Essays on Strategic Uncertainty with non-Subjective Expected Utility Agents

by

Evan M. Calford

B.Com(Hons), The University of New South Wales, 2008

B.Sci, The University of New South Wales, 2008

A THESIS SUBMITTED IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF

Doctor of Philosophy

in

THE FACULTY OF GRADUATE AND POSTDOCTORAL STUDIES

(Economics)

The University of British Columbia

(Vancouver)

July 2016

Abstract

This thesis contains three distinct chapters that contribute to our understanding of how people respond, both theoretically and in controlled experimental environments, to uncertainty that results from the strategic decisions of others. The standard framework for studying strategic interactions involves agents with Subjective Expected Utility preferences (Savage, 1954) interacting in an environment where, in equilibrium, all strategies are known to all agents. This thesis studies the effects of relaxing preferences to allow for ambiguity aversion, regret minimization, and approximate optimization.

The first chapter 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 chapter introduces a novel methodology for identifying the effects of risk and ambiguity preferences on behaviour 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 second chapter studies, experimentally, a simple dynamic entry game in both continuous and discrete time. We introduce new laboratory methods that allow us to eliminate natural inertia in subjects' decisions in continuous time experiments. Using our novel continuous time setting and the standard discrete time setting as benchmarks, we study the effects of inertia (caused by naturally occurring reaction lags) on behaviour. We demonstrate that the observed patterns of behaviour are consistent with standard models of decision making under uncertainty,

and that the degree of inertia affects subject responses to strategic uncertainty.

The third chapter examines, theoretically, the role of mixed strategies for agents with ambiguity averse preferences. This chapter demonstrates how a well known result from cooperative game theory, that a non-additive measure over a set of states can be equivalently represented by an additive measure over the set of events, can be used to introduce mixed strategies (in an equilibrium preserving fashion) to existing pure strategy equilibrium concepts.

Preface

Chapters 2 and 4 are the unpublished, original and independent work of the author, E. Calford. Chapter 3 is the unpublished and original work of E. Calford and R. Oprea.

Chapters 2 and 3 include experimental data from human subjects that was collected at the Experimental Lab at Vancouver School of Economics (ELVSE). Chapter 2, excluding Table 2.12, is covered by UBC Ethics Certificate number H13-02107. Table 2.12 of Chapter 2 is covered by UBC Ethics Certificate number H15-03133. Chapter 3 is is covered by UBC Ethics Certificate number H14-00449.

Table of Contents

List of Tables

List of Figures

Acknowledgments

This dissertation would not have been possible without the fantastic support of many people. Foremost, I would like to thank my supervisor, Yoram Halevy, for all of his support, guidance and encouragement. Yoram introduced me to decision theory, and his careful style of analytical thinking has improved this dissertation in numerous ways. I also thank Ryan Oprea, who complements Yoram perfectly. Ryan has been extremely generous with his time, energy and even his home (where Chapter 3 was born). If this dissertation reflects even a fraction of both Yoram's careful rigor and Ryan's artful storytelling then it must surely be a success. I would also like to thank my other committee members, Mike Peters and Wei Li, who have challenged and improved my thinking and writing greatly over the years. Thank you also to Li Hao and Vitor Farinha Luz for helpful discussions.

Thank you also to my family: my parents, who perhaps unwittingly set me on this path by teaching me the value of both education and incentives when they instituted pay-for-performance in high school, have provided me with unending and unquestioning emotional and financial support throughout my entire journey; my sister and brother, who motivate me both with their direct support and the drive that can only arise from having two fantastically successful younger siblings chasing after you. And, finally, my deepest gratitude must go to Kiri, who has supported me wholeheartedly throughout my entire PhD. Without her support and love I would likely never have reached this point and it is only fitting that she can now share in the rewards, just as she has shared in the journey.

To the ever-loving Kiri.

Chapter 1

Introduction

Strategic uncertainty is a fundamental property of any strategic interaction between agents. Different agents, in different situations, will respond to strategic uncertainty in different fashions. An agent with subjective expected utility preferences will always behave as if they can form a precise belief regarding the behaviour of others. Other agents may be unable or unwilling to form such beliefs, and these agents will necessarily have non-subjective expected utility preferences. This dissertation studies such agents using both theoretical and experimental methodologies.

Chapter 2 is an experimental study of the effects of uncertainty aversion in normal form games. The experimental design is straightforward: first risk and ambiguity preferences, as well as beliefs regarding others' preferences, are elicited via a series of "classification" games. Second, subjects play a normal form "testing" game. This structure allows us to directly observe the mapping between preferences, and beliefs over others' preferences, to behaviour in a strategic setting. The payoffs of the testing game are designed so that we can isolate the effects of preferences from the effects of beliefs over preferences, thereby teasing apart two key sources of strategic uncertainty.

The original contributions of Chapter 2 are threefold. First, we develop and implement a procedure for measuring subject preferences and beliefs over others' preferences in normal form games. Second, the chapter is the first paper to isolate the effects of ambiguity from the effects of risk in a normal form game. Third, the chapter is the first to study the role of beliefs over others' preferences for uncertainty in a strategic environment.

Chapter 3 is a joint work with Ryan Oprea. This chapter studies the effects of various timing protocols on behaviour in a simple dynamic entry game. The chapter is made possible by an innovative experimental design which allows us to implement the theoretical premises of a continuous time game directly in the laboratory. Previous experimental studies of continuous time games have used settings where subjects' natural reaction lags drive a wedge between the experimental setting and the theory of continuous time games (in which agents can respond instantaneously to changes in their environment). Our innovation is to introduce what we call the "freeze time protocol", whereby the game clock is frozen at the moment of entry. During the freeze other subjects can formulate their response to the entry, and implement it, allowing for truly instantaneous responses as measured by the game clock and satisfying key tenets of the theory.

Using our freeze time protocol as well as the standard discrete time protocol as benchmarks, we study the effects of reaction lags on behaviour in our entry game. The existence of reaction lags generates inertia in subject actions, and we manipulate the severity of inertia by speeding up and slowing down the game clock across treatments. We propose ex-ante hypotheses motivated by ε-equilibrium in continuous time games (Simon and Stinchcombe, 1989; Bergin and MacLeod, 1993) which are strongly borne out by the data. In our setting ε-equilibrium leads to a large multiplicity of equilibria in most treatments, and this multiplicity motivated an ex-post exploration of theories that were able to explain the data with greater precision. The best descriptor of our data is a theory of strategic uncertainty known as minimax regret aversion, which posits that when faced with an uncertain environment a decision maker should seek to choose that strategy that minimizes their maximal potential deviation from optimality. We find that changes in the degree of inertia change the payoff effects of strategic uncertainty and that minimax regret aversion captures subject responses to changing strategic uncertainty across treatments very well.

