Uncertainty Aversion in Game Theory: Experimental Evidence

Evan M. Calford*

Purdue University Economics Department Working Paper No. 1291

April 2017

Abstract

This paper experimentally investigates the role of uncertainty aversion in normal form games. Theoretically, risk aversion will affect the utility value assigned to realized outcomes while ambiguity aversion affects the evaluation of strategies. In practice, however, utilities over outcomes are unobservable and the effects of risk and ambiguity are confounded. This paper introduces a novel methodology for identifying the effects of risk and ambiguity preferences on behavior in games in a laboratory environment. Furthermore, we also separate the effects of a subject's beliefs over her opponent's preferences from the effects of her own preferences. The results support the conjecture that both preferences over uncertainty and beliefs over opponent's preferences affect behavior in normal form games.

Keywords: Ambiguity Aversion, Game Theory, Experimental Economics, Preferences

JEL codes: C92, C72, D81, D83

*Department of Economics, Krannert School of Management, Purdue University ecalford@purdue.edu; I am particularly indebted to Yoram Halevy, who introduced me to the concept of ambiguity aversion and guided me throughout this project. Ryan Oprea provided terrific support and advice throughout. Special thanks to Mike Peters and Wei Li are also warranted for their advice and guidance. I also thank Li Hao, Terri Kneeland, Chad Kendall, Tom Wilkening, Guillaume Frechette and Simon Grant for helpful comments and discussion. Funding from the University of British Columbia Faculty of Arts is gratefully acknowledged.

1 Introduction

In a strategic interaction a rational agent must form subjective beliefs regarding their opponent's behavior. But what form do these beliefs take? In a Nash equilibrium—resting on a bed of expected utility theory which does not allow for any uncertainty in beliefs–each agent has a consistent and precise belief over others behavior. Allowing for uncertainty over an opponent's choice of mixed strategy, which we shall refer to as strategic uncertainty, arises as an intuitive alternative to expected utility in games. It then becomes natural that an agent's attitude toward ambiguity will play a role in determining equilibrium behavior. We begin from the assertion that strategic uncertainty is the natural condition of strategic interactions.

There is a well-developed theoretical literature on ambiguity aversion in games (see Lo (2009), Dow and Werlang (1994), Epstein (1997) or Eichberger and Kelsey (2000), for example) that provides guidance on how agents *should* respond to strategic uncertainty. But how *do* people respond to strategic ambiguity? Do people behave as if they have unique probabilistic beliefs over their opponent's strategies, or do they behave as if they are ambiguity averse (perhaps in a fashion consistent with evidence found in individual decision making experiments)? Can people identify when their opponent is facing strategic ambiguity and, if so, do they respond rationally? If not, why not? We use experimental methods to provide answers to these questions; answers that have important implications for both applications, and development, of the theory of ambiguity aversion in games.

The form of strategic uncertainty studied here goes beyond uncertainty over which of several possible equilibria may be played, or uncertainty regarding types. The main testing game used in the experiment has, under an assumption that preferences are commonly known to satisfy subjective expected utility, a unique rationalizable outcome. The same game, with the same payoffs over *outcomes* but an assumption that preferences are commonly known to exhibit ambiguity aversion, has multiple rationalizable outcomes.

The experimental approach requires, therefore, a measurement of subject preferences. A concern for potential framing effects necessitates using only games to measure preferences. Risk and ambiguity preferences are measured using a pair of classification games, and the design of the testing game allows the effects of risk and ambiguity aversion to be separated without the use of additional elicitations of beliefs.[1]

[1] Ivanov (2011) was the first paper to separate the effects of risk and ambiguity in games, via the use of non-incentivized elicitation data to establish beliefs over opponent strategies. Here we provide an alternative approach

There is, however, a fundamental problem with direct inference of ambiguity aversion from behavior in games: risk aversion. How does a subject's ambiguity preference affect their behavior in the presence of risk aversion? In the standard theory, at least, the separation is straightforward: for a game where subjects earn monetary payoffs, we first take a monotonic transformation of the payoffs to move from a money space to a utility space. Then, ambiguity aversion acts to affect the way in which a subject evaluates her strategies which earn utility denominated payoffs. In practice, however, the two effects are often much more difficult to disentangle, and this paper is the first to tackle this separation in a game theoretic setting without eliciting any data beyond the choice of strategies in games. Because risk and ambiguity aversion have similar effects in games (making 'safe' strategies appear relatively more attractive), and are positively correlated, studies that focus only on risk aversion *or* ambiguity aversion in games will be prone to omitted variable bias.

The assumption of common knowledge of preferences that is embedded in the use of rationalizability in this paper is strong. Fortunately, the design of the testing game does not require the full weight of common knowledge: first order beliefs are sufficient to form sufficiently strong hypotheses. To facilitate inference, we include a treatment where first order beliefs are induced by *showing* some subjects their opponent's decisions in the preference measurement tasks. This credible preference information anchors the beliefs of subjects (without compromising incentive compatibility) and allows for clean inference of the effects of beliefs regarding other's uncertainty preferences on behavior.

The testing game used in the experiment is designed so that the set of rationalizable strategies varies with preferences, allowing for a partial separation of behavior as a function of preferences. The testing game also allows for a separation of the effects of ambiguity preferences from the effects of beliefs over an opponents' ambiguity preferences. These dual separations allow a detailed investigation of the role of uncertainty in normal form games, providing answers to our primary research questions.

Taken as a whole, our findings provide support for the claim that uncertainty preferences are an important determinant of behavior in games. We find evidence of both a first order effect (subjects own preferences affect their behavior) and second order effect (subjects opponent's preferences affect the subject's behavior) of uncertainty aversion on behavior in games. Furthermore, we find that the second order effect of risk aversion is stronger than the second order effect of ambiguity aversion, while the magnitude of the first order effects of risk aversion and ambiguity aversion are comparable.

that does not require subject beliefs over strategies to be observed.

There is a wealth of both theoretical and empirical evidence, tracing back to Knight (1921) and Keynes (1921) via Ellsberg (1961) and Halevy (2007), of ambiguity affecting decisions in individual decision making environments. The relative paucity of experimental evidence on the role of ambiguity aversion in strategic environments was a key motivation for this study. The previous literature provides a series of snapshots into how subjects behave in the face of strategic uncertainty, and suggests that ambiguity aversion plays a key role in strategic decision making. Camerer and Karjalainen (1994) provides evidence that subjects, on average, prefer to avoid strategic uncertainty by betting on known probability devices rather than on other subjects' choices. Eichberger et al. (2008) establish that subjects find "grannies" to be a greater source of strategic ambiguity than game theorists. Kelsey and le Roux (2015) find that subjects exhibit higher levels of ambiguity aversion in games than in a 3-color Ellsberg urn task. The current paper is the first to give a complete picture of the role of uncertainty aversion in games: we document the first-order effects (how do subjects respond to strategic uncertainty?) and second-order effects (how do subjects respond to opponents who face strategic uncertainty?) of both risk aversion and ambiguity aversion.

The closest paper to this one is Ivanov (2011), which asks the dual of our research question by estimating ambiguity preferences from behavior in games (rather than focusing on the effects of preferences on behavior as is the case here). Methodologically, we have two key points of departure from Ivanov (2011). First, Ivanov requires non-incentivized elicited beliefs to separate risk aversion from ambiguity aversion, while the current paper collects data only from strategy choices in incentivized games. Second, Ivanov (2011) is able to identify ambiguity (and risk) aversion, neutrality and seeking preferences and, in fact, finds non-trivial levels of ambiguity seeking. The current paper does not distinguish between ambiguity seeking and ambiguity neutrality, instead choosing to focus on the role of ambiguity aversion.

This paper is also the first to provide a procedure for measuring preferences using discrete choice tasks in a framing that is consistent with typical normal form game experiments. Heinemann et al. (2009) also recognize the importance of using frame-consistent tasks to measure preferences and strategic uncertainty. In their case, they used a modified coordination game that was framed as a multiple price list to study strategic uncertainty through the lens of global games. Other papers that have elicited preferences to study behavior in games include Healy (2013) and Brunner et al. (2017), although both papers measure preferences over outcomes (i.e. ordered pairs of payments to each player in a game) rather than preferences over uncertainty. Brunner et al. (2017) also displays elicited preferences to subjects in a treatment similar in style to the other-preference treatment in this paper.

This paper proceeds as follows. Section 2 presents the experimental design and hypotheses. Section 3 presents the experimental results, and section 4 provides a discussion and conclusion. Proofs, additional results, and the instructions to subjects are gathered in an appendix.

2 Experimental design and hypotheses

The heart of the experimental design is straightforward: in the first stage we measure (in the own-preference treatment) or display (in the other-preference treatment) preferences, and in the second stage we examine whether preferences affect behavior in a carefully chosen testing game. We use rationalizability, rather than equilibrium concepts, to identify the expected relationship between preferences and play in the testing game because of the stronger epistemic assumptions required to justify the use of equilibrium. Epstein (1997) rationalizability, which we use here and discuss in further detail below, provides an extension of Pearce (1984) and Bernheim (1984) rationalizability that allows for non-neutral ambiguity preferences.

	X	Y
A	25,20	14,12
B	14,20	25,12
C	18,12	18,22

Figure 1: Testing game. Payoffs are in Canadian Dollars.

