive preferences in two or more triplets. For 8, we could not impute any ranking since there were too many missing data; 3 were Ellsbergian ambiguity averse but there was no conclusive evidence for strict preference between bets on One and Two.

Another 6 subjects made choices inconsistent with monotone or transitive preferences in only a single triplet. Three of them were Ellsbergian ambiguity averse, and two of the three made choices consistent with One  $\succ$  Two (the other 3 made choices consistent with probabilistic beliefs in both domains).

Five other subjects made choices inconsistent with transitive preferences across triplets. None of them was Ellsbergian ambiguity averse, but 3 violated probabilistic choices in at least one pair of triplets.

<sup>10</sup>The code checks for all such cycles and is available from the authors.

Ellsbergian Ambiguity Attitude	Pairwise Choice		Choice List	
	#	%	#	%
<b>strictly averse (nonPS)</b>	60	46.5	30	40.5
<b>strictly seeking (nonPS)</b>	3	2.3	10	13.5
<b>weakly averse</b>	37	28.7	5	6.8
<b>weakly seeking</b>	14	10.9	5	6.8
<b>neutral</b>	–	–	20	27.0
<b>non-symmetric beliefs (PS)</b>	15	11.6	4	5.4
<b>Total</b>	<b>129</b>	<b>100</b>	<b>74</b>	<b>100</b>

Table 4.1: Ellsbergian ambiguity attitude

## 4.1 Ellsbergian ambiguity

Table 4.1 presents the attitude towards Ellsbergian ambiguity as derived from choices between bets on the color of a ball drawn from the ambiguous urn and bets on the color of the ball drawn from the risky urn.

In the pairwise choice, 63 out of 129 subjects (48.8%) exhibit strict ambiguity attitude and hence are not probabilistically sophisticated (nonPS); all but 3 are ambiguity averse. The corresponding numbers in the choice list elicitation are 40 out of 74 (54%), and 75% of them are strictly ambiguity averse.

The choices made by the remaining subjects are consistent with probabilistic sophistication (PS). It is important to remember that in the pairwise choice elicitation method we cannot elicit indifference, but only weak preference. It is possible that some of the 51 subjects who revealed weak ambiguity preference, (those who chose the risky bet when prizes are equal and otherwise chose the bet with the higher prize), are indifferent between bets on risky and ambiguous events, and that for others the premium we used (\$1) was too large to reveal strict preference. In particular, it is possible that these 51 subjects all violate probabilistic sophistication.<sup>11</sup> The categories labeled as “weak preference” in the choice list correspond to a difference of \$1 in the switching lines (as noted in Section 3.2), while ambiguity neutrality

---

<sup>11</sup>The fact that 72.5% of the subjects who exhibited weak preference chose the risky bets when the prizes were equal is consistent with the hypothesis that the proportion of ambiguity averse subjects among them is higher than the proportion of ambiguity seeking subjects.

<b>One vs Two</b>	<b>Pairwise Choice</b>		<b>Choice List</b>	
	#	%	#	%
<b>One <math>\succ</math> Two (nonPS)</b>	25	19.4	25	33.8
<b>Two <math>\succ</math> One (nonPS)</b>	10	7.7	18	24.3
<b>One <math>\succsim</math> Two</b>	30	23.3	5	6.8
<b>Two <math>\succsim</math> One</b>	15	11.6	1	1.3
<b>One <math>\sim</math> Two</b>	15	11.6	17	23.0
<b>non-symmetric beliefs (PS)</b>	34	26.4	8	10.8
<b>Total</b>	129	100	74	100

Table 4.2: One vs Two

corresponds to exactly the same switching line.

The final category in Table 4.1 (non-symmetric beliefs) includes subjects who, when comparing bets on one color preferred the ambiguous urn, and when comparing bets on the other color preferred the risky urn.<sup>12</sup> Seven subjects made choices consistent with strictly non-symmetric beliefs and another eight subjects made choices consistent with weakly non-symmetric beliefs. The frequency of this behavior in the choice list design is much lower (only 4 out of 74 subjects).

Comparing the two elicitation methods, they detect ambiguity aversion and neutrality at roughly similar rates, pairwise choice classifies proportionally more subjects as having non-symmetric beliefs over the composition of the ambiguous urn, and the choice list design finds relatively more evidence for ambiguity seeking.

## 4.2 One vs Two

Table 4.2 presents the attitude in One vs Two as derived from choices between bets on the color of a ball drawn from a single ambiguous urn and bets on the colors of the balls drawn from the two ambiguous urns.

In the pairwise choice design, 35 subjects (out of 129) made choices in One vs Two that are inconsistent with probabilistic sophistication, 71% of them strictly preferring a bet on the bias to a bet on the relation between

---

<sup>12</sup>Note that these choices must be consistent with choices made in the first four decision problems that study rankings in One vs Two.

the urns. Another 45 subjects made choices that reveal weak preference to bet on bias or on correlation (two thirds of them weakly preferred bets on bias). As discussed above, some of them may hold probabilistic beliefs, but for others the \$1 difference in prizes may have been too high to reveal strict preference. Altogether, between 35 and 80 subjects (out of 129) made choices consistent with violation of probabilistic sophistication in One vs Two.

In the choice list elicitation, probabilistic sophistication was contradicted by 43 subjects (out of 74), of whom 58% strictly preferred betting on bias rather than on correlation between the urns. An additional 6 subjects showed weak preference (\$1 difference in the switching points), and may have shown strict preference if the difference between lines in the choice list were smaller.

The category One~Two in Table 4.2 includes subjects who did not exhibit any strict preferences between the 4 ambiguous bets and who were not included in the two previous lines. This includes 15 subjects in the pairwise choice and 17 subjects in the choice list.

The last category in Table 4.2 includes subjects who made choices consistent with probabilistic beliefs that assign a strictly higher likelihood that the urns are positively rather than negatively correlated (or vice versa); for example,  $Same \succeq R_1$  and  $B_1 \succeq Diff$  (where at least one ranking is strict). This pattern of choices is quite common in the pairwise choice design as 34 subjects (26.4% of 129) are classified in this way. It is not as frequent in the choice list treatment (only 8 subjects, 10.8% of 74).

Comparing the two elicitation methods, some differences emerge. First, the proportion of subjects that are not probabilistically sophisticated is much higher in the list elicitation than in the pairwise choice (58.1% vs. 27.1%). However, the difference is much smaller when considering also subjects that exhibited weak preference in One vs Two. Even among the probabilistically sophisticated subjects there are differences between the two elicitation methods. The pairwise choice treatment yields relatively more subjects exhibiting non-symmetric attitude between  $Same$  and  $Diff$  (26.4% vs. 10.8%), and the choice list treatment identifies more subjects as being indifferent between all ambiguous bets (23% vs 11.6%). We find these differences, like the differences in the elicitation of Ellsbergian ambiguity attitude, intriguing and deserving independent experimental investigation. We conjecture that some of the differences can be attributed to the fact the choice list compares (approximation to) certainty equivalents (which integrates payments and beliefs), while in pairwise choice no such integration is necessary. It could also be that the relatively large “tick” of \$1 contributed to these differences. Fi-

nally, it is possible that the pairwise choice method we have proposed here is cognitively more taxing relative to the choice list method, leading to more noise in measurement of underlying preferences.<sup>13</sup>

### 4.3 Comparison

In the pairwise choice design, the proportion of subjects whose choices are consistent with probabilistic sophistication is higher in One vs Two (72.9%) than in the standard Ellsberg problem (51.2%). This difference may not be surprising, in hindsight, because, as we suggest next, *there is more scope for probabilistically sophisticated behavior in One vs Two than in Ellsberg*, even for a subject who dislikes (or alternatively likes) ambiguity. In the choice list design, however, we did not observe such a difference in the relative frequency of probabilistic sophistication: choices consistent with probabilistic sophistication were made by 34 subjects in the Ellsberg problems and by 31 subjects in One vs Two.

In the classic Ellsberg problem, the choice of the risky urn leaves the subject with a purely risky bet in which the probability of winning is 50%. Thus, the preference to bet on the risky urn, and hence violation of probabilistic sophistication, arises given only aversion to uncertainty about the bias of the ambiguous urn. In contrast, in the choice problem One vs Two all alternatives (the bets  $R_1$ ,  $B_1$ , *Same* and *Diff*) are ambiguous. Accordingly, as explained when discussing our thought experiment, choices depend on which is more important—ambiguity about bias or ambiguity about how urns differ or are related. If these opposing motives exactly offset one another, then lack of strict preference between  $R_1$ ,  $B_1$ , *Same* and *Diff* (classified under One  $\succeq$  Two, Two  $\succeq$  One, or One  $\sim$  Two in Table 4.2) could result. This rationale includes 20 subjects in the pairwise choice design and 8 in the choice list design whose choices are not probabilistically sophisticated in the Ellsberg problems (total of 63 and 40 in the pairwise choice and choice list designs) but do not exhibit strict preference in One vs Two. Hence, this behavior is consistent with all bets being perceived equally ambiguous. Our categorization of this behavior as being probabilistically sophisticated in Table 4.2 reflects the conservative (and demanding) approach we are adopting throughout to identifying empirical support for the effects of ambiguity about correlation.