Chapter 4 is a contribution to the theoretical literature of game theory with ambiguity agents. In games with standard, subjective expected utility, agents there are multiple interchangeable interpretations of mixed strategies. First, we may interpret agents as submitting explicitly mixed strategies to the game maker. Second, we may restrict agents to pure strategies only, and then interpret a mixed strategy equilibrium as an equilibrium in beliefs. Third, agents may use a pre-play randomization device before reporting a pure strategy to the game maker.

For games with ambiguity averse agents the different interpretations do not generate observationally equivalent equilibrium because the agents may have a strict preference for mixing. It has been demonstrated that the first and second interpretations lead to different equilibrium: compare Lo (1996) to Dow and Werlang (1994), for example. It is not, however, known how the third interpretation of mixed strategies relates to the first two. The third interpretation is particularly important for experiments with ambiguity averse agents, such as the one presented in Chapter 2, given that subjects must enter a pure strategy into the experimental interface but are free to randomize across their strategies before selecting a pure strategy. Chapter 4 fills this gap in the literature by establishing conditions under which the third interpretation of mixed strategies is equivalent to the second interpretation.

1.1 Relationship to the Global Games Literature

The most common methodology for modeling strategic uncertainty is derived from the literature on global games (Carlsson and van Damme, 1993; Morris and Shin, 2003). The nature of strategic uncertainty considered in this dissertation is distinct from the strategic uncertainty considered in the global games literature.

In this dissertation, strategic uncertainty arises from agents not having a precise belief regarding the strategic behaviour of their opponent. The strategic uncertainty is first order strategic uncertainty (although higher order beliefs may also be important), it is Knightian uncertainty, and the uncertainty may exist even in games of complete information. A necessary condition for such strategic uncertainty to affect behaviour is that the agents have non-subjective expected utility preferences. The area of study in this thesis is, therefore, strategic uncertainty in games of complete information when agents have non-subjective expected utility preferences.

In the global games literature, strategic uncertainty arises from a relaxing of complete information while maintaining subjective expected utility preferences.

The key insight of the global games literature is that when complete information is relaxed then uncertainty is generated across the entire hierarchy of beliefs, and beliefs about beliefs, and so on. In the limiting case as first order uncertainty disappears, the effect of strategic uncertainty over the entire belief hierarchy can still influence behaviour. This is, however, a very different mechanism to the one studied in this dissertation.

Each approach has some advantages and disadvantages and it is important that we understand the effect of each mechanism on behaviour. It is well established that many people do not have subjective expected utility preferences (Table 2.4 confirms this to be true), and it is clearly also the case that many strategic interactions involve incomplete information. By isolating, and studying, each cause of strategic uncertainty we can generate a complete picture of the role of strategic uncertainty in human decision making.

Chapter 2

Uncertainty Aversion in Game Theory: Experimental Evidence

Strategic interactions are a source of subjective uncertainty for agents. A rational agent must, necessarily, form subjective beliefs regarding their opponent's behaviour. But what form do these beliefs take? The standard approach to modelling strategic interactions, Nash equilibrium, resides on a bed of expected utility theory: in an equilibrium each agent has a consistent and precise belief over other's behaviour, and the only remaining source of uncertainty is risk that stems from the possible use of mixed strategies. This approach, however, fails to consider that agents may retain some uncertainty regarding their opponent's choice of (mixed) strategy. We shall refer to this uncertainty over mixed strategies as strategic uncertainty. Strategic uncertainty is the natural condition of real-world 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 provide 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.

There is, however, a fundamental problem with direct inference of ambiguity aversion from behaviour in games: risk aversion. How does a subject's ambiguity preference affect their behaviour in the presence of risk aversion? In 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 affects are much more difficult to disentangle, and this chapter is the first to tackle this separation in a game theoretic setting. Because risk and ambiguity aversion have similar effects in games (making 'safe' strategies appear relatively more attractive), and are positively correlated[1], studies that focus only on risk aversion *or* ambiguity aversion in games will be prone to omitted variable bias.

To disentangle these effects we use a laboratory experiment to measure ambiguity and risk aversion in games, at the individual level, and study how behaviour in a carefully chosen normal form "testing" game is related to these measures. The testing game is designed so that the set of rationalizable strategies varies with preferences, allowing for a partial separation of behaviour 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. In a follow up treatment we allow a group of subjects to observe their opponents' responses to the preference measuring tasks, allowing us to study how subjects reason conditional on their opponents' preferences.

Subjects play 3 distinct games, and play each game as both the row player and the column player. The first two games, the 'classification' games, are used to measure ambiguity and risk preferences, respectively. The third game, the 'test-

[1]Although there is some debate on this topic, risk and ambiguity aversion are correlated in the data presented in this chapter. See footnote 22 for more details.

ing' game, is a 3×2 normal form game used to investigate the role of preferences on strategic behaviour and strategic reasoning. Across the three games, the behaviour of row players is used to study the effect of preferences on behaviour in normal form games. In comparison, the behaviour of column players across the three games is used to study the effect of beliefs over opponents' preferences on behaviour in normal form games. Details on the structure of the games are provided in section 2.2. Importantly, preferences are measured using games, rather than individual decision tasks, which allows the entire experiment to be conducted in a game environment.

Consider the testing game used in this experiment, presented in figure 2.1. For a risk neutral row player with subjective expected utility (SEU) preferences, C is never a best response. Climbing the chain of rationalizability[2], once C is eliminated then Y is never a best response for the column player and once Y is eliminated B is never a best response for the row player. The unique rationalizable outcome of the game is (A,X). Naturally, (A,X) is also the unique Nash equilibrium of the game.

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

Figure 2.1: Testing game. Payoffs are in Canadian Dollars.

In contrast, consider an agent who is ambiguity averse, with the maxmin expected utility (MEU) preferences of Gilboa and Schmeidler (1989). For the row player, C is now a best response for at least some feasible beliefs.[3] To see this, suppose that the row player faces complete uncertainty regarding her opponent's strategy. Then, applying MEU preferences, she evaluates strategy A by considering the worst possible scenario (her opponent playing Y) and values the strategy such

[2]Recall that the rationalizable set can be found via iterated elimination of never-best-response strategies or, equivalently for two-player games, iterated elimination of strictly dominated strategies. An implication of this is that if no strategies are strictly dominated then all strategies are rationalizable.