Consider the testing game as presented in Figure 1, noting some of the key features of the payoff structure. First, consider the case with risk and ambiguity neutral agents (i.e. use Pearce/Bernheim rationalizability): $\{A, X\}$ is the unique rationalizable outcome. To see this, notice that C is never a best response for the row player, and therefore eliminate C. Once C is eliminated, the unique best response for the column player is X, therefore eliminating Y. Finally, A is the best response to X.

Second, notice that the best response correspondence for the column player can be simply described as "play X if they believe the union of A and B is sufficiently more likely than C, and play Y otherwise." Furthermore, if the column player places 0 weight on their opponent playing C then X is the unique best response irrespective of the column player's risk and ambiguity preferences.

Third, note that although C is never a best response for a risk and ambiguity averse neutral row player, if the row player is sufficiently risk averse or ambiguity averse then C is a best response for at least some beliefs over the column player's behavior. To take an extreme example, a risk neutral ambiguity averse row player who has Gilboa and Schmeidler (1989) preferences and holds complete (subjective) uncertainty regarding the column player's action will value strategy C at 18 and strategies A and B at 14.

Taken as a whole, these three observations generate the intuition that underlies the experimental design: the rationalizable set, and hence expected behavior, is determined solely by the preferences of the row player and is unaffected by the preferences of the column player. We can therefore test for the own-preference effect by observing row player behavior in the testing game, and test for the other-preference effect by observing column player behavior in the same game. This logic is formalized in subsection 2.4.

2.1 The classification games: own-preference treatment

There are two classification games. The first game elicits ambiguity preferences, while the second game elicits risk preferences. In each game the row player selects between a set of prospects, whose payoffs depend only on exogenous random events, while the column player earns a positive payoff if and only if she correctly predicts the row player's choice. We are primarily interested in the row player behavior in these games: the column player is a necessary consequence of converting the preference measurement task into a game form. For completeness, column player behavior is presented in subsection B.2.

The first (ambiguity) classification game is shown in Figure 2. For the row player, this game is isomorphic to a standard Ellsberg task with a slight asymmetry in payoffs as recommended by Epstein and Halevy (2014). The payoff asymmetry ensures that an ambiguity neutral subject has a strict preference to play S.[2] The game involves two ball draws, one from the U urn (which contains red and yellow balls in unknown proportion, as depicted in Figure 3) and one from the K urn (which contains red and yellow balls in equal proportion, as depicted in Figure 4). Therefore there are four possible states of nature, but only two payoff tables. The left payoff table represents the state red ball drawn from the U urn and yellow ball drawn from the K urn: (R_U, Y_K). The right payoff table represents the state (Y_U, R_K). The payoffs for state (R_U, R_K) are found by adding the two payoff tables together, and the payoffs in state (Y_U, Y_K) are identically 0 for both players. The relationship between states and payoffs was carefully explained to the subjects, and understanding

[2] An implication of this is that agents with extremely slight ambiguity aversion may be erroneously classified as ambiguity neutral. This is superior to having the expected large number of ambiguity neutral subjects be indifferent between the two options. See Epstein and Halevy (2014) for further details.

was tested via a series of comprehension questions that are discussed in detail in subsection 2.2.

	S'	M'
S	$30.1, 15$	$30.1, 0$
M	$0, 0$	$0, 15$

Red ball drawn from U urn

	S'	M'
S	$0, 15$	$0, 0$
M	$30, 0$	$30, 15$

Red ball drawn from K urn

Figure 2: Classification game 1. This game is used to measure the row player's ambiguity aversion and the column player's belief of the row player's ambiguity aversion.

Figure 3: U urn. The U urn consists of 10 balls, each of which may be either red or yellow. The total number of red balls in the urn lies between 0 and 10.

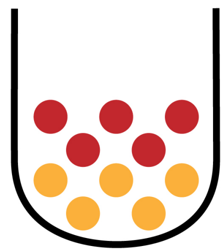

Figure 4: K urn. The K urn contains 5 red and 5 yellow balls.

Given that row player payoffs are independent of the column player strategy choice, we can view the row player as facing a choice between a bet that pays $30.10 if a red ball is drawn from the U urn and a bet that pays $30 if a red ball is drawn from the K urn. We assume that subjects hold symmetric beliefs about the distribution of balls in the U urn.[3] If a subject has Subjective Expected Utility (SEU) preferences, then they should strictly prefer strategy S (the bet on the U urn). A subject with ambiguity averse preferences should prefer strategy M (the bet on the K urn). We note that because the row player is indifferent to her opponent's strategy, the existence of the column player should have no effect on the row player's choices.

The second classification game is shown in Figure 5 and has a very similar structure to the first classification game, with the key difference being that the state is now determined by a single draw from the K urn. The row player chooses which risky prospect they would like to hold, and the column player attempts to predict the row player's preferences. Strategy L has both the highest

[3]Note that this is an assumption regarding the symmetry of beliefs, and not an assumption on the subjects preferences regarding red or yellow balls. If a subject does happen to prefer red balls over yellow balls, for whatever reason, there are still no confounding effects. Subjects may only bet on red balls, and a general preference for red would be equivalent to increasing the prize paid on a red ball being drawn an equal amount for each urn.

	L'	I'	H'
L	$25, 30$	$25, 0$	$25, 0$
I	$11, 0$	$11, 30$	$11, 0$
H	$15, 0$	$15, 0$	$15, 30$

Red ball drawn from K urn

	L'	I'	H'
L	$10, 30$	$10, 0$	$10, 0$
I	$23, 0$	$23, 30$	$23, 0$
H	$15, 0$	$15, 0$	$15, 30$

Yellow ball drawn from K urn

Figure 5: Classification game 2. This game is used to measure the row player's risk aversion and the column player's belief of the row player's risk aversion.

expected return and highest variance, while strategy H provides a lower but certain payoff. Strategy I has an intermediate expected return and variance.

Subjects that participated in the own-preference treatment played as both the row player and column player in both of the classification games as well as in the testing game. Only one game was chosen for payment, and the game to be used for payment was fixed at the beginning of each experimental session but not revealed to subjects until the end of the session. This payment protocol was used to mitigate any potential for subjects to hedge across games (see Azrieli et al. (2016) and Baillon et al. (2014) for details). In contrast, subjects in the other-preference treatment played only a single game (the testing game as the column player) and therefore had no means by which to hedge across games.

2.2 Comprehension questions

On the experimental screen, underneath each of the normal form games, a series of dynamic drop down menus were included for each game. Before a subject could confirm their strategy choice in a game, they were required to fill in the drop down menus correctly. To ensure that subjects took the drop down menus seriously they were paid a bonus of $1 for each game where they filled in the drop down menus correctly on their first attempt. Each incorrect attempt reduced the bonus payment, for that game, by $0.25.

The drop down menus were designed in such a way that the subjects recreated the worded decision problem that describes the relevant game. For example, for the game in Figure 2, the worded problem for option S would read "Your earnings for this choice [are/**are not**] affected by your counterpart's strategy. Your earnings for this choice will be [**$30.10**/$30/$15/$0] if a red ball is drawn from the [**U urn**/K urn] and nothing otherwise." For each set of terms that are square bracketed, the subjects were required to select the correct term (shown in bold here) from a drop down menu. Subjects were required to fill in drop down menus that described the possible outcomes for each of their strategies.

There were very few subjects that had substantial difficulty with the drop down menus, with just under 3% of subjects in the own-preference treatment earning half or less of the available bonus payment. We used the level of bonus payment earned by the subject as a measure of comprehension, and removed subjects that performed poorly from our sample. For our data analysis in section 3 we remove subjects that failed to achieve a perfect score for the games in which we used their data.[4] This requirement is fairly strict, was determined ex-ante, and was chosen because subjects face each game only once and therefore have no opportunity to learn the game structure via experience. The standard approach, to play the game multiple times and discard the first few rounds of data, cannot be used here because repeated play will cause subjects to learn something about the distribution of play which will alter the degree of strategic uncertainty involved in the game.

2.3 Other-preference treatment

The other-preference treatment was run subsequent to the own-preference treatment, and each subject was matched with a data point from the own-preference treatment. Each subject in the other-preference treatment was *shown* their opponent's choices as the row player in the classification games. Some subjects were shown their opponent's choices in both the risk and ambiguity games and some subjects were only shown their opponent's choice in the ambiguity classification game. The subjects were shown figures similar to Figure 2 and Figure 5 with the column player payoffs suppressed with their opponent's choice highlighted in light blue. After subjects had viewed their opponent's choices in the preference measuring tasks, subjects were asked to play the main testing game as the column player against this same opponent. The matching protocol was explained to subjects with the aid of the diagram in Figure 6: the subject is matched with a previous participant, and therefore the subject cannot influence the payoffs that are being awarded to any other player. The subjects were also required to fill in the drop down menu comprehension questions for their role as a column player, could earn up to a $1 bonus for completing the comprehension questions correctly, and lost $0.25 for each mistake made.

Subjects were informed that they were to be matched with a previous experimental participant, but were not given any information regarding how their opponent was selected. The actual

[4]For the own-preference treatment, a subject was removed if they made any mistakes as the row player in either the risk measurement game, the ambiguity measurement game or the testing game. In the other-preference treatments, where subjects only played one game, any mistake eliminated them from the sample.