---

<sup>13</sup>Noise should be distinguished from random choices, as the latter will lead to cyclical choice patterns - that we do not observe here.

There is another possibility consistent with probabilistic sophistication in One vs Two, namely the rankings  $Diff \succ R_1$  and  $B_1 \succ Same$ , which is a variant of (2.2). Such rankings can arise from a strong belief about how the urns were constructed. As an extreme example, suppose that the subject’s hypothesis is that the experimenter drew two balls without replacement from an auxiliary urn containing one red ball and one black ball, and if a red (black) ball was drawn first, then urn 1 was filled with 2 red (black) balls; urn 2 was then filled with two balls of the opposite color. Then it is certain that the balls drawn from urns 1 and 2 have different colors and thus  $Diff$  pays  $x$  (25 or 26) and  $Same$  pays 0, each with certainty, so that  $R_1$  is ranked between them. The latter might arise also more generally from the feeling that “there are only so many red balls to go around,” say because the urns are thought to have been constructed by drawing without replacement from an auxiliary urn containing  $n \geq 2$  balls of each color. The description of urns given to the subjects does not suggest this perception but there is no reason to rule it out. Probabilistic sophistication is consistent also with the rankings (2.2) as indicated at the end of Section 2. This might arise if urns are perceived to have a common component, so that a red draw from urn 1 indicates that a red draw is more likely also from urn 2; for example, if the preceding construction is modified so that draws from the auxiliary urn are made with replacement. Out of the subjects who are not probabilistic sophisticated in the Ellsberg problem (63 and 40 in the pairwise choice and choice list elicitation, respectively) there are many subjects who exhibit this pattern (23 and 4 in the two methods, respectively). The substantial difference in the patterns across the methods suggests that at least part of the interpretations above depend on the specific incentive system used to elicit rankings.

In a similar vein, not all subjects who violate probabilistic sophistication in One vs Two, violate it also in the Ellsberg problem. Out of 35 and 43 in the pairwise choice and choice list methods, respectively, only 20 and 25 violate probabilistic sophistication in the Ellsberg problem. For example, some subjects are ambiguity neutral or even ambiguity seeking when betting on bias, but are averse to the ambiguity concerning the correlation between the urns.

For the reasons given above, we do not expect behavior in One vs Two to mirror behavior in the Ellsberg problems. Indeed, although the association between probabilistic sophistication in the two problems is significant in the choice list elicitation (Fisher exact test  $p$ -value 3.4%), the association is not

significant when measured in the pairwise choice design (Fisher exact test  $p$ -value  $> 10\%$ ).<sup>14</sup> When partitioning the set of probabilistically sophisticated subjects in One vs Two into those who expressed or did not express strict preference, we find a significant association between the two domains (Fisher exact test  $p$ -value 0.3%): subjects who did not express strict preference in one domain (Ellsberg / One vs Two) tended not to express strict preference in the other domain as well.<sup>15</sup>

## 5 Models

Consider how the behavior exhibited in the experiment can be modeled. As shown earlier, both One $\succ$ Two and Two $\succ$ One contradict probabilistic sophistication, and hence, in particular, subjective expected utility theory. Thus we are led to consider generalizations of the latter that were developed in response to Ellsberg’s experiments; the maxmin expected utility model (Gilboa and Schmeidler 1989) is one prominent example, and others include Choquet expected utility (Schmeidler 1989), the smooth model (Nau 2006; Klibanoff, Marinacci and Mukerji 2005; Seo 2009), and variational utility (Maccheroni, Marinacci and Rustichini 2006). All of these models have many free parameters and thus can accommodate a broad range of ambiguity-sensitive behaviors. For an extremely simple example, both Ellsbergian ambiguity aversion and One $\succ$ Two are implied by a Choquet expected utility preference if we specify the representing capacity (or non-additive probability measure)  $\nu$  so that

$$\nu(R_1) = \nu(B_1) = .3 \text{ and } \nu(\text{Same}) = \nu(\text{Diff}) = .2; \quad (5.1)$$

and Two $\succ$ One is accommodated instead if .2 is replaced by .4. Similarly the new behavior identified here does not pose a challenge for the other models.

However, our objective is not merely to model the behavior observed in the idealized laboratory setting with Ellsberg urns. We wish also to determine what the experimental results suggest about modeling behavior outside the laboratory when uncertainty is due to multiple random events (as we have seen, portfolio choice is one example; statistical decision problems are typically of this form; see Section 6 for more). These applications suggest a

---

<sup>14</sup>This is due to the fact that many subjects who are Ellsbergian ambiguity averse make choices in One vs Two that may be rationalized by the rankings (2.2).

<sup>15</sup>See further discussion in Appendix A.2.3.

Cartesian product state space *and* the desirability of *preference models that are designed to address and exploit this special structure*. Rationalization through (5.1), or similar exercises, do not satisfy the latter requirement. The Bayesian benchmark is the special case of subjective expected utility, due to de Finetti (1937), where the Savage predictive prior on the state space has the well-known “conditionally i.i.d.” form.<sup>16</sup>

With the preceding motivation, we consider two classes of models in this section. The first builds on de Finetti by adding a role for ambiguity. For the second, we adapt an approach to modeling Ellsberg-style behavior that centers on multiple “sources” (Tversky and Fox 1995; Tversky and Wakker 1995). This model captures naturally and simply a main point of the paper, namely the distinction between three kinds (or sources) of uncertainty: risk, bias and correlation.

## 5.1 Conditionally i.i.d. models

The four possible outcomes of the two draws lie in the state space  $S_1 \times S_2$ , where, for  $i = 1, 2$ ,  $S_i = \{R_i, B_i\}$ . For any subset  $A$  of outcomes,  $A$  denotes also the corresponding bet with prizes  $x$  and 0. According to the maxmin model, the utility of the bet on  $A$  is given by

$$U(A) = \min_{P \in \mathcal{P}} P(A), \quad (5.2)$$

where  $\mathcal{P}$  is a set of probability measures on  $S_1 \times S_2$ , the set of predictive priors.<sup>17</sup> (The model includes also a vNM utility index  $u$ , which here we have normalized to satisfy  $u(x) = 1$  and  $u(0) = 0$ .)

More generally, consider all utility functions over bets of the form

$$U(A) = W((P(A))_{P \in \mathcal{P}}), \quad (5.3)$$

where  $\mathcal{P}$  is a set of predictive priors and where  $W$  is weakly increasing in the sense that

$$U(A') \geq U(A) \text{ if } P(A') \geq P(A) \text{ for all } P \in \mathcal{P}.$$

---

<sup>16</sup>See Kreps (1988, Ch. 11) for a description and for a discussion of its importance as a normative guide to decision-making.

<sup>17</sup>We follow common terminology in that “predictive prior” refers to beliefs about the payoff relevant state space and “prior” refers to beliefs about unknown probability laws or “parameters.”

We consider alternative specifications for  $\mathcal{P}$  and  $W$ .

In the benchmark Bayesian model, each urn is parametrized by a common unknown parameter  $p$ ,  $0 \leq p \leq 1$ , representing the proportion of red balls in the urn. In addition, the two draws are taken to be i.i.d. conditional on the true but unknown  $p$ , and uncertainty about  $p$  is modeled by a (single) prior over its possible values. Because beliefs are probabilistic, this model cannot accommodate the preference for *One* over *Two* (or the reverse) – as noted, one needs to permit a role for ambiguity. The simplest way to do so is to generalize the preceding by positing a nonsingleton set  $\mathcal{M}$  of priors about the value of  $p$ . Each prior  $\mu$  in  $\mathcal{M}$  induces a predictive prior  $P$  in the familiar way described by

$$P = \int (p \otimes p) d\mu(p), \quad (5.4)$$

where  $p \otimes p$  denotes the i.i.d. product of the measure on  $\{R, B\}$  that assigns the probability  $p$  to the outcome  $R$ . By varying over all priors in  $\mathcal{M}$ , one obtains the set  $\mathcal{P}^{exch}$  of predictive priors. Using this set of predictive priors in (5.3), without any further restrictions on  $W$ , constitutes a seemingly natural generalization of de Finetti’s model. (Models of this form are studied in Epstein and Seo (2010, Model 1), Al Najjar and De Castro (2014), Cerreia-Vioglio et al (2013) and Klibanoff, Mukerji and Seo (2014)).