[3]Under MEU preferences an agent has a set of beliefs, and evaluates the utility of a prospect with respect to the worst possible belief in the set. MEU preferences are not critical here, any of the standard models of ambiguity aversion could be used with only minor modifications throughout the chapter.

that $U(A) = 14$. Similarly, when evaluating B she considers the worst possible scenario (her opponent playing X) so that $U(B) = 14$ as well. However, $U(C) = 18$ and therefore C is rationalizable for the row player. Obviously A and B are also potential best responses (to the belief the column player is using pure X or pure Y, respectively), and both column player strategies are also potential best responses. Therefore, under MEU preferences the rationalizable set is the full strategy set for both players.

Finally, note that if the row player has SEU preferences with a high enough level of risk aversion then the rationalizable set is also the complete set of outcomes. For a subject with SEU preferences, C is never a best response to any beliefs for the row player if and only if C is not a best response to an opponent who mixes 50-50 over X and Y. When C is a best response to the 50-50 mix then all strategies for both players are best responses to some beliefs and, therefore, all strategies are rationalizable.

The size of the rationalizable set hinges critically on the row player's preferences, but not the column player's preferences: the rationalizable set is (A, X) if and only if the row player can eliminate C in the first round. This provides the separation of the role of preferences and beliefs, which is one of the key innovations of this chapter. If the row player has SEU preferences and low risk aversion then she should never play C *regardless of her beliefs regarding her opponent's preferences*. The column player should never play Y if he believes that his opponent has SEU preferences and low risk aversion *regardless of his own preferences*. The separation of the role of preferences and beliefs is a novel design feature that allows for an investigation of the structural underpinnings of ambiguity averse solution concepts.

Preferences organize the data: in the game in figure 2.1 ambiguity neutral subjects with low risk aversion play A more than twice as often as ambiguity averse subjects with high risk aversion. On the other hand, measured beliefs over opponent's preferences are independent of column player behaviour: amongst subjects who passed a series of incentivized comprehension tasks, there is no difference in behaviour between subjects who believe their opponent to be ambiguity neutral or ambiguity averse: both groups play X approximately three-quarters of the time. A follow-up treatment, where subjects were shown their opponents' behaviour in the classification games before playing the testing game, establishes that this null result is not driven by an inability of subjects to use preference information to predict opponent behaviour. A subject that observes his opponent choosing options consistent with ambiguity and risk neutrality in the preference measuring games chooses X 100% of the time, while a subject who observes his opponent choosing options consistent with ambiguity and risk aversion in the preference measuring games chooses X 30% of the time.

Taken as a whole, our findings provide support for the claim that ambiguity preferences are an important determinant of behaviour in games. However, the data also suggest that mutual knowledge of preferences is not satisfied in the subject population and an intervention that provides credible preference information to subjects induces large behavioural changes. Nash equilibrium, therefore, is rejected on two accounts. First, Nash fails to take account of ambiguity preferences. Second, subjects do not appear to hold a consistent set of beliefs across games. Furthermore, the results suggest that, empirically, both risk and ambiguity aversion are important factors in determining how subjects respond to the natural uncertainty that arises in strategic interactions.

There is a wealth of evidence, tracing back to Knight (1921) and Keynes (1921) via Ellsberg (1961), 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 (2015a) find that subjects exhibit higher levels of ambiguity aversion in games than in a 3-colour Ellsberg urn task. This chapter is the first to give a complete picture of the role of uncertainty aversion in games: we document the first-order (how do subjects respond to strategic uncertainty?) and second-order (how do subjects respond to opponents who face strategic uncertainty?) effects of both risk aversion and ambiguity aversion.[4] This chapter 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.[5]

This chapter proceeds as follows. Section 2.1 provides a brief overview of some relevant theoretical considerations, while section 2.2 presents the experimental design in detail. Section 2.3 presents the experimental results, and section 2.4 presents an additional treatment motivated by these results. Section 2.5 provides a discussion and conclusion. Appendix B presents additional results not included in the main text as well as proofs and the instructions for subjects. Appendix A contains some additional background on the theory of games with ambiguity averse agents, a topic that is also explored in chapter 4.

2.1 Theoretical Considerations

The identification of the relationship between preferences and "reasonable" strategies in the testing game can be motivated either using rationalizability arguments or equilibrium concepts (i.e. Nash equilibrium for SEU subjects and an appropriately chosen ambiguity averse equilibrium concept for ambiguity averse subjects). We use rationalizability rather than equilibrium concepts in the body of the chapter because of the stronger epistemic assumptions required to justify the use of equilibrium; however, we provide two equilibrium derivations (in the style of Lo (2009) and Dow and Werlang (1994)) in Appendix A for the interested reader.

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 typically imply a strict preference for mixed strategies or are not able to define a utility level for

[4]Interestingly, Ivanov (2011) asks the dual of our research question by estimating ambiguity aversion from behaviour in games. A more detailed literature review can be found in appendix B.1.5.

[5]Heinemannn et al. (2009) also recognized 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.

mixed strategies at all. Chapter 4 and Eichberger and Kelsey (2000) contain extensive discussion on the role of mixed strategies in games with ambiguity averse agents. As Chapter 4 demonstrates it is possible to extend the theory in this section to allow for subjects to use mixed strategies if the 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 by the theory described below.

2.1.1 Rationalizability

In this section we formalize the interactions between preferences and rationalizability in the game in figure 2.1. We use Epstein's (1997) notion of rationalizability, which allows for a generalization of SEU preferences to the 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 preferences of the agents*, and therefore generalizes the more familiar Pearce (1984)/Bernheim (1984) rationalizability 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:

$$U(a_R) = \min_{\phi_R \in \Phi_R} \sum_{a_C \in \{X,Y\}} u_R(m_R(a_R, a_C))\phi_R(a_C) \quad \forall a_R \in \{A,B,C\}, \tag{2.1}$$

where $u_R(.)$ is the row player's utility over monetary outcomes, $m_R(a_R, a_C)$ is the monetary payoff for outcomes as shown in figure 2.1, and $a_R \in \{A, B, C\}$ and $a_C \in \{X, Y\}$. 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 2.1. *Consider the normal form game in figure 2.1. 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 2.1. See Appendix B.2. □

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 2.1 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. Furthermore, Epstein (1997) provides a fully worked application of his rationalizability concept using MEU preferences, and it is also the preference structure underlying Lo (2009) (which is discussed in detail in Appendix A). Alternatives such as Choquet Expected Utility (Schmeidler (1989)), which underlies Dow and Werlang (1994) (also discussed in discussed in detail in Appendix A), or Klibanoff et al.'s (2005) smooth ambiguity aversion preferences would work just as well. In fact, in a recent working paper Battigalli et al. (2015) independently[6] establish a related result, illustrated using an example on a similar

[6]The experimental design in this chapter was created in late 2013, and experiments were conducted in early 2014. The first version of this chapter that was circulated outside UBC was in November 2014. The earliest version of Battigalli et al. (2015) available on Google Scholar is dated September 2014, and the first version that I read was dated May 2015.

game to the one studied here, for the case of smooth ambiguity aversion preferences. Their example also highlights and reaffirms that "the risk and ambiguity attitudes of the (column player) are immaterial."