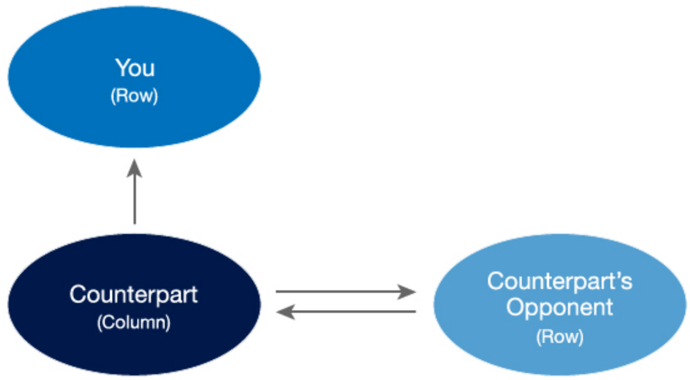

Figure 6: The matching protocol diagram, as shown to subjects. The subject was designated "You", and the subject's opponent was designated "Counterpart". Arrows are used to represent strategic interactions: the originator's choice influences the target's payoff. Note that the subjects were shown a transposed version of the game in Figure 1, so the row and column labels are reversed.

selection process used was two-tiered. In the first stage, a subset of data points from the own-preference treatment was identified for use in the other-preference treatment. The subsample was chosen to improve statistical power (i.e. appropriately balancing the number of ambiguity averse versus ambiguity neutral opponents) and was biased towards opponents that performed well on the comprehension tasks. Within this subsample, each subject was randomly allocated an opponent.

2.4 Hypothesis

In this section we formalize the interactions between preferences and rationalizability in the game in Figure 1. Although we use rationalizability rather than equilibrium concepts, similar conclusions could be derived from equilibrium models.[5]

Throughout we assume that subjects choose from a set of pure strategies, and do not play mixed strategies. This is consistent with the experimental implementation where subjects were required to select a pure strategy choice for each game. The reason for this choice, which is also common amongst the ambiguity aversion in game theory literature, is that models of ambiguity aversion with well defined preferences over mixed strategies typically generate a strict preference for mixed strategies. Calford (2016) and Eichberger and Kelsey (2000) contain extensive discussion on the role of mixed strategies in games with ambiguity averse agents. As Calford (2016) demonstrates it is possible to extend the theory in this section to allow for subjects to use mixed strategies if the

[5]Compare Lo (2009) and Dow and Werlang (1994) with Epstein (1997), for examples.

mixed strategy randomization device is resolved before the game is played. Therefore, a subject who rolls a dice or otherwise randomizes their choice, before clicking on a pure strategy, could also be accommodated.

We use Epstein's (1997) notion of rationalizability, a generalization of the more familiar Pearce (1984)/Bernheim (1984) rationalizability, which relaxes SEU preferences to the Maxmin Expected Utility (MEU) preferences of Gilboa and Schmeidler (1989). Theorem 3.2 in Epstein (1997) establishes that the set of rationalizable strategies can be found by iteratively eliminating all strategies that are not best responses *given the (possibly ambiguity averse) preferences of the agents* for two player games. The only conceptual difference between the Pearce-Bernheim framework and the Epstein framework is the treatment of mixed strategies: Epstein (1997) restricts the feasible set of strategies to consist only of pure strategies, although agents may still hold beliefs over the mixed strategies of their opponents (following a population or belief-based interpretation of mixed strategies).

Consider Gilboa and Schmeidler's (1989) MEU preferences (for ambiguity averse agents), in which an agent's beliefs regarding her opponent's strategy is a closed and convex subset of probability measures over her opponent's strategy set. The agent then evaluates her utility for a given prospect with respect to her 'worst case' belief, and may use different beliefs to evaluate different prospects. For example, in our game, the set of feasible row player beliefs is a subset of the probability simplex ($\Phi_R \subseteq \Delta(\{X, Y\})$) and, given Φ_R, the row players' preferences can be represented by:

$$U(a_R) = \min_{\phi_R \in \Phi_R} \sum_{a_C \in \{X,Y\}} u_R(m_R(a_R, a_C))\phi_R(a_C) \qquad \forall a_R \in \{A, B, C\}\,, \tag{1}$$

where $u_R(.)$ is the row player's utility over monetary outcomes and $m_R(a_R, a_C)$ is the monetary payoff for outcomes as shown in Figure 1. Given these preferences, we interpret risk aversion, in the standard manner, as curvature of the utility function. We model ambiguity aversion in a binary fashion. For subjective expected utility subjects Φ_R is constrained to be a singleton ($\Phi_R \in \Delta(\{X, Y\})$), while for ambiguity averse subjects Φ_R is unconstrained ($\Phi_R \subseteq \Delta(\{X, Y\})$. Column player preferences are defined analogously.

Proposition 1. *Consider the normal form game in Figure 1. If the row player has preferences such that $u(25)+u(14) > 2u(18)$ and Φ_R is a singleton for all feasible beliefs then the rationalizable set is $\{(A, X)\}$. If $u(25)+u(14) \leq 2u(18)$ or Φ_R is unrestricted then all pure strategies are rationalizable.*

Proof of proposition 1. See appendix A. □

The condition $u(25)+u(14) \leq 2u(18)$ provides a necessary and sufficient condition on the utility function of a SEU agent for C to be undominated.

Notice that Proposition 1 does not restrict the column player's preferences: the column player's preferences play no role in determining the size of the rationalizable set. Also note that the use of MEU preferences to model ambiguity aversion is not essential: any of the standard models of ambiguity aversion could be used. We use MEU here because it is arguably the simplest model of ambiguity averse preferences for this game, and is emphasized in Epstein (1997). Alternatives such as Choquet Expected Utility (Schmeidler (1989)), which underlies Dow and Werlang (1994), or Klibanoff et al.'s (2005) smooth ambiguity aversion preferences would work just as well: in a recent paper Battigalli et al. (2016) independently illustrate this, for the case of smooth ambiguity aversion preferences, using an example on a similar game to the one studied here.

Proposition 1 establishes the rationalizable set for our testing game as a function of preferences. To finalize our hypothesis we require a mapping from the revealed preferences of the risk classification game to the antecedents of Proposition 1.

Lemma 1. *Suppose that an agent has a utility function that satisfies constant absolute risk aversion. If* $u(25)+u(10) \geq \max\{u(23)+u(11), 2u(15)\}$ *then* $u(25)+u(14) > 2u(18)$*. If* $2u(15) \geq \max\{u(23)+u(11), u(25)+u(10)\}$ *then* $u(25)+u(14) < 2u(18)$*.*

Proof of Lemma 1. See appendix A. □

The implications of Lemma 1 are straightforward. If a SEU subject prefers H over L and I in the risk classification game then it follows from Proposition 1 that their set of rationalizable strategies in the testing game is $\{A\}$. Additionally, if a SEU subject prefers L over I and H in the risk classification game then it follows from Proposition 1 that their set of rationalizable strategies in the testing game is $\{A, B, C\}$. While we have assumed CARA utility functions—which have the desirable property that the proposition would still hold if we assumed that subjects integrated laboratory earnings into their lifetime wealth—in deriving Lemma 1 the first part of the lemma also holds under CRRA utility and the second part almost holds.[6,7]

The relationship between choices made in the classification games and the set of rationalizable strategies in the classification game for the own-preference treatment is summarized in Table 1.

[6] Under CRRA utility, if $2u(15) \geq \max\{u(23)+u(11), u(25)+u(10)\}$ then $u(25)+u(14) < 2u(18.04)$.

[7] I thank Simon Grant for highlighting, in private correspondence, this desirable feature of CARA preferences.

Note that subjects that choose I in the risk classification game are not included in the table; for these subjects the preference ordering between the lottery that pays \$25 and \$14 with equal probabilities and the lottery that pays \$18 with certainty is ambiguous under our assumption of CARA utility. It is not, therefore, possible to make a statement regarding the rationalizable set for these subjects with any degree of confidence, and these subjects are excluded from the analysis.[8]

Row Player Strategies	S	M
L	$\{A\}$	$\{A, B, C\}$
H	$\{A, B, C\}$	$\{A, B, C\}$

Table 1: A summary of the relationship between choices in the classification games and the set of rationalizable strategies in the own-preference treatment. A subject that plays L in the risk classification game and S in the ambiguity classification game has a unique rationalizable strategy $\{A\}$. All strategies are rationalizable for subjects that play H in the risk classification game or M in the ambiguity classification game. Subjects that select I in the risk measurement game are not classified.

There are two important takeaways from Table 1. First, the theoretical prediction hinges on whether A is the unique rationalizable strategy, or whether all strategies are rationalizable. For this reason, we will analyse the data with a focus on whether or not A is played. Second, the effects of risk aversion can only clearly be identified from the rationalizable sets by comparing subjects that chose L and S in the measurement games to subjects that chose H and S. Similarly, the effect of ambiguity aversion is identified by comparing subjects that chose L and S with subjects that chose L and M. This will be important when we identify and compare our measure of the effects of risk aversion with previous work. Finally, note that Proposition 1 immediately implies that the relevant rationalizable sets for column are $\{X\}$ whenever A is uniquely rationalizable for the row player and $\{X, Y\}$ otherwise.

2.5 Experimental conditions

All sessions were held in the ELVSE lab at the Vancouver School of Economics. The own-preference treatment consisted of 20 sessions held between March and September 2014, with between 8 and 12 subjects per session and a total of 206 subjects. The instructions for the own-preference treatment are included in the appendix. Sessions lasted between 60 and 90 minutes, and the average payment was C\$26.60. The other-preference treatment was conducted in January 2015 (92 subjects across

[8] Under an assumption of CRRA utility we could claim that the rationalizable set for these subjects is $\{A\}$. The results presented below are robust to this alternative classification, which is unsurprising given that only 10% of subjects selected I in the risk classification game.