However, this generalized conditionally i.i.d. model *cannot* simultaneously accommodate both *One* $\succ$ *Two* and Ellsbergian (strict) ambiguity aversion: from the fact that  $p^2 + (1 - p)^2 \geq \frac{1}{2}$  for all  $p$  in  $[0, 1]$ , infer that for every  $P$  of the form (5.4),

$$P(\text{Same}) = \int [p^2 + (1 - p)^2] d\mu(p) \geq \frac{1}{2}.$$

It follows, given only that  $W$  is monotonic, that *Same* is weakly preferred to betting on red in a 50-50 urn. Adding Ellsbergian ambiguity aversion yields *Same*  $\succ$   $R_1$ , contrary to *One* $\succ$ *Two*.

The intuitive reason for this limitation is clear: though there is ambiguity about  $p$ , the fact that every predictive prior in  $\mathcal{P}^{exch}$  has the form in (5.4) expresses certainty that draws are conditionally i.i.d., which is contrary to the rationale for preferring to bet on the single urn. This intuition suggests, and it is readily confirmed,<sup>18</sup> that the conditionally i.i.d. model can accommodate

---

<sup>18</sup>Let  $0 < q < 1/2$  be fixed and let  $\mathcal{P}$  consist of the two i.i.d. predictive priors  $q \otimes q$  and  $(1 - q) \otimes (1 - q)$ . Let  $W$  be as in maxmin. Then  $U(R_1) = U(B_1) = q$ ,

both Two>One and Ellsbergian ambiguity aversion.<sup>19</sup> The reason is that Two>One indicates the predominance of ambiguity about the composition of each urn and this can be captured adequately by multiple priors about  $p$ .

## 5.2 A source-based model

We consider next a model that relaxes the restriction that the set  $\mathcal{P}$  contains only conditionally i.i.d. measures, while retaining structure that reflects the distinction between risk, uncertain bias and uncertain correlation.

A streamlined version of the model is described diagrammatically in Figure 5.1. (A more general and precise specification is provided in Appendix A.3.) Notationally, the composition of a single urn is described by a probability vector of the form  $(p, 1 - p)$ , where  $p$  denotes the proportion of red, and the joint composition of the two urns is described by a probability vector of the form  $(p_{RB}, p_{BR}, p_{RR}, p_{BB})$ . If urns 1 and 2 are described by  $(p, 1 - p)$  and  $(q, 1 - q)$  respectively, then  $(p, 1 - p) \otimes (q, 1 - q)$  denotes the joint distribution given by the product measure,

$$(p, 1 - p) \otimes (q, 1 - q) \equiv (p(1 - q), q(1 - p), pq, (1 - p)(1 - q)).$$




Figure 5.1: Hierarchical beliefs on the two ambiguous urns

The evaluation of a bet may be influenced by the following three issues:

$U(\text{Same}) = q^2 + (1 - q)^2 > 1/2$  and  $U(\text{Diff}) = 2q(1 - q) > q$ , which implies all the desired rankings.

<sup>19</sup>Of the 60 (30) subjects in the pairwise choice (choice list) design who were strictly ambiguity averse in the Ellsberg sense, 12 (9) exhibited One>Two and 8 (13) exhibited Two>One.

- Issue 1 Uncertainty about the relation between urns takes the form of two alternative hypotheses. One possibility entertained by the decision-maker is that the urns are i.i.d. according to  $(p, 1 - p)$  for some (unspecified)  $p \in \{0, \frac{1}{2}, 1\}$ , that is, the two urns can be described by a measure in the set  $\mathcal{P}^{ident}$  in the figure. The alternative hypothesis is that the urns are “complementary” in the sense that their joint distribution lies in the set  $\mathcal{P}^{compl}$  in the figure. This hypothesis is justified, for example, by the following perception of how the urns are constructed: there is a set of two red and two black balls—two are drawn without replacement to fill urn 1 and the remaining two are put into urn 2. The two hypotheses are assigned subjective probabilities  $\sigma$  and  $1 - \sigma$  respectively.
- Issue 2 Conditioning on either of the above hypotheses, the composition, or bias, of each urn is uncertain. The decision-maker assigns (conditional) probability  $\frac{1}{3}$  to each possible value of  $p$ .
- Issue 3 After conditioning on both the relation between urns and on the bias, there remains uncertainty about the colors of the two drawn balls. However, resolution of Issues 1 and 2 implies a unique probability distribution over  $\{R_1B_2, B_1R_2, R_1R_2, B_1B_2\}$ . Thus the last issue concerns risk.

Each bet (or act)  $f$  associates a (dollar) payoff to each terminal node, depending on the colors of the two balls drawn. Thus it induces a 3-stage lottery, denoted  $D_f$ , which is evaluated recursively along the lines of Kreps and Porteus (1978): as Appendix A.3 details, an expected utility function is used at each stage, with different utility indices for different issues, and compound lotteries are evaluated recursively.<sup>20</sup> The utility of a bet on Ellsberg’s risky urn is computed by identifying the (single-stage) lottery induced by the bet with a three-stage lottery that is resolved completely at the third stage. Denote the utility indices by  $u_3$  (used to evaluate risk or one-stage lotteries),  $u_2$  (used to address uncertainty due to the unknown bias), and  $u_1$  (used to address the first issue). For simplicity, take  $u_3$  to be linear.

The model is a special case of (5.3), where the set  $\mathcal{P}$  is equal to the union of the sets  $\mathcal{P}^{ident}$  and  $\mathcal{P}^{compl}$  defined in the figure; hence,  $\mathcal{P}$  contains measures (all those in  $\mathcal{P}^{compl}$ ) that do not conform to (5.4). The recursive structure of utility defines the corresponding function  $W$ .

---

<sup>20</sup>For more on such recursive models see, for example, Segal (1987, 1990), and Ergin and Gul (2009). Segal advocates using non-expected utility functions at each stage.

Turn to behavior. Because the bets of interest depend on different issues, the model’s predictions depend largely on the curvatures (both absolute and relative) of  $u_1$  and  $u_2$ . In particular, Ellsbergian ambiguity aversion (affinity) is implied if and only if  $u_2$  is concave (convex). In comparing One vs Two, the bets  $R_1$  and  $B_1$  are subject to uncertainty about the bias, but their payoffs do not depend on the first issue because they concern only a single urn. In contrast, the payoffs to both *Same* and *Diff* depend crucially on the relation between urns. This suggests that the choices in *One* vs *Two* depend on both the relative curvatures of  $u_1$  and  $u_2$  and on the magnitude of  $\sigma$ . If we take  $\sigma = \frac{1}{2}$ , then  $\text{One} \succ (\prec) \text{Two}$  is implied if  $u_1$  is more (less) concave than  $u_2$ . Consequently, the model can handle all preference patterns observed in our experiments by allowing sufficient heterogeneity in the curvatures of subjects’ utility indices and in beliefs over the relation between the urns (see Appendix A.3).

## 6 Concluding Discussion

The literature stimulated by Ellsberg, particularly the experimental literature, has focused on the two-fold distinction between risk and ambiguity, or “Knightian uncertainty.” In Ellsberg’s two-urn experiment, the latter is embodied in the uncertain bias (or composition) of the unknown urn. In a setting with multiple ambiguous urns, we have introduced a second source of ambiguity—the relation between urns. Thus we have studied the three-fold distinction between risk, bias and the relation between biases. The results of our laboratory experiments provide support for the empirical relevance of the new dimension, and its relation to Knightian uncertainty.

To conclude, we describe concrete instances of decision-making where ambiguity about correlation is potentially relevant, and then we consider some related literature.

### 6.1 More on economic significance

As noted in the introduction, betting on the draws from a sequence of urns is intended as a canonical example of choice problems where payoffs to an action depend on multiple random events (or variables). Portfolio choice is one such problem as discussed in the introduction. Here we describe other instances where ambiguity about correlation is plausibly a concern for a decision-maker.

Suppose, for example, that the outcome  $s_i$  of the  $i$ -th random event is given by an equation of the form

$$s_i = \beta \cdot x_i + \epsilon_i, i = 1, 2, \dots, I. \quad (6.1)$$

There may be differences between the underlying mechanisms—these are captured by the vectors  $x_i$  which describe the observable heterogeneity (in the urns context, there is no observable heterogeneity). The decision-maker chooses between bets on the outcomes of the  $I$  random variables. Her choices depend on her model of the residuals or unobserved heterogeneity  $\epsilon_i$ , which are the source of the uncertainty she faces. If all sources of heterogeneity of which she is aware are included in the  $x_i$ s, then it is natural that she be indifferent between any two bets on the realization of residuals that differ only in a reordering of the random variables. However, the individual may not be confident that the  $x_i$ s describe all relevant differences, in which case she may not be certain that residuals are identical, or that they are related in any particular way. Though she may not be able to describe further forms of heterogeneity, she may be worried that there are gaps in her understanding that could be important and thus she may wish to take into account their possible existence when making choices.