For the remainder of the chapter, we shall restrict ourselves to four possible preference types that can be held by the row player. The row player may either satisfy or violate $u(25)+u(14) > 2u(18)$ and the row player may either have SEU or MEU preferences, as shown in table 2.1. Applying proposition 2.1 to this classification structure, it is immediate that the rationalizable set is $\{A, X\}$ if and only if the row player has Type 1 preferences.

	SEU	MEU
$u(25)+u(14) > 2u(18)$	Type 1	Type 2
$u(25)+u(14) \leq 2u(18)$	Type 3	Type 4

Table 2.1: The type space for preferences.

Implicit in proposition 2.1, because of the use of rationalizability as the solution concept, is an assumption that the column player knows the row player's preferences. Under this assumption the column player has no role to play in determining the rationalizable set. We can, however, relax the link between the row player's preferences and the column player's beliefs by modeling the game as a Bayesian game with one sided private information.[7] Suppose there is a common prior, $\alpha \in [0,1]$, that the row player has Type 1 preferences. The sum of the probabilities of Type 2, 3 or 4 row players is, therefore, $1-\alpha$.[8]

Proposition 2.2. *Consider the normal form game in figure 2.1. Suppose that the game is a Bayesian game with one-sided private information and the row player has Type 1 preferences with prior probability* $\alpha \in [0,1]$. *It follows that:*

1. *if* $\alpha > \bar{\alpha}$, *then the rationalizable set is* $\{(A,X)\}$,

[7] It is possible to relax the knowledge requirements even further by modeling the game as a Savage game ala Grant et al. (Forthcoming), or by allowing both the row and column to have private preference information. These richer models will generate a similar result, replacing the $\alpha > \frac{5}{9}$ condition in the proposition with an object that is harder to interpret.

[8] Note that in this structure the prior is a singleton, and any ambiguity is captured through the mapping from types to actions; that is, all ambiguity is ambiguity with respect to strategies while types introduce only objective risk.

2. *if $\alpha \leq \bar{\alpha}$, then all actions are part of a rationalizable profile, but a Type 1 row player will never be observed to play C.*

where $\bar{\alpha} = \frac{u(22)-u(12)}{u(22)+u(20)-2u(12)}$. $\bar{\alpha} = \frac{5}{9}$ for a risk neutral column player. Furthermore, $\bar{\alpha} > \frac{1}{2}$ and is decreasing in the column player's level of risk aversion.

Proof of proposition 2.2. See Appendix B.2. □

The key implication of proposition 2.2 is that Y is not rationalizable for the column player if $\alpha > \frac{5}{9}$.

In the next section we shall discuss the procedure that was used to classify subjects into types, and to elicit subjects beliefs regarding their opponent's type. We close this section by noting that although both risk aversion and ambiguity aversion have identical effects on the rationalizable sets, there are some differences in the effects of risk and ambiguity aversion in the equilibria of the game in figure 2.1. It is straightforward to establish that the Nash equilibrium set for SEU subjects with high risk aversion (preferences that satisfy $u(25)+u(14) < 2u(18)$) includes equilibria where the row player mixes over exactly two of their strategies. There are no fully mixed equilibria, and no equilibria that involve pure C for the row player. However, for ambiguity averse subjects there are fully mixed strategies and, for some equilibrium concepts, pure C can be sustained in an equilibrium. As demonstrated in Appendix A the equilibria of both Lo (2009) and Dow and Werlang (1994) can support fully mixed equilibria and the latter may also support equilibria where the column player plays pure C.[9]

2.2 Experimental Design

The experimental design involves 3 two-player normal form games[10], with subjects playing each of the games as both the row player and the column player. The

[9]I thank Simon Grant (private correspondence) for demonstrating that pure C can be supported under ambiguity averse equilibrium concepts that do not impose consistency of the support of beliefs with the support of the actions used, such as Eichberger and Kelsey (2000) and Dow and Werlang (1994).

[10]Subjects were also presented with a fourth game, which was intended to be used as an additional diagnostic test in the case of contingencies that did not arise in the data. The fourth game is available in an earlier working paper version of this chapter.

first two games were used to measure the risk and ambiguity preference of the row player, and the column player's beliefs regarding his opponent's risk and ambiguity preference; these games are referred to as the classification games. The remaining game was then used to test whether the subject's preferences and beliefs, coupled with an appropriate solution concept, are associated with behaviour in normal form games; this game is referred to as the testing game. The testing game was constructed so that behaviour as the row player is expected to depend only on the subject's preferences (and not her beliefs regarding her opponent's preferences) and behaviour as the column player is expected to depend only on the subject's beliefs regarding his opponent's preferences (and not his own preferences).

The experiment was designed to mitigate a number of factors that have been identified as affecting the elicitation of ambiguity preferences. Recent research has identified that subject confusion can significantly lower measured ambiguity aversion (Chew et al. (2013)), that experiments with ambiguity can be particularly susceptible to violations of incentive compatibility across tasks (Baillon et al. (2014), Azrieli et al. (2014)) and that framing effects may be particularly strong (Chew et al. (2013)). 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. 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. This chapter is the first to elicit ambiguity preferences in a framing that is consistent with normal form games.

A range of other important, yet technical, experimental considerations, including a description of both the random payment mechanism used and the detailed and incentivized comprehension questions, are discussed in detail in the experimental methodology appendix B.1.

2.2.1 Classification games

There are two classification games. The first game elicits ambiguity preferences (for row players) and beliefs regarding the opponent's ambiguity preferences (for column players), while the second game elicits risk preferences and beliefs regarding the opponent's 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.

The first (ambiguity) classification game is shown in figure 2.2. For the row player, this game is isomorphic to a simplified version of the ambiguity elicitation process used in Epstein and Halevy (2014).[11] The game involves two ball draws, one from the U urn (figure 2.3) and one from the K urn (figure 2.4). 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 was tested via a series of comprehension questions that are discussed in detail in appendix B.1.3.

	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 2.2: 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.