8 sessions) and January 2016 (39 subjects across 2 sessions). The other-preference sessions were significantly shorter (because subjects faced fewer games), lasting around 45 minutes, and subjects earned an average of C$22.31. With one exception, which was removed from the data, no subject participated in more than one session. Subjects were recruited from the ELVSE implementation of the ORSEE subject pool (Greiner, 2015), which is overwhelmingly made up of UBC undergraduate students. The experiments were run using custom built software programmed in Javascript.

Recent research has identified that subject confusion can significantly lower measured ambiguity aversion (Chew et al., 2017), that experiments with ambiguity can be particularly susceptible to violations of incentive compatibility across tasks (Baillon et al., 2014; Azrieli et al., 2016) and that framing effects may be particularly strong (Chew et al., 2017). The experimental design presented here mitigates these factors by using extensive, incentivized, comprehension tasks to test for subject understanding, realizing objective randomizations prior to realizing subjective randomizations, and presenting all tasks in a unified normal-form game framing.[9] The use of classification *games* to measure preferences, rather than more traditional individual preference measuring tasks, ensures that the measurement of preferences is performed in the same domain as the testing game thereby reducing the chance of framing effects across domains contaminating the data.

3 Results

This section presents the experimental results. There are three key dimensions to the results presented:

- Risk and ambiguity preferences are positively correlated;
- There is an own-preference effect of preferences on row player behavior in the testing game, as predicted by Epstein (1997) rationalizability;
- There is an other-preference effect of opponent's preferences on column player behavior in the testing game, also as predicted by Epstein (1997) rationalizability.

We also emphasize one key implication of these results: Because risk and ambiguity preferences are correlated, and the expected effects of risk and ambiguity on behavior in games are typically

[9] See Eichberger et al. (2016), Baillon et al. (2014), Azrieli et al. (2016) and Johnson et al. (2014) for variations on the argument that objective randomizations should be realized prior to subjective randomizations. That is, the choice of game that is being used for payment should be determined and revealed to subjects prior to subjects learning the strategy choices of their opponent.

similar, previous studies that have studied only the effects of risk (or ambiguity) aversion on behavior in games have likely overestimated the parameter of interest because of an omitted variable bias effect. That is, in games where behavior appears biased towards "safe" strategies, relative to risk neutral Nash equilibrium, attributing this behavior entirely to risk aversion instead of a combination of risk and ambiguity aversion will lead to overestimates of risk aversion parameters.

3.1 Preferences

		Ambiguity preference		
		Neutral	Averse	
Risk aversion	Low ($\rho < 0.05$)	N=57 (E=49)	N=19 (E=27)	76
	High ($\rho > 0.12$)	N=18 (E=26)	N=22 (E=14)	40
		75	41	116

Table 2: Number of subjects, classified by preferences, restricted to subjects that passed the comprehension tasks. Low risk aversion subjects selected L in the risk classification game, implying a CARA parameter of $\rho < 0.05$, while high risk aversion subjects selected H, implying a CARA parameter of $\rho > 0.12$. Ambiguity neutral subjects selected S in the ambiguity classification game, while ambiguity averse subjects selected M. Expected values, assuming independence of risk and ambiguity preferences, are in brackets. The null of independence is rejected at the 1% level, using Pearson's χ^2 test ($p = 0.001$).

Table 2 presents the classification of subjects into types based on their responses as row players in the preference measuring games.[10] 35% of the subjects were classified as ambiguity averse, a figure that is at the lower end of the level of ambiguity aversion reported in previous individual decision making papers, and lower than that measured previously in 2-urn Ellsberg tasks.[11] Ivanov (2011) also found relatively lower levels of ambiguity aversion (22% of his subject population), and

[10] In the own-preference treatment we drop 69 out of 185 subjects (another 21 subjects chose I in the risk measurement game, and are dropped on that basis), while in the other-preference treatment we drop 9 out of 124 (with another 6 dropped because they faced an opponent who chose I in the risk measurement game, and 1 subject dropped because he participated in both the own-preference and other-preference treatments). In subsection B.1 we present the data including subjects that failed the comprehension tasks and find no significant differences in the results.

[11] Chew et al. (2017) provide an overview of previous Ellsberg urn experimental results in individual decision making environments. For the 2-urn case, as used in this paper, previous studies have found that between 47% and 78% of subjects are ambiguity averse (with a weighted mean of 66%). For the 1-urn (3-color) Ellsberg task, previous studies have found that between 79% and 8% of subjects are ambiguity averse (with a weighted mean of 27%).

a significant amount of ambiguity seeking, when inferring preferences from behavior in games. We find that 25% of subjects with low risk aversion are ambiguity averse, while 55% of the subjects with high risk aversion are also ambiguity averse.[12] A Pearson's χ^2 exact test rejects the null hypothesis of independence of risk and ambiguity preference at the 1% level ($p = 0.001$), and the Φ coefficient, a measure of association between two binary variables, is 0.30.[13]

Result 1. *Risk and ambiguity preferences are positively correlated.*

The previous experimental literature on the correlation between risk and ambiguity preferences has produced mixed results. Some studies have found positive correlation, others have found no correlation, and others have found correlations under some circumstances but not others, although the weight of evidence tends to favour a correlation.[14] From a theoretical viewpoint, axiomatic models of preferences (Schmeidler (1989) and Gilboa and Schmeidler (1989), for example) are typically agnostic regarding the relationship between risk and ambiguity aversion. There are, however, some non-axiomatic theoretical models of behavior under uncertainty that suggest a positive correlation between risk and ambiguity preferences (Halevy and Feltkamp (2005), for example). The evidence provided here is of a very base variety - once the framing of the games is stripped away, subjects were faced with either 2 or 3 element choice sets. Given that decisions over small choice sets form the building blocks of decision theory under uncertainty, there is a strong case to be made that decisions over small choice sets are the correct domain for examining preferences.

		Ambiguity preference		
		Neutral	Averse	
Risk aversion	Low	0.74	$0.58 (p = 0.194)$	0.70
	High	$0.61 (p = 0.307)$	$0.41 (p = 0.006)$	0.50
		0.71	0.49	

Table 3: Proportion of subjects that play A in the game in Figure 1 by preference type, restricted to subjects that passed the comprehension task. Restricted to subjects that passed comprehension. p-values shown are calculated using a χ^2 test under the null hypothesis that the proportion of subjects playing A in the relevant cell is the same as the proportion of Low Risk Aversion Ambiguity Neutral subjects playing A. $N = 116$

3.2 Own-preference results

Table 3 displays the results from the own-preference treatment. A low risk aversion ambiguity neutral subject is 1.8 times more likely to play A in the testing game than a high risk aversion ambiguity averse subject, and this difference in behavior is statistically significant at standard levels of significance ($p = 0.006$, χ^2 test). This represents clear evidence that preferences toward uncertainty, as measured in the classification games, affect behavior in the testing game.

Table 3 also highlights the inference problems that can occur in data sets that measure only one of risk or ambiguity aversion. Consider the effects of risk aversion in our testing game (Figure 1). As Table 1 makes clear, there is only a difference in the rationalizable set of strategies between low and high risk aversion subjects when the subject is also ambiguity neutral. Therefore, we can only identify the effect of risk aversion on rationalizable behavior by comparing the behavior of low risk aversion ambiguity neutral and high risk aversion ambiguity neutral subjects. Table 3 indicates

[12]Female subjects are significantly more likely to be risk averse than male subjects (Pearson's χ^2 $p = 0.025$). After conditioning on preferences, gender does not play a significant role in determining behavior. Complete demographics were not collected for two sessions because of a network failure in one session and a powerbank being knocked out of a wall in another session. The change in sample has a larger effect on the results than controlling for demographic factors does, so demographic controls are not used.

[13]The choice of statistical tests for categorical data is an oft-debated topic. Throughout this paper, every p-value could be calculated in multiple ways, with the statistical inferences almost always being the same under all alternatives. A general preference for non-parametric tests is displayed throughout, and regression (or logit) analysis is avoided in favour of χ^2 tests over appropriately defined sub-populations where possible. Fisher exact tests are used for cases where the sample sizes are small or severely unbalanced. An earlier version of the paper used Fisher exact tests throughout and generated the same conclusions.

[14]A non-exhaustive list of papers that have found a positive correlation between risk and ambiguity aversion includes: Abdellaoui et al. (2011), Bossaerts et al. (2009) and Dean and Ortoleva (2014) . A similar list on the other side of the debate includes: Curley et al. (1986) and Di Mauro and Maffioletti (2004).

that the difference in the rate of A being played between these groups is 13 percentage points ($p = 0.307$). If ambiguity preferences were not available, however, the only measured effect of risk aversion would be to compare the behavior of low risk subjects to high risk subjects aggregated across ambiguity preferences. Table 3 indicates that, using the marginal distribution of ambiguity preferences, the effect of ambiguity aversion on the rate of A being played is 20 percentage points, and is also statistically significant ($p = 0.036$).