A number of studies have argued for the importance of such a lack of confidence. In all the examples below, the decision-maker should be thought of as a *policy maker*, where policies can be identified with (Savage-style) acts that pay off according to the outcomes of the  $I$  random variables. In the context of the cross-country growth literature where each random variable corresponds to the growth rate for one of  $I$  countries, Brock and Durlauf (2001) point to the open-endedness of growth theories as a reason for skepticism that all possible differences between countries can be accounted for (p. 231), and they emphasize the importance of “heterogeneity uncertainty.” King (2001) makes a similar critique in an international relations context where each random variable corresponds to a pair of countries and its outcome is conflict or lack of conflict; he refers to “unmeasured heterogeneity.”<sup>21</sup> This paper complements these critiques by translating them into behavioral terms and thus giving more precise meaning to a concern with “heterogeneity uncertainty” or “unmeasured heterogeneity.”

---

<sup>21</sup>He writes (p. 498) that “in international conflict data are neither powerful nor even adequate summaries of our qualitative knowledge”, so that the common assumption of exchangeability is usually violated.

The applied IO literature provides an example of a different sort.<sup>22</sup> Here there is a cross-section of markets in each of which an entry game is played. Thus each random variable corresponds to one of  $I$  markets and its outcome is the number and identity of entrants in a pure strategy Nash equilibrium. The difficulty faced by the policy maker is that there may be multiple equilibria and she has little understanding of how equilibria are selected, and accordingly how selection mechanisms may differ or be related across markets. A fourth example arises in repeated English auctions when, as in Haile and Tamer (2003), because of the free-form nature of most English auctions in practice, one makes weak assumptions about bidders' behavior. Then equilibrium behavior in each auction is multiple-valued and can be narrowed down and related across auctions only via heroic assumptions. This has implications for an auctioneer who is choosing reserve prices (Aryal and Kim 2013). Though our laboratory experiment does not investigate behavior in the above specific settings, the results lend support to the hypothesis that decision-makers may care about poorly understood differences across markets or auctions.

## 6.2 Related literature

We have already contrasted Ellsberg's classic two-urn experiment with ours. Eliaz and Ortoleva (2016) and Eichberger et al (2015) modify Ellsberg's experiment in order to study different dimensions of ambiguity; in addition to ambiguity about the probability of winning, both papers consider also ambiguity about the prize to be won, and the former considers in addition ambiguity about the time delay for receipt of the prize. Their experiments have in common with ours the investigation of multiple dimensions of ambiguity and the association between (ambiguity averse) behaviors in different dimensions, but the "dimensions" studied are different; this leads also to completely different experimental designs.

Multiple urns are used in experiments exploring learning and dynamic consistency. These issues are not involved in our study because we consider only *ex ante* choice.

Epstein and Seo (2010, Model 2) and (2015) present axiomatic models that generalize de Finetti's Bayesian model in a way different from the

---

<sup>22</sup>We are referring to the literature on entry games and partial identification (see Tamer 2010, and the references therein). For an explicit choice-theoretic perspective, see Epstein and Seo (2015).

generalizations described in Section 5. Their capacity to accommodate the behavior studied here is intermediate between the capacities of the two models examined above; more detail would take us too far afield. We add only that though they are axiomatic, their axioms are not nearly as simple and transparent as the choice between betting on one urn or on two. Another important difference is that the axioms are formulated in terms of horse-race/roulette wheel (Anscombe-Aumann) acts which are less natural objects of choice than are the Savage-style acts considered here.<sup>23</sup> Correspondingly, while the former acts are prominent in axiomatic work and are used in some laboratory experiments, we are not aware of any applied studies where they have been used to model behavior in the field or to explain field data—objects of choice are universally formalized as lotteries or as Savage acts. This disconnect in the literature between Anscombe-Aumann acts and descriptive modeling in the field suggests (to us) that tests of preference models that refer only to Savage-style acts are more relevant to the potential usefulness of these models outside the laboratory.<sup>24</sup>

Finally, the role played in decision-making by the individual’s perception of correlation is the focus also in a literature on “correlation neglect,” including, for example, experimental studies (Eyster and Weizsacker 2010, Enke and Zimmerman forthcoming), applications (De Marzo et al 2003, Levy and Razin 2015), and decision-theoretic foundations (Ellis and Piccione, 2017). These papers concern environments where there is a “true” probability law describing the uncertainty (that is, the law observed by the modeler), and where the decision-maker is Bayesian but where she has wrong beliefs. Because of limited attention or other cognitive constraints, she misperceives the connections between key random variables (for example, between returns to different securities in financial decision-making). A common hypothesis is that she wrongly treats key correlated variables as independent. There is no sense in these models that awareness of the complexity of her environment and self-awareness of her cognitive limitations lead the decision-maker to doubt that her wrong beliefs are correct. Such doubts are what we are try-

---

<sup>23</sup>Savage-style acts can be thought of as consisting of the horse-race alone followed by monetary prizes as opposed to the spin of a roulette-wheel; see Kreps (1988), particularly Chs. 4 and 7.

<sup>24</sup>In discussing if/how the use of axioms involving Anscombe-Aumann acts makes sense, Kreps (p. 101) writes “In descriptive applications, axioms are supposed to concern behavior that is observable, so what sense does it make to pose axioms about preferences/choices that are never observed ...?”

ing to detect here, through behavior (such as One $\succ$ Two or Two $\succ$ One) that contradicts probabilistic sophistication.<sup>25</sup> Note that the Eyster-Weizsacker and Enke-Zimmerman experiments are not designed to test for violations of probabilistic sophistication.

In a very recent paper, Levy and Razin (2017) model agents who are self-aware in the above sense and accordingly are averse to ambiguity about the correlation between different sources of information. Correlation aversion is modeled with a specific single-parameter (“correlation capacity”) functional form which the authors apply to models of financial investments and CDO ratings. Thus preferences are assumed known to the modeler and some effects of changes in preferences are explored in specific settings. We see such work as complementary to ours, which takes behavior alone to be observable and asks “which behavior would reveal a concern with ambiguous correlation?”

## References

- [1] M. Abdellaoui, A. Baillon, L. Placido and P.P. Wakker, The rich domain of uncertainty: source functions and their experimental implementation, *Amer. Econ. Rev.* 101 (2011), 695-723.
- [2] M. Abdellaoui, P. Klibanoff and L. Placido, Experiments on compound risk in relation to simple risk and to ambiguity, *Manag. Sc.* 61 (2015), 1306-1322.
- [3] Y. Ait-Sahalia, J. Fan and D. Xiu, High frequency covariance estimates with noisy and asynchronous financial data, *J. Amer. Stat. Assoc.* 105 (2010), 1504-1517.
- [4] N. Al Najjar and L. De Castro, Parametric representation of preferences, *J. Econ. Theory* 150 (2014), 642-667.
- [5] G. Aryal and D. H. Kim, A point decision for partially identified auction models, *J. Bus. Econ. Statist.* 31 (2013), 384-397.
- [6] A. Baillon, Y. Halevy and C. Li, Incentive compatibility of the random incentive system in ambiguity elicitation, 2014.

---

<sup>25</sup>We emphasize that violation of probabilistic sophistication rules out beliefs that can be represented by *any* probability measure, including one expressing stochastic independence or being “wrong” in any other way.

- [7] G. M. Becker, M. H. DeGroot and J. Marschak, Measuring utility by a single response sequential method, *Behavioral Science* 9 (1964), 226-232.
- [8] W. Brock and S. N. Durlauf, Growth empirics and reality, *The World Bank Review* 15 (2001), 229-272.
- [9] S. Cerreia-Vioglio, F. Maccheroni, M. Marinacci and L. Montrucchio, Ambiguity and robust statistics, *J. Econ. Theory* 148 (2013), 974-1049.
- [10] S.H. Chew, B. Miao and S. Zhong, Partial ambiguity, *Econometrica* 85 (2017), 1239-1260.
- [11] B. De Finetti, La prevision: ses lois logiques, ses sources subjectives. *Ann. Inst. H. Poincare* 7 (1937), 1-68. English translation in *Studies in Subjective Probability*, 2nd edition, H.E. Kyburg and H.E. Smokler eds., Krieger Publishing, Huntington NY, 1980, pp. 53-118.
- [12] M. Dean and P. Ortoleva, Is it all connected: a testing ground for unified theories of behavioral economics phenomena, 2016.
- [13] P.M. De Marzo, D. Vayanos, and J. Zweibel, Persuasion bias, social influence and unidimensional opinions, *Quart. J. Econ.* 118 (2003), 909-968.
- [14] J. Dow and S.R. Werlang, Uncertainty aversion, risk aversion, and the optimal choice of portfolio, *Econometrica* 60 (1992), 197-204.
- [15] J. Eichberger, J. Oechssler, and W. Schnedler, How do subjects view multiple sources of ambiguity? *Theory and Decision* 78 (2015), 339-356.
- [16] K. Eliaz and P. Ortoleva, Multidimensional Ellsberg, *Manag. Sc.* 62 (2016), 2179-2197.
- [17] A. Ellis and M. Piccione, Correlation misperception in choice, *Amer. Econ. Rev.* 107 (2017), 1264-1292.
- [18] D. Ellsberg, Risk, ambiguity, and the Savage axioms, *Quart. J. Econ.* 75 (1961), 643-669.
- [19] B. Enke and F. Zimmermann, Correlation neglect in belief formation, *Review of Economic Studies*, forthcoming.