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

[11] Epstein and Halevy (2014) use four tasks to elicit ambiguity preferences, whereas we have simplified this to one task. We assume symmetry of beliefs, as discussed below, which allows us to drop two of the four tasks. Furthermore, we are only interested in ambiguity aversion here (and not ambiguity seeking behaviour), allowing us to remove one additional task.

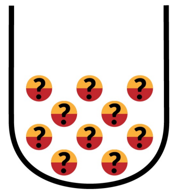

Figure 2.3: U urn. The U urn consists of 10 balls, each of which may be either red or yellow. The total number of red balls in the urn lies between 0 and 10.

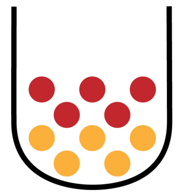

Figure 2.4: K urn. The K urn contains 5 red and 5 yellow balls.

\$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, which appears reasonable because red and yellow are interchangeable labels.[12] If a subject has 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.

Now, consider the column player in the game in figure 2.2. The column player is tasked with predicting the row player's action. If the outcome of the game is (S, M') or (M, S') then the column player receives \$0 in all states. So the rational column player will play S' if they believe that the row player is more likely to choose S, and will play M' if they believe that the row player is more likely to choose M.

The second classification game is shown in figure 2.5 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

[12] Note that this is an assumption regarding the symmetry of beliefs, and not an assumption on the subjects preferences regarding red or yellow balls. However, even 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 the red balls in this formulation, 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 2.5: 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.

predict the row player's preferences. A highly risk averse row player will choose H, while a risk neutral (or low risk aversion) row player will choose L. Subjects with intermediate levels of risk aversion will choose I. Similarly, a subject that chooses L' believes his opponent to have low risk aversion, a subject that chooses I' believes his opponent to have intermediate risk aversion, and a subject that chooses H' believes his opponent to be rather risk averse.

We now consider the mapping from responses in the preference measurement tasks to the type structure that was introduced in section 2.1.1. For an extensive discussion of the auxiliary assumptions underlying the classification procedure see Section B.1.2. The mapping is straightforward: a subject who selects S in the game in figure 2.2 is considered to be ambiguity neutral and therefore is either type 1 or type 3, and a subject who selects M is considered to be ambiguity averse and therefore is either type 2 or type 4. A subject who selects L in the game in Figure 2.5 is considered to have low risk aversion and therefore is either type 1 or type 2, and a subject who selects H is considered to have high risk aversion and is either type 3 or type 4. Subjects who selected I are not classified, on the basis that their risk preferences are such that it is not possible to determine whether they would prefer to play C or a 50-50 mix between A and B in the game in Figure 2.1. Column players are classified in an analogous fashion.

Table 2.2 presents a summary of the mapping from row player choices in the preference measuring games to types and table 2.3 presents a summary of the mapping for column players.

Row Player Strategies	S	M
L	Type 1 $\{A\}$	Type 2 $\{A,B,C\}$
I	N/A	N/A
H	Type 3 $\{A,B,C\}$	Type 4 $\{A,B,C\}$

Table 2.2: The classification procedure for row players, as a function of responses to the classification games. Subjects that select I in the risk measurement game are not classified. The rationalizable set in the testing game for each type is also indicated.

Column Player Strategies	S'	M'
L'	Type $1'$ $\{X\}$	Type $2'$ $\{X,Y\}$
I'	N/A	N/A
H'	Type $3'$ $\{X,Y\}$	Type $4'$ $\{X,Y\}$

Table 2.3: The classification procedure for column players, as a function of responses to the classification games. Subjects that select I' in the risk measurement game are not classified. A type $1'$ subject believes that their opponent is a type 1 subject. The rationalizable set in the testing game for each type is also indicated.

2.2.2 Testing game

The testing game, reproduced in figure 2.6, is the same game that was introduced in the introduction. To recap, the testing game will be used to test whether the strategies played by subjects change with preferences (and beliefs) in the same manner that the rationalizable set of strategies changes with preferences (and beliefs). For row players, subjects who are classified as type 1 (low risk aversion and ambiguity neutrality) have a unique rationalizable set of $\{A\}$ while for other subjects all strategies are rationalizable. For column players, subjects who are classified as type $1'$ (believe their opponent to be type 1) have a unique rationalizable set of $\{X\}$ while for other subjects both strategies are rationalizable.

The interaction of risk and ambiguity preferences in the testing game is of a precise nature: the row player needs only be sufficiently risk averse *or* ambiguity

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

Figure 2.6: The testing game. The row player's rationalizable strategy set varies with her preferences. The column player's rationalizable strategy set varies with his belief of the row player's preferences.

averse for the rationalizable set to be the complete action space. The hypothesis of subject level behaviour are summarized in tables 2.2 and 2.3. For row players, subjects that choose S and L in the classification games are uniquely expected to play A, while all other subjects may play any of A, B or C. For column players, subjects that choose S' and L' in the classification games are uniquely expected to play X, while all other subjects may either of X or Y.

2.2.3 Experimental conditions

All sessions were held in the ELVSE lab at the Vancouver School of Economics. There were 10 sessions held in March and April 2014, and a further 10 sessions held in September 2014, with between 8 and 12 subjects per session. There were some minor changes to the instructions made between the April and September sessions, otherwise all sessions were identical. The September instructions are included in the appendix. Sessions lasted between 60 and 90 minutes, and the average payment was $26.60. 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 the Redwood experimental software tool (Pettit et al. (2015)).

2.3 Results

This section presents the experimental results. We split this section in two: first we present the results and discuss the statistical methods used to produce them. Second, we discuss the economic implications of the results. There are three key dimensions to the results presented in this section:

- Risk and ambiguity preferences are positively correlated;

- Row player behaviour in the classification games is correlated with row player behaviour in the testing game, as predicted by theory;
- Column player behaviour in the classification games is independent of column player behaviour in the testing game.

2.3.1 Preferences

We begin our tour of the results with a look at the preferences of our subjects.

		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 2.4: 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 ($p = 0.002$).

Table 2.4 presents the classification of subjects into types based on their responses as row players in the preference measuring games. 38% 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 papers, and lower than that measured previously in 2-urn Ellsberg tasks.[13] We find that 30% of subjects with low risk aversion are ambiguity averse, while more than half of the subjects with high risk aversion are also ambiguity averse.[14] A Pearson's χ^2 exact test rejects the

[13]Chew et al. (2013) provide an overview of previous Ellsberg urn experimental results. For the 2-urn case, as used in this chapter, 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-colour) Ellsberg task, previous studies have found that between 79% and 8% of subjects are ambiguity averse (with a weighted mean of 27%). Pilot sessions suggest that the framing of the decision problem in a bi-matrix format may cause lower levels of ambiguity aversion to be measured, although this effect is mitigated by the use of comprehension questions. This supports the conclusions of Chew et al. (2013), who argue that less ambiguity aversion will be measured in more complex environments.