The point estimate of the marginal effect of risk aversion is therefore approximately 1.54 times larger than the conditional effect of risk aversion. Compare this to, say, Goeree et al. (2003) which compares risk preferences as measured in various matching pennies games to risk preferences measured using a Holt and Laury (2002) risk task using a within subject design. Using maximum likelihood estimation Goeree et al. find a constant relative risk aversion parameter of 0.44 (standard error 0.07) in the games and 0.31 (0.03) in the Holt-Laury task, a difference that may be attributable to unmeasured ambiguity aversion in the games.

Result 2. *Risk and ambiguity preferences have an affect on row player behavior in the testing game.*

Result 3. *The marginal effects of both risk and ambiguity aversion in the testing game are overstated relative to the (well-identified) conditional effects.*

Of independent interest is the behavior of ambiguity neutral low risk aversion subjects, who have a unique rationalizable strategy of A in the testing game. Table 3 shows that 26% of these subjects that passed the comprehension tasks, and thereby demonstrated a thorough understanding of the payoff structure of the games, failed to play the rationalizable strategy. One possible resolution can be found in models of source-dependent ambiguity preferences: if some subjects do not consider an Ellsberg urn to be a source of ambiguity, but do consider the behavior of other people to be ambiguous, then it would be possible to rationalize a choice of S in the ambiguity classification game and something other than A in the testing game for a risk neutral subject.

3.3 Other-preference results

As in the own-preference treatment, we exclude subjects that performed poorly on the comprehension questions: 121 of the 130 subjects in the other-preference treatment answered the comprehension questions correctly on their first attempt. The higher rate of success is a function of the number of tasks performed – subjects in the other-preference treatment only played one game and therefore faced only one set of comprehension questions.[15]

The payoff structure that was used in the preference measuring games was designed for the own-preference treatment in which an unobserved variable (the subject's true preferences) was expected to influence behavior in both the preference measuring games and the testing game. As a consequence, we would expect the results in the own-preference treatment to persist independently of whether the subject understood the relationship between the preference measuring games and the testing game. In the new treatment, however, we will only observe the predicted effects when the subject understands the relationship between games.

It is, of course, unrealistic to think that subjects will have knowledge of the formal relationship between risk preferences, ambiguity preferences and solution concepts across games. There are, however, some obvious and intuitive patterns that subjects may recognize. For example, in the risk measurement game (figure 5) it is clear that, in laymen's terms, strategies L and I are "dangerous", and strategy H is "safe". In the testing game (Figure 1) it is also clear that strategies A and B are "dangerous", and strategy C is "safe". If an opponent is observed to be willing to take risks in one situation, then it seems reasonable to suppose that they might take risks in another situation; Indeed, the data suggests that subjects agree.

The data for this treatment is summarized in tables 4 and 5, with each table showing the proportion of subjects that played X as a function of their *opponent's* preferences (i.e. as a function of the choices the subject observed in the classification games). Recall that we should expect to see subjects that observe a low risk aversion and ambiguity neutral opponent should only play X, while for other subjects both X and Y are rationalizable. Subjects represented in Table 4 observed both their opponent's risk and ambiguity preferences, while subjects in Table 5 observed only their opponent's ambiguity preferences.

Remarkably, subjects that observed a low risk aversion ambiguity neutral opponent played X, their unique rationalizable strategy, 100% of the time. There is a large and statistically significant effect of opponent's risk preferences on behavior, with X being chosen only 56% of the time by subjects with a high risk aversion ambiguity neutral opponent ($p < 0.001$). In contrast, there is very little evidence of any effect of observing an opponent's ambiguity preferences on behavior in the testing game: subjects that observe a low risk aversion ambiguity averse opponent play X 92% of the time ($p = 0.32$). As in the own-preference treatment, it is inappropriate to infer the effect of preferences from the marginal distributions.

[15] For the comparable game in the main treatment 179 of 206 subjects answered the comprehension questions correctly on the first attempt. The difference (93% vs. 87%) is not statistically significant.

		Opponent ambiguity preference		
		Neutral	Averse	
Opponent risk aversion	Low	1.00	0.92($p = 0.32$)	0.97
	High	0.56($p < 0.001$)	0.30($p < 0.001$)	0.41
		0.83	0.51	

Table 4: Percentage of subjects that play X in the game in Figure 1, by opponents' preference type, restricted to subjects that passed the comprehension task, for subjects that observed both their opponents' risk and ambiguity preference. p-values calculated using a Fisher exact test under the Null that the proportion of subjects playing X in the relevant cell is equal to the proportion of subjects playing X in the top left cell. $N = 77$.

Because subjects in Table 4 observed two signals of their opponent's preferences, it is not possible to determine whether subjects view the ambiguity preference signal as completely uninformative or merely less informative than the stronger risk preference signal. Consequently, we ran additional sessions where subjects only observed their opponent's ambiguity preferences. The results are reported in Table 5. Again, subjects that observe an ambiguity neutral opponent almost always (22 out of 23 subjects) play X. subjects that observe an ambiguity averse opponent play X less than 70% of the time, and the difference in behavior is significant at standard significance levels ($p = 0.03$, one-sided Fisher exact test).

	Opponent's ambiguity preference		
	Neutral	Averse	Total
$Pr(X)$	0.96	0.67	0.84

Table 5: Percentage of subjects that play X in the game in Figure 1, by opponents' preference type, restricted to subjects that passed the comprehension task, for subjects that observed their opponents' ambiguity preference only. The difference in proportion of subjects playing X when their opponent is ambiguity neutral or averse is significant under Fisher's exact test ($p = 0.03$). $N = 38$.

It is obvious that the results for subjects with risk and ambiguity neutral opponents are exactly equal to the pure strategy Nash equilibrium of the game. A natural follow up question is: are the results for subjects with risk averse ambiguity neutral opponents also rationalizable by a Nash equilibrium?[16] The answer is yes: the proportion of subjects playing X when they observe their opponent to be ambiguity neutral but risk averse (9 out of 16 subjects, see Table 4) forms a Nash equilibrium of the game (with row players mixing between A and C) when the row player has a

[16]Because of the large set of mixed equilibrium probabilities for subjects with ambiguity averse opponents we do not repeat the analysis for those subjects.

CARA utility function with risk aversion parameter $\rho \approx 0.15$. Given that the choice of H in the risk measurement game implies $\rho > 0.12$ the results are consistent with potential equilibrium play at reasonable levels of risk aversion.

Result 4. *When a subject observes a signal of her opponent's risk and ambiguity preferences, her behavior as the column player in the game in Figure 1 is affected by the signal in the direction predicted by theories of rationalizability.*

Result 5. *When subjects observes a signal of their opponent's risk and ambiguity preferences, and their opponent exhibits ambiguity neutrality, the subjects behavior as a column player in the game in Figure 1 is consistent with Nash equilibrium.*

4 Discussion and Conclusion

This section discusses some of the implications of the data, and then concludes the paper with subsection 4.4.

4.1 Irrational behavior or incomplete information?

A typical refrain upon the observation of experimental data that does not conform to the expectations of complete and perfect information game theory is that subjects must not be behaving in a "rational" fashion. The data presented here suggests an alternative explanation – it is the blind assumption of complete information by the observer, rather than the subjects, that is not rational.

In our other-preference treatment subjects respond in a very natural, and indeed rational, way to the signal of their opponent's preferences. This is in contrast with the supplemental data shown in subsection B.2, of column players in the own-preference treatment, who did not appear to show any relationship between their inferred beliefs regarding their opponent's preferences in the measurement games and behavior in the testing game. This independence of play is easily explained by subjects being very uncertain about their opponent's preferences. If subjects have only very weak beliefs regarding their opponent's preferences, then we should expect play in the testing game to be independent of these beliefs. On the other hand, in the measurement games, the only thing that a column can do is guess their opponent's preferences which implies that even weak beliefs should determine behavior.

The uncertainty regarding others risk and ambiguity preferences implies that the assumption of complete information is unlikely to be met in many experiments. This view is reinforced by the results of Healy (2013) which finds that subjects are generally poor at estimating their opponent's preferences over outcomes.

4.2 Risk or ambiguity?

As demonstrated in subsection 3.2, it is a mistake to identify risk aversion directly from behavior in a game without appropriately accounting for ambiguity aversion. Similarly, any attempt to identify ambiguity aversion from behavior in a game should account for the possibility of risk aversion.[17]

The data presented in subsection 3.2 also shows a remarkable similarity between the effects of risk aversion and ambiguity aversion on behavior in the testing game. This similarity is suggestive— particularly given that the experiment was designed around the classic notion that risk aversion affects the valuation of *outcomes* while ambiguity aversion affects the valuation of *strategies*— of subjects that do not consider risk and ambiguity to be structurally different.[18] In this light, applications of rank-dependent expected utility models (in which risk and ambiguity can be modeled under the same conceptual framework, see Diecidue and Wakker (2001) for a straightforward comparison) to games hold particular appeal. While there is already some work in this area (see Dow and Werlang (1994), Eichberger and Kelsey (2000) and Eichberger and Kelsey (2014) for examples built using CEU), there does not yet appear to be a consensus on how to define an equilibrium, and a significant focus on risk aversion in games under rank-dependent expected utility appears to be missing from the literature.