- [20] L.G. Epstein and M. Schneider, Ambiguity and asset markets, *Annual Rev. Finan. Econ.* 2 (2010), 315-346.
- [21] L. G. Epstein and K. Seo, Symmetry of evidence without evidence of symmetry, *Theor. Econ.* 5 (2010), 313-368.
- [22] L. G. Epstein and K. Seo, Exchangeable capacities, parameters and incomplete theories, *J. Econ. Theory* 157 (2015), 879-917.
- [23] H. Ergin and F. Gul, A theory of subjective compound lotteries, *J. Econ. Theory* 144 (2009), 899-929.
- [24] E. Eyster and G. Weizsacker, Correlation neglect in financial decision-making, 2010.
- [25] J.P. Fouquet, C.S. Pun, and H.Y. Wong, Portfolio optimization with ambiguous correlation and stochastic volatilities, *SIAM J. Control Optim.* 54 (2016), 2309-2338.
- [26] I. Gilboa and D. Schmeidler, Maxmin expected utility with non-unique prior, *J. Math. Econ.* 18 (1989), 141-153.
- [27] B. Gillen, E. Snowberg and L. Yariv, Experimenting with Measurement Error: Techniques with Applications to the Caltech Cohort Study, 2016.
- [28] B. Greiner, Subject pool recruitment procedures: organizing experiments with ORSEE, *J. Econ. Sc. Assoc.* (2015), 114-125.
- [29] P. A. Haile and E. Tamer, Inference with an incomplete model of English auctions, *J. Pol. Econ.* 111 (2003), 1-51.
- [30] Y. Halevy, Ellsberg revisited: an experimental study, *Econometrica* 75 (2007), 503-536.
- [31] Y. Halevy, Time consistency: stationarity and time invariance, *Econometrica* 83 (2015), 335-352.
- [32] H.H. Huang, S. Zhang and W. Zhu, Limited participation under ambiguity of correlation, *J. Finan. Markets* 32 (2017), 97-143.
- [33] J. Jiang and W. Tian, Correlation uncertainty, heterogeneous beliefs and asset prices, 2016.

- [34] G. King, Proper nouns and methodological propriety: pooling dyads in international relations data, *Internat. Organization* 55 (2001), 497-507.
- [35] P. Klibanoff, M. Marinacci and S. Mukerji, A smooth model of decision making under ambiguity, *Econometrica* 73 (2005), 1849-1892.
- [36] P. Klibanoff, S. Mukerji and K. Seo, Perceived ambiguity and relevant measures, *Econometrica* 82 (2014), 1945-1978.
- [37] D. Kreps, *Notes on the Theory of Choice*, Westview Press, 1988.
- [38] D. Kreps and E. Porteus, Temporal resolution of uncertainty and dynamic choice theory, *Econometrica* 46 (1978), 185-2000.
- [39] G. Levy and R. Razin, Correlation neglect, voting behaviour and information aggregation, *Amer. Econ. Rev.* 105 (4) (2015), 1634-1645.
- [40] G. Levy and R. Razin, Combining forecasts: why decision makers neglect correlation, 2017.
- [41] J. Liu and X. Zeng, Correlation ambiguity and Under Diversification, 2016, [https://papers.ssrn.com/sol3/Papers.cfm?abstract\\_id=2692692](https://papers.ssrn.com/sol3/Papers.cfm?abstract_id=2692692).
- [42] F. Maccheroni and M. Marinacci, A strong law of large numbers for capacities, *Ann. Prob.* 33 (2005), 1171-1178.
- [43] F. Maccheroni, M. Marinacci and A. Rustichini, Ambiguity aversion, robustness and the variational representation of preferences, *Econometrica* 74 (2006), 1447- 1498.
- [44] M. Machina and D. Schmeidler, A more robust definition of subjective probability, *Econometrica* 60 (1992), 745-780.
- [45] L.J. Savage, *The Foundations of Statistics*, Dover, New York, 1972.
- [46] D. Schmeidler, Subjective probability and expected utility without additivity, *Econometrica* 57 (1989), 571-587.
- [47] U. Segal, The Ellsberg paradox and risk aversion: an anticipated utility approach, *Int. Econ. Rev.* 28 (1987), 175-202.
- [48] U. Segal, Two-stage lotteries without the reduction axiom, *Econometrica* 58 (1990), 349-377.

- [49] K. Seo, Ambiguity and second-order belief, *Econometrica* 77 (2009), 1575-1605.
- [50] E. Tamer, Partial identification in econometrics, *Annual Rev. Econ.* 2 (2010), 167–195.
- [51] E. Tamer, Incomplete simultaneous discrete response model with multiple equilibria, *Rev. Econ. Stud.* 70 (2003), 147-165.
- [52] S. Trautman and G. van de Kuilen, Ambiguity attitudes, in G. Keren and G. Wu (eds.), *The Wiley Blackwell Handbook of Judgement and Decision Making*, Blackwell, 2015.
- [53] A. Tversky and C. R. Fox, Weighing risk and uncertainty, *Psych. Rev.* 102 (1995), 269-283.
- [54] A. Tversky and P. P. Wakker, Risk attitudes and decision weights, *Econometrica* 63 (1995), 1255-1280.

## A Appendix

### A.1 Design of earlier experiments (2013 and 2014)

We conducted two earlier experiments (in 2013 and 2014) to study the hypothesized behavior of One vs Two. The two experiments had a similar main treatment to the experiment reported in the body of the paper, although in the 2014 experiment we changed the way we handled symmetry (which became a redundant assumption due to the design of the 2015 experiment) and we improved the explanation of the incentive system. More importantly, in the 2014 experiment we added control treatments. In describing the design we will focus on the 2014 experiment as it relates to the 2015 experiment reported in the body of the paper, and note any differences from the 2013 experiment.

#### A.1.1 Subjects, stakes and number of balls

Subjects were recruited from UBC’s Vancouver School of Economics subject pool using ORSEE (Greiner, 2015). We made sure that no subject participated in more than a single treatment during all iterations of the experiment (2013-2015). In addition to a \$10 participation fee, they could earn up to \$51 (\$101 in the 2013 experiment) during the one hour experiment. That is, the absolute value of the stakes in the earlier experiments was twice (2014) and four times (2013) higher than in the 2015 experiment. However, the marginal incentive was constant (at \$1) through all the experiments. This implied that the proportional premium required to exhibit strict preference is higher in the later experiments than in the earlier one. In other words, the strict ambiguity sensitivity revealed in the later experiments is “stronger” (in relative terms) than in our earlier experiments.

A minor difference between the 2013 experiment and later experiments is that in the first experiment each urn contained 10 balls, while in later experiments we used only 2 balls.

#### A.1.2 Choice problems

Choice problems were organized in pairs which allowed us to infer strict ranking from choices by slightly varying the prizes. Compared to the pairwise choice design used in 2015, there were no pairs with equal payments (which allows us to infer weak preference). Instead, lack of revealed strict

preference (choosing the bets with the higher payment) is interpreted as indifference. The design establishes an upper (lower) bound on the indifference (strict preference) class. Complete instructions may be found in the Online Appendix.

In the 2014 experiment, subjects were presented with the following choice problems.<sup>26</sup>

***One vs Different:*** Choose between a bet on the color of the ball drawn from a single urn and a bet that pays if the balls drawn from the two urns have different colors.

***One vs Same:*** Choose between a bet on the color of the ball drawn from a single urn and a bet that pays if the balls drawn from the two urns have the same color.

***Same vs Different:*** Choose between a bet that pays if the balls drawn from the two urns have the same color and a bet that pays if they have different colors.

**Color symmetry:** Choose between bets on the two possible colors of the ball drawn from a single urn.

**Standard Ellsberg:** Choose between a bet on the color of the ball drawn from one of the two ambiguous urns and a bet on the color of a ball drawn from a third urn containing one red ball and one blue ball (a risky urn as in Ellsberg).