[14]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 behaviour and additional demographic analysis is relegated to the appendix. Due to a technology

null hypothesis of independence of risk and ambiguity preference at the 1% level ($p = 0.002$).[15]

Result 2.1. *Risk and ambiguity preferences are positively correlated.*

2.3.2 Row player behaviour

We now turn our attention to the testing game, beginning with row player behavior in the game in figure 2.6. Recall that Type 1 subjects (low risk aversion and ambiguity neutral) are expected to play the unique Bernheim/Pearce rationalizable strategy of A, while for other subjects all strategies are rationalizable.

For this analysis, we restrict our sample to those subjects who "passed" the comprehension tasks. A subject was considered to pass the comprehension task, as a row player, if they answered the comprehension questions correctly on the first attempt for all three games they played as the row player. This criteria was determined ex-ante[16], and if a subject made any mistakes they were considered to have failed. While this criteria may seem rather harsh, it is important to note that the study of ambiguity aversion in games prohibits the use of experience to train subjects: if the subjects were to play the game repeatedly, even using a stranger matching protocol, they would learn about the distribution of behaviour in the population and thereby reduce the degree of strategic ambiguity in later interactions. Effectively, repeated play between ambiguity averse subjects can change the equilibrium set even when there are no supergame effects present. A strict comprehension criteria is therefore used as a substitute for experiential learning to identify a

mishap, demographics were not collected for two sessions. The change in sample has a larger effect on the results than controlling for demographic factors does, so demographic controls are not used in the main analysis.

[15]The choice of statistical tests for categorical data is an oft-debated topic. Throughout this chapter, every p-value could be calculated in multiple different 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 severeley unbalanced. An earlier version of the chapter used Fisher exact tests throughout and generated the same conclusions.

[16]There are many different inclusion criteria that could be constructed from the comprehension data. See appendices B.1.3 and B.3.3 for a discussion of the comprehension data and inclusion criterion.

sub-population of subjects that thoroughly understand the structure of the games they are playing.

We present results for the full sample, as a robustness check, in appendix B.3.3. Moving to the full sample strengthens the row player results presented in this subsection, because the subjects who failed the comprehension tests display approximately the same effect of preferences on behaviour in the testing game and the increased sample size helps with statistical power. We do not combine the two samples, however, because subjects who failed the comprehension tests show a statistically significant level effect: across all preference types, subjects who failed the comprehension play A less often in the testing game.

Table 2.5 presents an aggregate view of the row player data, detailing the percentage of subjects that play A in the game in figure 2.6 as a function of their preference type, for subjects who passed the comprehension tasks.

	Type 1 LRA, AN	Type 2 LRA, AA	Type 3 HRA, AN	Type 4 HRA, AA	Total
Pr(A)	0.74 (N=57)	0.58 (N=19)	0.61 (N=18)	0.41 (N=22)	0.63 (N=116)

Table 2.5: Proportion of subjects that play A in the game in figure 2.6 by preference type, restricted to subjects that passed the comprehension tasks. LRA denotes low risk aversion, HR denotes high risk aversion, AN denotes ambiguity neutral and AA denotes ambiguity averse.

We now turn to analyzing the effects of preferences on the probability of a subject playing A in the testing game. We focus on the choice of A because the theory outlined in section 2.1.1 only provides predictions of A or not A and is agnostic over the choice of B or C. We shall, however, let the data speak regarding the choice of B or C in section 2.3.2.

Table 2.6 presents the estimates of the effects of preferences on behaviour in the testing game. The top two lines present the average, unconditional, effects of ambiguity aversion and risk aversion, respectively. The next four lines present conditional effects of risk and ambiguity aversion, and the final line presents the joint effect of both risk and ambiguity aversion (the difference in behaviour between Type 1 and Type 4 subjects). All p-values are calculated non-parametrically using Pearson's χ^2 test of independence.

As we can see from table 2.6 the average, unconditional, effect of both risk and

Effect of	Restricted to	$\Delta Pr(A)$
Ambiguity Aversion		-0.22*
		(0.020)
Risk Aversion		-0.20*
		(0.036)
Ambiguity Aversion	LRA Subjects	-0.16
		(0.194)
Ambiguity Aversion	HRA Subjects	-0.20
		(0.204)
Risk Aversion	AN Subjects	-0.13
		(0.307)
Risk Aversion	AA Subjects	-0.17
		(0.278)
Ambiguity Aversion & Risk Aversion		-0.33**
		(0.006)

Table 2.6: Change in proportion of subjects playing A in the testing game, as a function of subject preferences,restricted to subjects that passed the comprehension tests. $N = 116$. LRA denotes low risk aversion, HR denotes high risk aversion, AN denotes ambiguity neutral and AA denotes ambiguity averse. Pearson's χ^2 p-value shown in brackets. * indicates value is significantly different from 0 using a non-directional test at the 5% level and ** indicates significance at the 1% level.

ambiguity on the probability of choosing A are negative, large and statistically significant. These results are, however, biased estimates of the true effects: because both effects are negative, and risk and ambiguity preferences are negatively correlated, they suffer from omitted variable bias. Nevertheless, we report them here for comparison to other data sets that may not include measures of both risk and ambiguity preference.

There is clear evidence in table 2.6 for a joint effect of risk and ambiguity preferences on behaviour in the game in figure 2.6: the bottom row indicates that a highly risk averse and ambiguity averse (Type 4) subject who passed the comprehension tests is 33 percentage points less likely to play A than a slightly risk averse and ambiguity neutral (Type 1) subject who passed the comprehension tests. The effect is large and highly statistically significant. Decomposing the joint effect into the individual effect of risk and ambiguity aversion is more difficult, as the smaller sample sizes generate imprecise estimates of the effect size.[17] The effect of ambiguity aversion conditional on low risk aversion is identified from the difference in behaviour between Type 1 and Type 2 subjects, so that the small sample of 19 Type 2 subjects who passed the comprehension test reduces the power of the χ^2 test of independence.

Result 2.2. *A joint measure of risk and ambiguity aversion is associated with a lower probability of playing A in the game in figure 2.6.*

While restricting our analysis to the choice of A or not A has strong theoretical motivations, it does involve throwing away information regarding the choice of B and C. The next section makes use of this information to investigate if there are any differential effects of risk and ambiguity aversion in the testing game.