4.3 Rationalizability or equilibrium?

While this paper has relied on rationalizability as its core solution concept, other solution concepts are of course available. Rationalizability has the advantage of being relatively light on epistemic assumptions, but generates multiplicity of predictions in our testing game. Unfortunately, the use of appropriate equilibrium concepts doesn't refine the rationalizable set in a meaningful fashion. Typical ambiguity averse equilibrium concepts generate large equilibrium sets in the testing game: all of Dow and Werlang (1994), Eichberger and Kelsey (2000) and Lo (2009), for example, allow for equilibria where A is always played by the row player (which coincides with the Nash equilibrium)

[17] Kelsey and le Roux (2015) find some unexpected correlations between ambiguity aversion measured via 3-color Ellsberg urns and ambiguity aversion measured via a game, which could potentially be resolved by accounting for unobserved risk aversion.

[18] Notwithstanding the data in the other-preference treatment, which could be attributed to subjects having a greater familiarity with risk than ambiguity.

and for equilibria where A is never played by the row player.[19]

This provides an illustration of what is, perhaps, the biggest challenge for ambiguity averse equilibrium concepts. Moving from Epstein (1997) rationalizability to an equilibrium concept requires demanding epistemic assumptions regarding conjectures and priors (see Lo (2009)), but often fails to refine the rationalizable set in a meaningful fashion. While the same critique can, of course, be leveled at Nash equilibrium relative to Pearce/Bernheim rationalizability the problem seems more pressing for ambiguity aversion in part, at least, because of the multitude of available equilibrium concepts. Furthermore, in the interesting case where ambiguity averse equilibrium generalizes Nash equilibrium, there is multiplicity of equilibria that often generate a large convex hull of feasible behavior.

4.4 Conclusion

This paper presents an experimental study into the effects of both risk and ambiguity preferences on behavior in normal form games. Using a novel methodology for measuring preferences under a game-consistent framing, the results provide fresh evidence on the interplay of risk and ambiguity in games. We find that preferences over uncertainty affect behavior both directly and indirectly. That is, both a subject's own preferences and their understanding of their opponent's preferences affect behavior. One critical implication of the results is that measurements of risk preferences, in other work, that are inferred from behavior in games are overstated because of a classical omitted variable bias caused by the unmeasured effects of ambiguity aversion.

References

Mohammed Abdellaoui, Aurelien Baillon, Laetitia Placido, and Peter P. Wakker. The Rich Domain of Uncertainty: Source Functions and Their Experimental Implementation. *American Economic Review*, 101(2):695–723, April 2011.

Yaron Azrieli, Christopher P. Chambers, and P. J. Healy. Incentives in experiments: A theoretical analysis. Working Paper, 2016.

[19]The equilibrium concepts in the first two papers are equilibrium in beliefs with no restrictions on actions, and Lo is also best interpreted via a belief interpretation of mixing, which makes the above statement a little imprecise. More precisely, there exist equilibrium where the unique support of the column player's beliefs regarding the row players behavior is $\{A\}$ and equilibrium where the unique support does not contain A.

Aurelien Baillon, Yoram Halevy, and Li Chen. Experimental elicitation of ambiguity attitude using the random incentive system. Mimeo, 2014.

P. Battigalli, S. Cerreia-Vioglio, F. Maccheroni, and M. Marinacci. A note on comparative ambiguity aversion and justifiability. *Econometrica*, 84(5):1903–1916, 2016.

B. Douglas Bernheim. Rationalizable strategic behavior. *Econometrica*, 52(4):1007–1028, 1984.

Peter L. Bossaerts, Serena Guarnaschelli, Paolo Ghirardato, and William R. Zame. Ambiguity and asset prices: An experimental perpesctive. *Review of Financial Studies*, 23:1325–1359, 2009.

Christoph Brunner, T. Florian Kauffeldt, and Hannes Rau. Does mutual knowledge of preferences lead to more equilibrium play? experimental evidence. University of Heidelberg Discussion Paper Series No. 629, 2017.

Evan M. Calford. Mixed strategies in games with ambiguity averse agents. Mimeo, 2016.

Colin F. Camerer and Risto Karjalainen. Ambiguity aversion and non-additive beliefs in non-cooperative games: Experimental evidence. In Bertrand Munier and Mark J. Machina, editors, *Models and Experiments in Risk and Rationality*. Springer Netherlands, 1994.

Soo Hong Chew, Mark Ratchford, and Jacob S. Sagi. You have to recognize ambiguity to avoid it. *The Economic Journal*, Forthcoming, 2017.

Shawn P. Curley, J. Frank Yates, and Richard A. Abrams. Psychological sources of ambiguity avoidance. *Organizational Behavior and Human Decision Processes*, 38(2):230–256, October 1986. URL http://ideas.repec.org/a/eee/jobhdp/v38y1986i2p230-256.html.

Mark Dean and Pietro Ortoleva. Is it all connected? A testing ground for unified theories of behavioral economics phenomena. Mimeo, March 2014.

Carmela Di Mauro and Anna Maffioletti. Attitudes to risk and attitudes to uncertainty: experimental evidence. *Applied Economics*, 36(4):357–372, 2004.

Enrico Diecidue and Peter P. Wakker. On the intuition of rank-dependent utility. *The Journal of Risk and Uncertainty*, 23(3):281–298, 2001.

James Dow and Sergio Riberio da Costa Werlang. Nash equilibrium under Knightian uncertainty: Breaking down backward induction. *Journal of Economic Theory*, 64:305–324, 1994.

Jurgen Eichberger and David Kelsey. Non-additive beliefs and strategic equilibria. *Games and Economic Behavior*, 30:183–215, 2000.

Jurgen Eichberger and David Kelsey. Optimism and pessimism in games. *International Economic Review*, 55(2):483–505, 2014.

Jurgen Eichberger, David Kelsey, and Burkhard C. Schipper. Granny versus game theorist: Ambiguity in experimental games. *Theory and Decision*, 64:333–362, 2008.

Jurgen Eichberger, Simon Grant, and David Kelsey. Randomization and dynamic consistency. *Economic Theory*, 62(3):547–566, 2016.

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

Larry G. Epstein. Preference, rationalizability and equilibrium. *Journal of Economic Theory*, 73: 1–29, 1997.

Larry G. Epstein and Yoram Halevy. No two experiments are identical. Mimeo, December 2014.

Itzhak Gilboa and David Schmeidler. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18:141–153, 1989.

Jacob K. Goeree, Charles A. Holt, and Thomas R. Palfrey. Risk averse behavior in generalized matching pennies games. *Games and Economic Behavior*, 45(1):97–113, 2003.

Ben Greiner. Subject pool recruitment procedures: Organizing experiments with orsee. *Journal of the Economic Science Association*, 1(1):114–125, July 2015.

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

Yoram Halevy and Vincent Feltkamp. A Bayesian approach to uncertainty aversion. *Review of Economic Studies*, 72:449–466, 2005.

P. J. Healy. Epistemic foundations for the failure of Nash Equilibrium. Mimeo, 2013.

Frank Heinemannn, Rosemarie Nagel, and Peter Ockenfels. Measuring strategic uncertainty in coordination games. *Review of Economic Studies*, 76:181–221, 2009.

Charles A. Holt and Susan K. Laury. Risk aversion and incentive effects. *American Economic Review*, 92(5):1644–1655, 2002.

Asen Ivanov. Attitudes to ambiguity in one-shot normal-form games: An experimental study. *Games and Economic Behavior*, 71:366–394, 2011.

Cathleen Johnson, Aurelien Baillon, Han Bleichrodt, Zhihua Li, Dennie van Dolder, and Peter P. Wakker. Prince: An improved method for measuring incentivized preferences. Mimeo, 2014.

David Kelsey and Sara le Roux. An experimental study on the effect of ambiguity in a coordination game. *Theory and Decision*, 79(4):667–688, 2015.

John M. Keynes. *A Treatise on Probability*. MacMillan & Co., 1921.

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

Frank H. Knight. *Risk, Uncertainty and Profit*. Hart, Schaffner & Marx; Boston: Houghton Mifflin Company, The Riverside Press, Cambridge, 1921.

Kin Chung Lo. Correlated Nash Equilibrium. *Journal of Economic Theory*, 144:722–743, 2009.

David G. Pearce. Rationalizable strategic behavior and the problem of perfection. *Econometrica*, 52(4):1029–1050, 1984.

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

A Proofs

The proof of Proposition 1 uses the method of iterated elimination of strategies that are never a best response. This method produces the set of Pearce/Bernheim rationalizable strategies under SEU preferences, and produces the set of Epstein (1997) rationalizable strategies under MEU preferences.

Proof of proposition 1. Suppose that the row player has preferences such that $u(25) + u(14) > 2u(18)$ and Φ_R is restricted to be a singleton. Then the best response for the row player is A when $\Phi_R(X) > \frac{1}{2}$, is B when $\Phi_R(X) < \frac{1}{2}$ and is either A or B when $\Phi_R(X) = \frac{1}{2}$. C is never a best response.

Using the procedure of iterated elimination of never best response strategies, we eliminate C. Now, in the reduced game, Y is never a best response for the column player. Eliminate Y. Now B is never a best response.

Therefore $\{A, X\}$ is the unique rationalizable strategy.

Now, consider the case where $u(25) + u(14) \leq 2u(18)$. C is now a best response to the belief set $\Phi(X) = \{\phi : \phi(X) = \frac{1}{2}\}$. A and B are both clearly best responses for some beliefs, as are X and Y. Therefore all strategies are rationalizable.