### A.1.3 Control treatments (2014)

The 2014 experiment included two control treatments that were designed to investigate our hypothesis that the preference for *One* over *Two* indicates subjects' ambiguity concerning the relation between the two urns. An alternative hypothesis is that a bet on the ball drawn from a single urn is preferred because it is viewed as "simple," while a bet that depends on two balls drawn from the two urns is more "complex." In order to test this hypothesis, two urns – each with one red and one black ball, were presented to subjects.

---

<sup>26</sup>The 2013 experiment included two additional choice problems which are described in the Online Appendix, but did not include the color symmetry choice problems which were included in the 2014 experiment.

There is no ambiguity in this environment, and the probabilities of winning when betting on *One*, *Same* and *Different* are all equal to 0.5. However, the calculation of probabilities in the latter two bets is more involved and requires that the subject multiply probabilities correctly (reduction of compound objective lotteries - ROCL). There exists strong experimental evidence (Halevy 2007; Abdellaoui, Klibanoff and Placido 2015; Chew, Miao and Zhong 2017; Dean and Ortoleva 2016; Gillen, Snowberg and Yariv 2016) that violation of reduction is frequent in the general population, so choosing *One* over *Two* might be a manifestation of violation of ROCL. If the frequency of indifference in *One* vs *Two* is significantly higher in the risk control than in the two-urns ambiguity treatment, we conclude that the behavior of *One* over *Two* reflects something beyond the violation of reduction.

We implemented also a second control for two reasons, the first being that indifference in the risk control above is extremely fragile. Secondly, we wanted to examine whether our interpretation of the preference *One* over *Two* can be supported also when the environment is ambiguous. As noted above, there exists empirical evidence that relates ambiguity aversion and violations of ROCL, and such violations in turn may be responsible for a preference to bet on one urn. In this control, two draws were made with replacement from a single *ambiguous* urn containing two balls (with unknown composition). If a subject exhibits  $R_1 \succ \textit{Same}$  in this environment, it cannot be because of ambiguity about how the urns are related since the two draws are made from the same (ambiguous) urn and hence from the identical composition.<sup>27</sup> However, our interpretation of *One* over *Two* would be supported if the frequency with which subjects exhibit the ranking  $R_1 \succ \textit{Same}$  is significantly higher in the two-urns treatment than in this single-urn control.

#### A.1.4 Symmetry in colors

The 2013-2014 experimental designs relied on color symmetry ( $R_i \sim B_i$  for  $i = 1, 2$ ) to identify the behavior in *One* vs *Two*, since we asked subjects to pick a single color to bet on. As a result, a Bayesian subject who holds non-symmetric beliefs may strictly prefer *One* over *Same* and *Diff*. Though we believe that symmetry in colors is a very natural assumption in the current setting, we took two alternative approaches to handling the issue empirically. In the 2013 experiment, before subjects made any payoff-relevant choices,

---

<sup>27</sup>Note that in this environment,  $R_1 \succeq \textit{Diff}$  is easily rationalized, for example, by a Bayesian model. Thus we focus on the the choice between  $R_1$  and *Same*.

they were presented with four pairs of bets. The symmetry between urns and between colors in the experiment implied indifference within each pair. Following the presentation of the four pairs, each subject was asked, in a non-incentivized question, whether she agreed with the indifferences. Since symmetry serves as an identifying restriction to evaluate various models in light of the data, we thought it important to clarify this assumption to subjects. Their response (agree/disagree) indicates their understanding of the symmetry encoded in the experiment. We hypothesized, based on other studies, that most of the subjects will hold symmetric beliefs anyway. However, we cannot exclude the possibility that this form of communication manipulated some subjects with non-symmetric beliefs into symmetry.<sup>28</sup> Regardless, we believe that this is a desirable outcome since otherwise one would interpret behavior in light of an identifying assumption (symmetry in colors) that may not hold. In the 2014 experiment we adopted the more conservative approach of asking subjects to make pairwise choices between bets on the two colors, thus identifying possible non-symmetry directly and eliminating any potential contamination of the other choice problems.

### **A.1.5 Incentive structure and hedging**

As in the 2015, payment was determined by randomly choosing one choice problem *before* subjects made their choices. Also, each subject was asked to choose at the beginning of the experiment an urn (1 or 2) to bet on.

An important distinction from the 2015 experiment is that subjects did not bet on both colors: the subject chose a color (red or black) to bet on in the bets that involve a single urn. As a result, there is no risk that choices consistent with One  $\succ$  Two actually reflect hedging rather than preference for a bet on bias to a bet on relation. The empirical question whether subjects hedged in spite of our implementation is discussed further below in Section [A.2.4](#).

## **A.2 More Results (2013-2015)**

A total of 80 subjects participated in 4 sessions of the 2013 experiment. Subjects were paid a total of CA\$4,851 (an average of just over \$60 per sub-

---

<sup>28</sup>One may argue that the “instruction” of symmetry may have affected choices made elsewhere. However, we believe that such a scenario is highly improbable and we could not find evidence to that effect in the 2014 experiment.

ject). Out of the 80 subjects, a total of 24 were removed from the analysis reported immediately below. In particular, 18 subjects were removed due to choices violating either monotonicity or transitivity, and 4 were removed for disagreeing with symmetry over colors and urns (see below).<sup>29</sup> This leaves 56 subjects. We consider this retention rate to be high when taking into account the strong consistency (transitivity and monotonicity) that we imposed. We attribute this rate to the high stakes (more than \$100) employed in the experiment, which provided subjects sufficient incentive to minimize arbitrariness and to consider their choices seriously.

87 subjects participated in the main treatment of the 2014 experiment, of which 75 had no monotonicity or transitivity violations. The risk control included 47 subjects, of which 43 had no monotonicity violations. The single-urn control included 42 subjects, of which 37 had no monotonicity or transitivity violations. Subjects were paid a total of CA\$6,734 (an average of about \$38.26 per subject).

All the results reported are robust to employing a less strict retention criteria as reported in Appendix [A.2.5](#).

### A.2.1 Control treatments (2014)

In the risk control, 23 out of 43 subjects (53.5%) did not exhibit strict preference in choosing between *One* vs *Same* and *One* vs *Diff*. In the two-urn treatment we find that 28 out of 78 subjects (35.9%) exhibited this choice pattern in this two choice problems. The difference is significant at 5% ( $p$ -value of one-sided Z test is 0.03<sup>30</sup>). The finding that almost one half of subjects in the risk control did not reduce objective compound lotteries is consistent with the experimental evidence noted above. The fact that the frequency of indifference is significantly lower in the two-urn treatment indicates that choices in *One* vs *Two* are affected by more than the “simplicity” of a bet that depends on a single draw.

---

<sup>29</sup>We removed 11 for choosing in at least one pair of questions two lotteries with prizes of \$100, and 7 due, for example, to choices revealing that  $Same \prec R_1 \prec Diff$  based on pairwise choice between  $R_1$  and the other two bets but choosing consistently with  $Same \succ Diff$  in the direct pairwise choice. Two more subjects in the first session were caught cheating and their choices were excluded from the analysis (one of them had non-transitive choices, and would have been removed in any case).

<sup>30</sup>A one-sided Fisher exact test, which is very conservative in comparing two binomial samples, yields a  $p$ -value of 0.046.

Ellsbergian Ambiguity Attitude	2013		2014	
	#	%	#	%
<b>strictly averse (nonPS)</b>	37	66.1	29	52.7
<b>strictly seeking (nonPS)</b>	4	7.1	0	0
<b>neutral (PS)</b>	15	26.8	26	47.3
<b>Total</b>	<b>56</b>	<b>100</b>	<b>55</b>	<b>100</b>

Table A.1: Ellsbergian ambiguity attitude 2013-2014

In the single-urn control, 8 out of 37 subjects (21.6%) with monotone and transitive choices strictly preferred *One* to *Same*, while in the two-urn treatment 34 out of 75 subjects (45.3%) with monotone and transitive choices exhibited this pattern. This difference is significant at 5% ( $p$ -value of one-sided Z test is 0.007<sup>31</sup>).

In summary, the two control treatments indicate that although a bet on one urn is simpler and also may be attractive relative to the bets *Same* and *Diff* for subjects who do not reduce compound objective lotteries, its appeal is significantly strengthened when the relation between the compositions of the urns is ambiguous.

## A.2.2 Probabilistic sophistication

We now discuss the distribution of choices in the two behaviors–Ellsberg’s two-urn classic problem and *One* vs *Two*–that tested probabilistic sophistication. We restrict to subjects who did not exhibit strict color preference.<sup>32</sup> Although strict color preference does not necessarily imply non-symmetric beliefs, (even some subjects in the risk control exhibited this behavior), we decided to err on the side of caution and excluded such subjects when calculating adherence to probabilistic sophistication. Results including the other subjects are reported in Appendix A.2.5.