Analysing the choice between B and C

Table 2.7 breaks down the choice of "not A" and investigates how preferences influence the decisions of subjects choosing between B and C. The point estimates provided are mechanically equivalent to the point estimates from a multinomial logit regression where A is specified as the base outcome.[18] The p-values are calculated non-parametrically, using either Pearson's χ^2 test for larger samples or Fisher's exact test for smaller samples (p-values that were calculated using Fisher's exact test are denoted with a F).[19]

The coefficients in table 2.7 are interpreted as follows. Consider the top left estimate of 0.19: this indicates that an ambiguity averse subject is 19 percentage

[17] One potential method for decomposing the effects of risk and ambiguity aversion is to use an ANOVA. The ANOVA yields a main effect of ambiguity aversion on behaviour that is statistically significant at the 10% level ($p = 0.07$). The interpretation of this result is difficult, however. As discussed in section 2.3.4 the theoretical effect of ambiguity aversion on the rationalizable set is only identified for subjects with low risk aversion, which implies that any effect of ambiguity aversion for high risk aversion subjects must be driven by a different mechanism. The ANOVA exercise conflates these two different mechanisms.

[18] We specify A as the base outcome so that we can still interpret the decision as a choice between A and something that isn't A, keeping to the spirit of the theoretical predictions.

[19] A word of warning on sample sizes: only 11 of the 116 subjects who passed the comprehension tests chose B in the testing game (32 chose C and the remainder chose A). However, more than a quarter of Type 2 subjects (low risk aversion, ambiguity averse) chose B, and it is this effect which drives the results presented below. We use Fisher's exact test instead of Pearson's χ^2 test when considering the choice of B because of its favorable small sample properties.

points more likely to choose B, conditional on not choosing C, than an ambiguity neutral subject. The results suggest that the effect of ambiguity aversion, particularly amongst low risk aversion subjects, is to move subjects away from A and towards B: a Type 2 subject is 25 percentage points more likely to choose B, conditional on not choosing C, than a Type 1 subject. The joint effect of risk and ambiguity aversion, on the other hand, is to move subjects away from A and towards C (see the bottom row of the table).[20]

Effect of	Restricted to	$\Delta\frac{N_B}{N_B+N_A}$	$\Delta\frac{N_C}{N_C+N_A}$
Ambiguity Aversion		0.19*	0.16
		(0.033)F	(0.099)
Risk Aversion		-0.02	0.24*
		(1.000)F	(0.011)
Ambiguity Aversion	LRA Subjects	0.25*	-0.01
		(0.024)F	(1.000)F
Ambiguity Aversion	HRA Subjects	0.09	0.20
		(0.590)F	(0.231)
Risk Aversion	AN Subjects	0.02	0.13
		(1.000)F	(0.341)F
Risk Aversion	AA Subjects	-0.13	0.34
		(0.622)F	(0.050)
Ambiguity Aversion & Risk Aversion		0.12	0.33**
		(0.251)F	(0.007)

Table 2.7: Change in the values of $\frac{N_B}{N_B+N_A}$ and $\frac{N_C}{N_C+N_A}$ in the testing game as a function of subject type, where N_A is the number of subjects selecting A. Restricted to subjects that passed the comprehension tasks. $N = 116$. LRA denotes low risk aversion, HR denotes high risk aversion, AN denotes ambiguity neutral and AA denotes ambiguity averse. Figures in brackets are p-values calculated under a null hypothesis that the coefficient is equal to zero under either Pearson's χ^2 test or a Fisher exact test (denoted by F). The coefficients are equivalent to the probabilities implied by a saturated multinomial logit regression with a choice of A denoted as the base outcome.

Result 2.3. *A joint measure of risk and ambiguity aversion is associated with a higher probability of playing C in the game in figure 2.6.*

[20] An ANOVA on the choice between A and B finds a significant main effect of ambiguity aversion ($p = 0.041$). An ANOVA on the choice between A and C finds a significant main effect of risk aversion ($p = 0.012$). Again, care should be taken in interpreting the ANOVA results, see footnote 17.

Result 2.4. *There is evidence that ambiguity aversion is associated with a higher probability of playing B in the game in figure 2.6, for low risk aversion subjects who passed the comprehension tests.*

2.3.3 Column player behaviour

We now turn our attention to column player behaviour in the game in figure 2.6. Recall that Type $1'$ subjects have a unique rationalizable action of X, while both X and Y are rationalizable for all other subjects. We should therefore, expect subjects who believe their opponent to be risk and ambiguity averse to play X less often than subjects that believe their opponent to be ambiguity neutral and have low risk aversion.

Believe opponent to be	Type $1'$ LRA, AN	Type $2'$ LRA, AA	Type $3'$ HRA, AN	Type $4'$ HRA, AA	Total
$\Pr(X)$	0.74 (N=47)	0.83 (N=18)	0.65 (N=20)	0.59 (N=17)	0.72 (N=102)

Table 2.8: Proportion of subjects that play X in the game in figure 2.6, by belief type. LRA denotes low risk aversion, HR denotes high risk aversion, AN denotes ambiguity neutral and AA denotes ambiguity averse.

Table 2.9 shows the change in the proportion of subjects that play X as a function of preferences, restricted to subjects who passed the comprehension tests for their three games as the column player. The model structure is analogous to Table 2.6. Interestingly, there is no evidence of any effect of beliefs regarding an opponent's preferences on column player behavior in the testing game. In particular, the point estimate of the average effect of a subject believing he is facing an ambiguity averse opponent, relative to believing he is facing an ambiguity neutral opponent, is 0 to two decimal places. Subject beliefs regarding their opponent's preferences are statistically independent of their behaviour as the column player in the testing game.

Result 2.5. *Neither beliefs over risk preferences nor beliefs over ambiguity preferences explain behaviour for the column player in the game in figure 2.6.*

Effect of beliefs over opponent's	Restricted to subjects who believe their opponent is	$\Delta Pr(X)$
Ambiguity Aversion		-0.00
		(0.980)
Risk Aversion		-0.15
		(0.112)
Ambiguity Aversion	Low Risk Aversion	0.09
		(0.448)
Ambiguity Aversion	High Risk Aversion	-0.06
		(0.699)
Risk Aversion	Ambiguity Neutral	-0.10
		(0.431)
Risk Aversion	Ambiguity Averse	-0.25
		(0.109)
Ambiguity Aversion & Risk Aversion		-0.16
		(0.226)

Table 2.9: Change in proportion of subjects playing X in the testing game, as a function of reported beliefs over opponents' preferences, restricted to subjects that passed the comprehension tasks. $N = 102$. Pearson's χ^2 p-value shown in brackets. * indicates value is significantly different from 0 using a non-directional test at the 5% level and ** indicates significance at the 1% level.