Now, consider the case where $\Phi(X)$ may be set valued. C is now a best response to the belief set $\Phi(X) = \{\phi : 0 \leq \phi(X) \leq 1\}$. A and B are both clearly best responses for some beliefs, as are X and Y. Therefore all strategies are rationalizable. □

Proof of Lemma 1. Under the assumption of CARA the agent has a utility function of the form $u(c) = 1 - e^{-\rho c}$.

For the first part, $u(25) + u(10) \geq \max\{u(23) + u(11), 2u(15)\}$ is equivalent to both

$$1 - e^{-25\rho} + 1 - e^{-10\rho} \geq 1 - e^{-23\rho} + 1 - e^{-11\rho}$$

and

$$1 - e^{-25\rho} + 1 - e^{-10\rho} \geq 2(1 - e^{-15\rho})$$

and both inequalities are satisfied whenever $\rho \leq 0.05$ (correct to 2 decimal places).

$u(25) + u(14) > 2u(18)$ is equivalent to

$$1 - e^{-25\rho} + 1 - e^{-14\rho} \geq 2(1 - e^{-18\rho})$$

which holds whenever $\rho < 0.10$.

For the second part, $2u(15) \geq \max\{u(23) + u(11), u(25) + u(10)\}$ is equivalent to both

$$2(1 - e^{-15\rho}) \geq 1 - e^{-23\rho} + 1 - e^{-11\rho}$$

and

$$2(1 - e^{-15\rho}) \geq 1 - e^{-25\rho} + 1 - e^{-10\rho}$$

and both inequalities are satisfied whenever $\rho \geq 0.12$. $2u(18) > u(25) + u(14)$ is satisfied whenever $\rho > 0.10$.

B Additional Results

B.1 Row player results

This section contains analogues of tables 2 and 3 that include subjects that failed to pass the comprehension tasks.

		Ambiguity preference		
		Neutral	Averse	
Risk aversion	Low	Type 1 N=85 (E=75)	Type 2 N=36 (E=46)	121
	High	Type 3 N=30 (E=40)	Type 4 N=34 (E=24)	64
		115	70	185

Table 6: Number of subjects, classified by type. Expected values, assuming independence of risk and ambiguity preferences, are in brackets. The null of independence is rejected at the 1% level, using Pearson's χ^2 test ($p = 0.002$).

B.2 Column player results

This section presents the column player results from the original own-preference treatment.

The results in this section can be interpreted in a similar fashion to the results from the other-preference treatment, with one key distinction: instead of subjects observing a signal of their opponent's preferences, we instead observe the beliefs of column players regarding their opponent's

		Ambiguity preference		
		Neutral	Averse	
Risk aversion	Low	0.67	$0.56(p = 0.229)$	0.64
	High	$0.60(p = 0.485)$	$0.32(p = 0.001)$	0.45
		0.65	0.44	

Table 7: Proportion of subjects that play A in the game in Figure 1 by preference type. p-values shown are calculated using a χ^2 test under the null hypothesis that the proportion of subjects playing A in the relevant cell is the same as the proportion of Low Risk Aversion Ambiguity Neutral subjects playing A. $N = 185$

		Opponent ambiguity preference		
		Neutral	Averse	
Opponent risk aversion	Low	1.00	$0.93(p = 0.33)$	0.98
	High	$0.63(p = 0.001)$	$0.33(p < 0.001)$	0.47
		0.85	0.55	

Table 8: Percentage of subjects that play X in the game in Figure 1, by opponents' preference type for subjects that observed both their opponents' risk and ambiguity preference. p-values calculated using a Fisher exact test under the Null that the proportion of subjects playing X in the relevant cell is equal to the proportion of subjects playing X in the top left cell. $N = 85$.

	Opponent's ambiguity preference		
	Neutral	Averse	Total
$Pr(X)$	0.96	0.69	0.85

Table 9: Percentage of subjects that play X in the game in Figure 1, by opponents' preference type, restricted to subjects that passed the comprehension task, for subjects that observed their opponents' ambiguity preference only. The difference in proportion of subjects playing X when their opponent is ambiguity neutral or averse is significant under Fisher's exact test ($p = 0.03$). $N = 39$.

preferences. We find only null results. In light of the strong results from the other-preference treatment presented in the main text, it appears that subjects hold only very weak beliefs regarding their opponent's preferences.

While the role of the column player was introduced into the preference measuring games to provide a consistent frame across tasks, the column player behavior still contains information. In those games, the column player earns a positive payoff if and only if they correctly predict their opponents behavior. In other words, the column player is incentivized to try and guess their opponent's behavior. For example, if a column player chooses L' in the game in Figure 5 we interpret this as implying that the column player believes that their opponent is most likely to have low risk aversion (consistent with a CARA parameter $\rho \leq 0.05$). Similarly, if a column player chooses M' in the game in Figure 2 we interpret this as implying that the column player believes that their opponent is most likely to be ambiguity averse.

Following the same reasoning as in the other-preference treatment in the main text, we should expect subjects that believe their opponent to be highly risk averse or ambiguity averse to be more likely to play Y in the testing game of Figure 1. Table 10 outlines the data, which shows no meaningful relationship between beliefs (as implied by column player behavior in the classification games) and column player behavior in the testing game.

Belief of opponent's		Ambiguity preference		
		Neutral	Averse	
Risk aversion	Low	0.74	0.83($p = 0.488$)	0.77
	High	0.65($p = 0.431$)	0.59($p = 0.226$)	0.62
		0.72	0.71	

Table 10: Proportion of subjects that play X in the game in Figure 1 by beliefs over opponent's preferences. Restricted to subjects that passed the comprehension tasks. p-values shown are calculated using a χ^2 test under the null hypothesis that the proportion of subjects playing X in the relevant cell is the same as the proportion of Low Risk Aversion Ambiguity Neutral subjects playing A. $N = 102$

C Online Appendix: Instructions

The instructions presented to subjects in the September sessions subjects are reproduced below. The instructions in the March and April sessions had a different example game presented as "Game X". Otherwise, the two sets of instructions were identical. The instructions were originally written in HTML, CSS and Javascript and were presented to students on their computer in an Internet browser. The instructions were interactive, and subjects could highlight strategies in the example games and practice filling in the drop down menus for "Game X". The instructions presented below have been reformatted to enhance readability in the pdf format, but are materially the same as those presented to the subjects.

C.1 Instructions

This is a research experiment designed to understand how people make economic decisions. To assist with our research, we would greatly appreciate your full attention during the experiment. Please do not communicate with other participants in any way and please raise your hand if you have a question.

You will participate in a series of 7 games. In each game you will make one decision, and a subject you will be paired with, called your counterpart, will make one decision. **One, and only one** of the games will be chosen, in a random fashion described below, as the game for which you will be paid. The amount that you earn will depend on a combination of your decision, your counterpart's decision and, for some games, the colour of a ball drawn from a bag. Your counterpart will be one of the other participants in this room and the identity of your counterpart has been randomly pre-determined.

C.1.1 How your payment will be determined

As mentioned above, only one game will be chosen as the game for which you will be paid. Each of the 7 games is equally likely to be chosen. You will notice that there are 7 pieces of paper labelled from 'A' to 'G' stuck to the wall. On the back of each of them is a number. At this point, I would like to ask one of you to choose one of the pieces of paper labelled from 'A' to 'G'. The number that is written on the back of the chosen piece of paper will determine which of the 7 games is chosen for payment. Although the choice of game has already been fixed by the choice of letter, we will not reveal the choice to you until the end of the experiment. This procedure suggests that you should treat each game independently (i.e. as if it was the only game in the experiment). After the experiment you may come and check that each of the games actually is represented on the back of one of the cards.

Your payment may also depend on the colour of a ball drawn from a bag. In some games there will no balls drawn, in some other games there will be only one ball drawn from a bag, and in some games there will be two balls drawn (each from a different bag). The 'Unknown' bag (denoted by the letter U) will contain only RED and YELLOW balls in an unknown ratio, while the 'Known' (denoted by the letter K) bag will contain RED and YELLOW balls in an equal ratio. The total number of balls in each bag will always be 10. A graphical representation of the two bags is shown below.

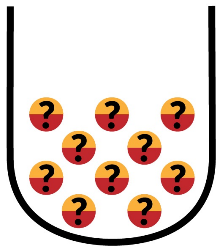

Figure 7: U bag.

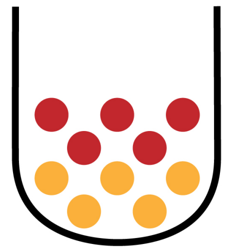

Figure 8: K bag.

When you are playing the games there will be a picture of the relevant bag(s) on the screen to remind you of the composition of the bag(s).

Before the experiment begun, I asked a graduate student in economics, who has no knowledge of this experiment, to place 10 balls into the U bag. I instructed the student to place exactly 10 balls in the bag, and that only RED and YELLOW balls are to be placed in the bag. I have no knowledge of the composition of the balls in the bag, other than what is described above. There might be 10 RED balls and 0 YELLOW balls, or 9 RED balls and 1 YELLOW ball, or 8 RED balls and 2 YELLOW balls, and so on up to 0 RED balls and 10 YELLOW balls. After the experiment you may come and check that the bag satisfies the requirements outlined above.

I shall now create the K bag, by placing 5 RED balls and 5 YELLOW balls into a second bag.

At the end of the experiment, after the chosen game has been revealed, I will ask one of you to draw a ball from the K bag (if required), and another one of you to draw a ball from the U bag (if required).