### Ellsbergian ambiguity (2013-2014)

Out of 111 subjects, almost 60% (66 subjects) exhibited strict Ellsbergian

<sup>31</sup> $p$ -value of one-sided Fisher exact test is 0.012.

<sup>32</sup>The fraction of subjects who exhibited strict color preference in the single ambiguous urn control is similar (9 out of 39 subjects) and is lower in the risk control (6 out of 44 subjects).

<b>One vs Two</b>	<b>2013</b>		<b>2014</b>	
	<b>#</b>	<b>%</b>	<b>#</b>	<b>%</b>
<b>One <math>\succ</math> Two (nonPS)</b>	23	41.1	18	32.7
<b>Two <math>\succ</math> One (nonPS)</b>	4	7.1	6	10.9
<b>One <math>\sim</math> Two</b>	20	35.7	26	47.3
<b>non-symmetric beliefs (PS)</b>	9	16.1	5	9.1
<b>Total</b>	<b>56</b>	<b>100</b>	<b>55</b>	<b>100</b>

Table A.2: One vs Two: 2013-2014

ambiguity aversion when asked to choose between betting on a red draw from the risky urn (with probability 0.5 of winning) and betting on a chosen color from one of the urns with unknown composition. Almost 37% (41 subjects) chose in a way that does not reveal ambiguity aversion or seeking, and the remaining 4 subjects exhibited ambiguity seeking. That is, about 63% of the subjects were not probabilistically sophisticated in a standard Ellsberg experiment. This proportion is consistent with existing studies that use certainty equivalent elicitation or choice data.<sup>33</sup>

### One versus Two (2013-2014)

There were 41 subjects (37% of 111 subjects) who exhibited One $\succ$ Two, and 10 subjects who exhibited Two $\succ$ One (7 of them ambiguity averse in Ellsberg). These 51 subjects (46% of 111) violated probabilistic sophistication.<sup>34</sup> The choices of the remaining 60 subjects (54% of the 111 subjects) can be rationalized by probabilistic beliefs. 2 subjects exhibited *Same*  $\succ$   $R_1$   $\succ$  *Diff*, and 12 subjects exhibited *Diff*  $\succ$   $R_1$   $\succ$  *Same* (11 of the latter exhibited Ellsbergian ambiguity aversion).

Combining both the Ellsberg and One vs Two choice problems, we find that only 27 out of the 111 subjects made choices that are consistent with probabilistic beliefs. That is, more than 75% were not probabilistically sophisticated in at least one of the choices.

<sup>33</sup>Note that 41 is an upper bound on the number of ambiguity neutral subjects—the proportion of ambiguity neutral subjects may be even smaller since the increment of \$1 used in the experiment to detect strict preference may have been too big for some subjects.

<sup>34</sup>Remember that we restricted the sample to subjects who agreed with the symmetry statements in the first experiment and were consistent with color symmetry in the second experiment.

Ellsbergian Ambiguity	One vs Two			Total
	nonPS	Symmetric PS	non-symmetric PS	
nonPS	20	20	23	63
PS	15	40	11	66
<b>Total</b>	35	60	34	129
Fisher exact test $p$ -value=0.003				

Table A.3: Ellsbergian ambiguity and One vs Two: Pairwise Choices (2015)

### A.2.3 Relation between attitude to sources of ambiguity

#### 2015

Tables A.3 and A.4 summarize the association between attitude to Ellsbergian ambiguity and choices made in One vs Two in the two elicitation methods. Each cell counts the number of subjects who exhibit the profile of behaviors.

The rows in both tables correspond to probabilistically sophisticated behavior in the standard Ellsberg experiment, where the top row counts the subjects who are either strictly Ellsbergian ambiguity averse or seeking (nonPS), while the second row counts all the rest (PS).

The columns correspond to probabilistic sophistication in One vs Two. The leftmost column corresponds to subjects who are not probabilistically sophisticated (nonPS, exhibiting either One $\succ$ Two or Two $\succ$ One). The other columns include subjects that made choices consistent with probabilistic sophistication. The middle column in Table A.3 includes only subjects who did not exhibit any strict preference in One vs Two, while the right column includes probabilistically sophisticated subjects who chose at least one bet with a lower prize (and hence revealed strict preference) in One vs Two. Fisher exact test for association reveals a tight association between behaviors in the two problems, but it is important to note that it is not that failure of probabilistic sophistication in one problem is associated with similar behavior in the other problem (this effect is not significant here); rather, probabilistically sophisticated behavior in the Ellsberg problem is associated with *symmetric* probabilistically sophisticated behavior in the One vs Two. In other words: subjects who always chose the bet with the higher prize in one problem tended to do the same in the other problem as well.

Table A.4 presents the association based on choice list elicitation. Since

Ellsbergian Ambiguity	One vs Two		Total
	nonPS	PS	
nonPS	28	12	40
PS	15	19	34
<b>Total</b>	43	31	74
Fisher exact test $p$ -value=0.034			

Table A.4: Ellsbergian ambiguity and One vs Two from Choice Lists (2015)

Ellsbergian Ambiguity	One vs Two			Total
	nonPS	Symmetric PS	non-symmetric PS	
nonPS	37	20	13	70
PS	14	26	1	41
<b>Total</b>	51	46	14	111
Fisher exact test $p$ -value<0.0001 <sup>1</sup>				

<sup>1</sup>  $p$ -values of Fisher exact tests in the 2013 and 2014 experiments (separately) are 0.008 and 0.113, respectively.

Table A.5: Ellsbergian ambiguity and One vs Two in 2013-2014 experiments

the tendency to make probabilistically sophisticated non-symmetric choices in One vs Two is much lower, we can combine the two probabilistically sophisticated columns. Here we find that failure of probabilistic sophistication in Ellsberg is significantly associated with (but distinct from) violation of probabilistic sophistication in One vs Two.

### 2013-2014

As in 2015, when distinguishing between forms of probabilistic sophistication in One vs Two the association is significant, but when pooling the two forms of probabilistic sophisticated behavior the association is less strong (one sided  $p$ -value=0.043).

When pooling all 240 subjects (2013-2015) who participated in the pairwise choice experiment, the association between probabilistic sophisticated behavior in the two domains is significant (one side  $p$ -value = 0.008), but the behaviors are far from being identical.

#### A.2.4 RIS and hedging

As discussed earlier, the standard use of RIS in ambiguity experiments (where the selection of the choice problem that counts for payment is performed *after* choices have been made and balls have been drawn) is open to the criticism that it could provide the subject with an opportunity to hedge the uncertainty using the randomization device employed in the incentive system. This implementation of the RIS induces a choice problem in which the subject faces Anscombe-Aumann acts and thus where she can hedge ambiguity using the randomization device employed in the incentive system by choosing certain combinations of ambiguous bets.<sup>35</sup> For this reason we performed the randomization *before* subjects made their choices and the draws from the urns were made. It is an empirical question whether subjects conformed with this order when evaluating lotteries, and if they did not, whether they hedged. These aspects of our experimental design as well as the empirical question whether there is evidence that subjects hedged in spite of our theoretically incentive compatible implementation are discussed in this Appendix.

Hedging might be revealed in different ways. For example, consider the two choices comprising One vs Two (in 2013-2014), and suppose, for simplicity, that the individual attaches probability 1/2 to payoffs being dependent on each of these questions. Then, a subject who is ambiguity averse and acts as if the randomization occurs after the balls are drawn, could choose *Same* over  $R_1$  in one question and *Diff* over  $R_1$  in another, and be left with the same state-independent payoff (in expected utility units) that she would obtain from betting on drawing red from Ellsberg’s risky urn. Accordingly, if she is ambiguity averse in the Ellsberg problem, then the preceding combined choices would be preferable to  $R_1$  and we would observe  $\text{Two} \succ \text{One}$ . However the latter is observed for only 15 subjects (6.25% of subjects) in the three experiments (2013, 2014 and 2015).

Another possibility for hedging that existed only in 2015 was to make

---

<sup>35</sup>The significance we are attaching to when the randomization takes place runs counter to Anscombe and Aumann’s well-known “reversal of order axiom.” However, the descriptive appeal of that axiom has been put into question by evidence that subjects who are ambiguity sensitive usually do not satisfy the reduction of compound lotteries assumption (Halevy 2007; Abdellaoui, Klibanoff and Pacido 2015; Chew, Miao and Zhong, 2017; Dean and Ortoleva 2016; Gillen, Snowberg and Yariv 2016). In a more normative vein, Seo (2009) argues for a relaxation of the “reversal of order” axiom and shows how it can be part of an axiomatic characterization of the “smooth ambiguity model.”

choices consistent with One $\succ$ Two (choose  $R_1$  and  $B_1$ ). This possibility was not available to subjects in the 2013-2014 experiments, where they were asked to choose a color to bet on in advance. The fact that the frequency of Ellsbergian ambiguity averse subjects making choices consistent with One $\succ$ Two in 2013-2014 is almost three times higher than in 2015 (28/111 compared with 12/129) suggests that this cannot be a major concern in our experiment.