The combination of the game chosen, your choices, your counterpart's choices and the colour of the ball(s) will determine how much money you make during the experiment. The amount that you make during the experiment will be added to your $5 show up fee and will be paid to you, in cash, at the end of the experiment.

One of the conditions of the ethics approval for this experiment is that I do not deceive the subjects (i.e. you). If you feel that I have deceived you in any way, you may contact either my thesis supervisor or the UBC Behavioural Ethics Review Board to lodge a complaint. Their contact details are included on the consent form that you have read and signed.

C.1.2 The structure of the games

In each game, you will need to make a choice of either A or B or C (in some games you will only have 2 options, A and B). You may only ever choose one option per game. Your counterpart will make a choice of either X or Y or Z (in some games your counterpart will only have 2 options, X and Y). Your counterpart may also only ever choose one option per game. The amounts that each of you can earn will be presented to you in a table format, as seen below.

	X	Y	Z
A	20,10	8,0	0,0
B	30,0	6,10	4,0
C	30,0	15,0	9,10

Your options will always be shown on the rows of the table. Your counterpart's options will always be shown on the columns of the table (the computer flips the game so that everyone can look at the table from the same perspective). Within each cell of the table your earnings will always be shown first, ¡b¿in bold¡/b¿, and your counterpart's earnings will always be shown second. Values shown are always shown dollar amounts. For example, in the above game, if you chose action 'A' and your counterpart chose action 'Y' then you would receive a $8 and your counterpart would receive $0. On the other hand, if you chose 'A' and your counterpart chose 'Z' then you would each earn $0. Games that have only a single payoff table, like the above example, do not require a ball to be drawn.

In other games, your earnings will also depend on the colour of a ball which will be drawn from a bag **after you have both made your decisions**. These games will have two tables, and the colour of the ball(s) will determine which table is used to calculate your earnings. In the example given below, the left table will be used if a RED ball is drawn and the right table will be used if a YELLOW ball is drawn.

	L'	I'
L	$12, 10$	$18, 0$
I	$17, 0$	$9, 10$

Red ball drawn from K bag

	L'	I'
L	$12, 10$	$19, 0$
I	$8, 0$	$4, 10$

Yellow ball drawn from K bag

Suppose, in the above example, that you chose option 'B' and your counterpart chose option 'Y'. Then your payment will be $9 if the ball drawn is RED, and $4 if the ball drawn is YELLOW. Your counterpart's payment will be $10 if the ball drawn is RED, and $10 if the ball drawn is

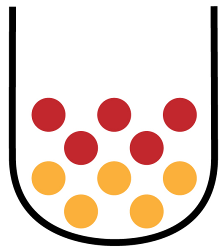

Figure 9: K bag.

YELLOW.

The game shown above was a K game. In a K game there is only one ball drawn, from the K bag, and the colour of the ball will determine which table will be used. There will also be U games, where there is also only one ball drawn, from the U bag. The other type of game, a U and K game, is shown below. It will always be clear which type of game you are playing from the labels on the tables and the pictures of the bags underneath the game.

	L'	I'
L	$12, 10$	$18, 0$
I	$17, 0$	$9, 10$

Red ball drawn from U bag

	L'	I'
L	$12, 10$	$19, 0$
I	$8, 0$	$4, 10$

Red ball drawn from K bag

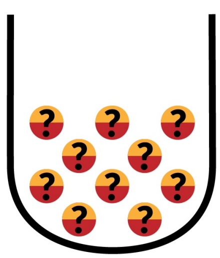

Figure 10: U urn.

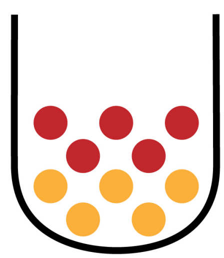

Figure 11: K bag.

In a U and K game there will be two balls drawn - one from each of the bags. If a RED ball is drawn from the U bag then you will be paid according to the left hand table. If a RED ball is drawn from the K bag then you will be paid according to the right hand table. It is also possible that a RED ball will be drawn from both bags. In this case, you would be paid according to both tables (the payments will be added together). However, it is also possible that a RED ball will not be drawn from either bag. In this case you would receive no payment (other than the $5 show up fee and any bonus payments you may earn).

C.1.3 Bonus payments and drop down menus

Before you can confirm your choices in the game itself, you will need to fill in a series of dynamic drop down menus to confirm that you have understood the game. You should pay careful attention to the drop down menus, becuase you will earn bonus payments that depend on whether you have filled the drop down menus in correctly.

For each game, you should fill in the drop down menus first. Once you are happy that your choices adequately describe the game, you should click on the "Check Description" button. If you fill in the drop down menus correctly on your first attempt you will earn a $1 bonus for that game. As there are 7 games, you may earn up to $7 in bonus payments. If you click "Check Description", but have filled in the dropdown menus incorrectly or have not filled the dropdown menus in at all, your bonus payment will decrease by $0.25. After four incorrect attempts your bonus payment *for only that game* will reach zero.

WARNING: The bonus payment system relies on 'alerts' that will pop up on your screen. Sometimes the Chrome browser will give you option of turning the alerts off. Do not do this; if you do then the computer may not record your bonus payments. If you accidentally turn the alerts off then please raise your hand and we will reset your browser.

Before we continue shall work through an example of the drop down menus together.

Note that most of the games appear in pairs - each game has two players, and you will play each game in each role. (Recall that there are seven games - the seventh game does not have a pair). Below is an example of Game X (with the drop down menus removed), viewed from the other role. Notice that while in Game X your earnings were not affected by your counterpart's choices, in Game XT your counterpart's earnings are not affected by your choices.

C.1.4 How to use the game interface

Once your description of the game is correct (and you have verified this by clicking the "check description" button) the drop down menus will inactivate, and the "lock" button will activate. At this point you can enter your choice by clicking on the desired option in the table. When you click on a choice, the computer will highlight the row that you have chosen (you may try this in Game X above). If you want to change your mind, you can simply click on a different choice. Once you are happy with your choice you should click on the "lock" button before moving on to the next

Counterpart's Choices

Your Choices		X	Y	Z
	A	9, 10	9, 0	9, 0
	B	4, 0	4, 10	4, 0
	C	12, 0	12, 0	12, 10

RED ball drawn from K bag

Counterpart's Choices

Your Choices		X	Y	Z
	A	9, 10	9, 0	9, 0
	B	15, 0	15, 10	15, 0
	C	6, 0	6, 0	6, 10

YELLOW ball drawn from K bag

Lock Unlock

Choice A: Your earnings for this choice [ ] affected by your counterparts' strategy.

Choice B: Your earnings for this choice [ ] affected by your counterparts' strategy.

Choice C: Your earnings for this choice [ ] affected by your counterparts' strategy.

Check Description

You may type whatever you want here. This textbox may be enlarged by dragging the lower right corner.

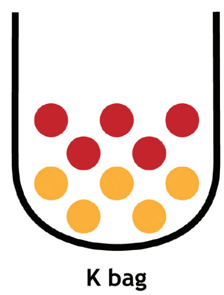

Figure 12: Game X.

game. If, later, you wish to change your mind you can always click on the "unlock" button and then change your choice.

You may also highlight any of your counterpart's options by clicking on the label for that choice (as long as the game is unlocked) - this feature is provided to you as a visual aid to assist in your decision making process, but will have no impact on the earnings received by either you or your counterpart. There is another decision making aid that has been provided to you: a textbox. In the past, some subjects have found it useful to write down the reasoning behind their decisions. You do not have to write anything in the textbox, and nothing you do write in the textbox will be

	X	Y	Z
A	$10, 9$	$0, 4$	$0, 12$
B	$0, 9$	$10, 4$	$0, 12$
C	$0, 9$	$0, 4$	$10, 12$

Red ball drawn from K bag

	X	Y	Z
A	$10, 9$	$0, 15$	$0, 6$
B	$0, 9$	$10, 15$	$0, 6$
C	$0, 9$	$0, 15$	$10, 6$

Yellow ball drawn from K bag

Figure 13: Game XT.

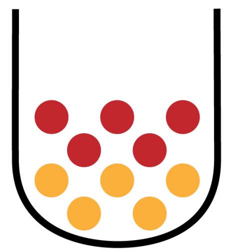

Figure 14: K bag.

used in the determination of your earnings. Nothing you write will be shown to your counterpart. Likewise, nothing your counterpart writes will be shown to you. When you lock a game, the textbox associated with that game will lock as well.

Once you have entered your choices for each of the 7 games, and locked each of the 7 games, you can click the button to submit all of your answers. Once you have pressed the "submit" button you can no longer go back and change your answers. Once everyone has clicked submit, the computer will match your responses with your counterpart's choices. The computer will not match you with your counterpart until everyone has finished the experiment, so there is no advantage to rushing. Take your time and make sure you are happy with your choices.

C.1.5 Summary of the order of events

- Read instructions, and choose a letter on the wall that will determine which game will be paid.
- Play the 7 games. For each game, begin by filling in the drop down menus, and then make your choice in the game. Lock each game after you have made your choice. If you want to change your mind, you can always unlock any game and change your choice.
- Once you are happy with all of your choices, press the 'submit' button. Once you have pressed the 'submit' button your choices are final.
- Enter some demographic information into the computer.
- Reveal the game that will be paid.
- Draw ball(s) from the bag(s).
- Calculate your payment, based on the chosen game, colour of ball, your action and your counterpart's action.

- Receive payment and leave the experiment.