Similarly, in the questions that compare *Same* to *Diff* directly (in 2013-2014), hedging would imply the absence of a strict preference between *Same* and *Diff*. This follows since if the subject had perceived the problem as if the randomization used for incentives occurred after the balls are drawn from the urns, then choosing the bet with a prize of \$51 (\$101 in the 2013 experiment) in each of the two questions would leave the subject with a lottery that pays \$51 (\$101) with probability 1/2 independently of the state, thus completely hedging the ambiguity. However, empirically about 36% of the subjects exhibit a strict preference in these questions,<sup>36</sup> and close to half (about 49%) of the subjects who are ambiguity sensitive in the Ellsberg choice problem (and thus may wish to hedge), made choices inconsistent with indifference between *Same* and *Diff*.

We conclude that we do not find convincing empirical evidence that subjects used the RIS to hedge any source of ambiguity in our experiment.

### A.2.5 Weaker inclusion criteria (2013-2014)

Throughout the paper we excluded subjects who made choices that were not consistent with transitive and monotone preferences in one of the choice problem. In 2013-2014 we also excluded subjects who did not agree with the symmetry statements (2013) or exhibited strict preference to betting on one of the colors.<sup>37</sup> The number of subjects excluded decreased as we progressed, since the instructions they received improved. Footnote 9 in the paper details the 24 subjects excluded in 2015, and below we describe the subjects excluded from the analysis of the previous two experiments. In this Appendix we show that their inclusion does not substantially affect our conclusions.

As noted above, out of 80 subjects who participated in the 2013 experiment, 24 were removed from the analysis: 11 were removed due to violations of monotonicity or transitivity in at least one pair of questions, 7 due to

---

<sup>36</sup>7 subjects strictly preferred *Same*, 33 subjects strictly preferred *Diff*, and the remaining 71 subjects made choices consistent with indifference between *Same* and *Diff*.

<sup>37</sup>This is not a necessary assumption in the 2015 experiment.

Ellsbergian Ambiguity Attitude	2013		2014	
	#	%	#	%
<b>strictly averse (nonPS)</b>	51	68.9	45	52.9
<b>strictly seeking (nonPS)</b>	6	8.1	3	3.5
<b>neutral (PS)</b>	17	23	37	43.6
<b>Total</b>	<b>74</b>	<b>100</b>	<b>85</b>	<b>100</b>

Table A.6: Ellsbergian ambiguity attitude 2013-2014: maximal inclusion

One vs Two	2013		2014	
	#	%	#	%
<b>One <math>\succ</math> Two (nonPS)</b>	31	40.2	30	38.5
<b>Two <math>\succ</math> One (nonPS)</b>	6	7.8	10	12.8
<b>One <math>\sim</math> Two</b>	25	32.5	28	35.9
<b>non-symmetric beliefs (PS)</b>	15	19.5	10	12.8
<b>Total</b>	<b>77</b>	<b>100</b>	<b>78</b>	<b>100</b>

Table A.7: One vs Two 2013-2014: maximal inclusion

cyclic choices between  $R_1$ , *Same* and *Diff*, 4 for disagreeing with the symmetry over colors and urns, and two more subjects in the first session were caught cheating and their choices were excluded from the analysis. Out of 87 subjects who participated in the main treatment of the 2014 experiment, 75 had no monotonicity or transitivity violations and 20 showed strict color preference.

In order to include as many subjects as possible from the 2013 experiment, several relaxations of the inclusion criteria were employed. First, the answer to the non-incentivized question concerning symmetry was ignored. Second, in case of cyclic choices between  $R_1$ , *Same* and *Diff*, the direct comparison between *Same* and *Diff* was not taken to invalidate the choices made in One vs Two (which relies on comparing  $R_1$  to *Same* and  $R_1$  to *Diff*). Third, if a subject had non-monotone choices in only one pair of questions (assuming transitivity), (s)he was not removed from the analysis. Instead, choices in other questions, together with transitivity, were used in order to

extend the preferences to the suspect direct comparison.<sup>38,39</sup> We used the first two methods for the 2014 experiment as well. The distributions of Ellsbergian ambiguity and One vs Two are presented in Tables A.6 and A.7.

The marginal distributions of ambiguity attitude expressed in the standard Ellsberg experiment and in One vs Two are remarkably similar to those for the smaller sample, taking into account that the bigger sample includes subjects with non-symmetric color preference, which may affect the choices that comprised ambiguity attitude (seeking) and One  $\succ$  Two.

### A.3 Details for the source model

This appendix provides supporting details for the model in Section 5.2.

To define utility precisely, let  $u_1$ ,  $u_2$  and  $u_3$  be strictly increasing vNM indices that will apply to the three issues respectively. Let the cdf  $F$  describe conditional beliefs about  $p$ , generalizing from the uniform distribution in Figure 5.1. Then, for any bet (or act)  $f$  over  $S_1 \times S_2$ , its utility (in certainty equivalent units) is computed recursively by:<sup>40</sup>

$$U(f) = u_1^{-1}(\sigma u_1(V_2(f)) + (1 - \sigma)u_1(W_2(f))),$$

$$V_2(f) = u_2^{-1}\left(\int_p u_2(V_3(f;p)) dF\right), W_2(f) = u_2^{-1}\left(\int_p u_2(W_3(f;p)) dF\right),$$

$$V_3(f;p) = u_3^{-1}\left(\int_{S_1 \times S_2} u_3(f) d[(p, 1-p) \otimes (p, 1-p)]\right), \text{ and}$$

$$W_3(f;p) = u_3^{-1}\left(\int_{S_1 \times S_2} u_3(f) d[(p, 1-p) \otimes (1-p, p)]\right).$$

---

<sup>38</sup>This applies to 5 subjects. An extreme example is provided by Subject 315: (s)he did not agree with the suggested symmetry in colors and urns, and her/his choices in  $R_1$  vs *Same* were inconsistent with monotone preferences. However, the choices in  $R_1$  vs *Diff* and *Diff* vs *Same* were consistent with  $R_1 \succ \text{Diff}$  and  $\text{Diff} \succ \text{Same}$ , so this subject was classified as exhibiting *One*  $\succ$  *Two*.

<sup>39</sup>We omitted only the two subjects who were caught cheating and another subject who made choices inconsistent with monotone preferences in both  $R_1$  vs *Same* and  $R_1$  vs *Diff*.

<sup>40</sup>Symmetry in urns is built into the model, and symmetry in colors is implied if we assume, as we do, that  $F$  is suitably symmetric (the compositions  $(p, 1-p)$  and  $(1-p, p)$  are equally likely).

Now specialize to uniform  $F$ ,  $\sigma = \frac{1}{2}$  and linear  $u_3$ , with  $u_3(x) = 1$  and  $u_3(0) = 0$ . Then the utility of the bet  $R_3$  on drawing red from the risky Ellsberg urn is given by

$$U(R_3) = u_3^{-1}\left(\frac{1}{2}\right) = \frac{1}{2},$$

and the utility of the bet  $R_1$  on drawing red from a single ambiguous urn is

$$U(R_1) = u_2^{-1}\left(\int_p u_2(p) dF\right).$$

Thus  $R_3$  is preferred if  $u_2$  is concave, as noted in the text. Moreover, One  $\succ$  Two is implied if  $u_1$  is more concave than  $u_2$ :

$$\begin{aligned} u_1 \circ U(\text{Same}) &= \frac{1}{2}u_1 \circ u_2^{-1}\left[\frac{2}{3}u_2(1) + \frac{1}{3}u_2\left(\frac{1}{2}\right)\right] + \frac{1}{2}u_1 \circ u_2^{-1}\left[\frac{2}{3}u_2(0) + \frac{1}{3}u_2\left(\frac{1}{2}\right)\right] \\ &\leq u_1 \circ u_2^{-1}\left[\frac{1}{2}\left(\frac{2}{3}u_2(1) + \frac{1}{3}u_2\left(\frac{1}{2}\right)\right) + \frac{1}{2}\left(\frac{2}{3}u_2(0) + \frac{1}{3}u_2\left(\frac{1}{2}\right)\right)\right] \\ &= u_1 \circ u_2^{-1}\left[\frac{1}{3}u_2(0) + \frac{1}{3}u_2\left(\frac{1}{2}\right) + \frac{1}{3}u_2(1)\right] = u_1 \circ U(R_1). \end{aligned}$$