2.3.4 Discussion of results

Result 2.1 established a correlation between risk and ambiguity preferences.[21] 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.[22] From a theoretical viewpoint, axiomatic models of pref-

[21] The data also presents a strong relationship between a subject's preferences and his beliefs regarding his opponent's preferences. Further investigation of this relationship is relegated to appendix B.3.2, however, as there is evidence (discussed in section 2.4) that subjects may have had a low level of confidence in their predictions of their opponent's behaviour. In this context of uncertainty, projecting their own preferences onto their opponent provides a sensible focal point that was followed by most subjects.

[22] 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).

erences (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 behaviour 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.

In the context of measuring the effects of risk and ambiguity aversion in games, correlation between risk and ambiguity preferences can lead to an omitted variable bias if only one factor is considered. Given that the current chapter is the first to consider the effects of both risk and ambiguity, we have evidence that previous estimates of risk aversion in games (Goeree et al. (2003), for example) and ambiguity aversion in games (Kelsey and le Roux (2015a), for example) may overestimate effect sizes. This effect is demonstrated clearly in table 2.6. The unconditional effect of risk aversion on the probability of playing A is estimated to be a statistically significant -0.20. Because of the omitted variable bias, this figure is larger than both the estimated effect of risk aversion for ambiguity neutral subjects (-0.13) and the estimated effect of risk aversion for ambiguity averse subjects (-0.17).

Result 2.2 established a joint effect of risk and ambiguity preferences on behaviour in the testing game, but did not identify the effect of ambiguity aversion separately from the effect of risk aversion. To understand why we do not try and pull apart average affects of risk and ambiguity, recall table 2.2 which summarizes the mapping from our Typology to rationalizable sets. The rationalizable sets differ between Type 1 subjects and subjects of any other Type. This generates a clear prediction that we should expect different behaviour between Type 1 subjects and Type 2 subjects (i.e. an effect of ambiguity aversion conditional on low risk aversion). We do not, however, have a clear prediction of a difference in behaviour between Type 3 and Type 4 subjects (i.e. an effect of ambiguity aversion conditional on high risk aversion). This differential prediction means that many of the standard statistical tests for categorical data of this nature, such as ANOVA, do not identify the effects that we are theoretically interested in here. The only theoretically motivated test for the effect of ambiguity aversion on behaviour is to test the behaviour between Type 1 and Type 2 subjects (which is the test that is conducted in the third row of table 2.5).

We have, so far, focused our analysis on the effect of type on behaviour in the testing game. However, also of interest is the behaviour of Type 1 subjects (who were measured to have low risk aversion and are consistent with ambiguity neutrality), who have a unique rationalizable strategy of A in the testing game. Table 2.5 shows that 26% of Type 1 subjects who passed the comprehension tasks, and thereby demonstrated an understanding of the payoff structure of the games, failed to play a rationalizable strategy. As discussed in appendix B.1.2 to the extent that there is misclassification we expect to overclassify subjects as Type 1, which may explain some of the non-rationalizable behaviour observed here: some of these subjects may perceive their opponent to be a source of ambiguity but do not perceive an Ellsberg urn to be a source of ambiguity.

Results 2.3 and 2.4 established that ambiguity aversion appears to be driving a movement away from A towards B, and the joint effect of risk and ambiguity aversion is a movement towards C. A choice of C is the natural 'safe' strategy in the game in figure 2.6, and it is not surprising that risk and ambiguity averse subjects show a strong affinity for choosing C. The choice of B amongst ambiguity averse subjects, however, is more interesting. While any explanation of why ambiguity averse subjects with low risk aversion have a tendency to choose B must be taken with caution given the small sample sizes involved, we provide a speculative hypothesis. Identifying and avoiding ambiguous prospects while being comfortable to take on risks at small stakes is a sophisticated set of preferences for a subject to hold.[23] For an ambiguity averse subject, a choice of C in the testing game is the obvious choice – perhaps a little too obvious. Expecting their opponent to best respond to C, a sophisticated ambiguity averse subject may apply 2 levels of reasoning and best respond to the best response to C: a choice of B. There is no evidence of such reasoning amongst high risk aversion subjects. An interesting hypothesis

[23]Chew et al. (2013) provide evidence that, in more complicated situations, the expression of ambiguity aversion is more likely amongst sophisticated subjects precisely because sophistication is required to identify and avoid ambiguity.

for future work is, therefore: do ambiguity averse subjects with low levels of risk aversion engage in higher levels of strategic reasoning than ambiguity neutral or highly risk averse subjects?

Result 2.5 is perhaps the most intriguing result. Why are subjects, who demonstrated a sound understanding of the games being played by passing the comprehension test, not behaving in a fashion that is consistent with their self reported beliefs regarding their opponent's preferences? There are at least three plausible explanations. First, as discussed in section 2.1.1, subjects may have very imprecise beliefs regarding their opponents preferences. Second, subjects may be unable to implement the higher order reasoning needed to understand the relationship between opponent preferences and opponent behaviour. Third, subjects might believe that their opponents are not rational or do not understand the game structure.

There is evidence that the first explanation may hold. A natural reaction for a subject who is very unsure about their opponent's preferences is to imagine that their opponent is just like them, and best respond to what they would play as the other player. Given that we observe each subject play the game as both the row and column player, we can check if this is the case. Table 2.10 presents the data (for all subjects that are classified a type as both the row and column player) and clearly demonstrates that column players have a tendency towards best responding to their own row player action. Notice, however, that the effect does not appear to be symmetric: if row players had a tendency to best respond to their own column player action then we should expect to see more subjects in the $\{B,Y\}$ cell, and less subjects choosing C.[24]

	X	Y
A	85	16
B	7	8
C	28	30

Table 2.10: The number of subjects that chose each outcome, under the counter factual that each subject was matched with themselves. Subjects who selected either I or I' in the risk measurement game are excluded. A Pearson's χ^2 test rejects the null hypothesis of independence at the 1% level ($p = 0.000$).

[24] In fact, if subjects were always responding to what they would do in the other role then we would see everyone playing the unique Nash equilibrium.

In the next section, we introduce an additional treatment that tests between the first two explanations. If subjects are unable to reason at a high enough level to respond to their opponent's preferences, then there are serious implications for the standard set of game theoretic tools. To conduct this test, we provide subjects with a signal about both their opponent's preferences and comprehension scores.

2.4 Unpacking the Results: A New Treatment

In the previous section we provided evidence that subjects' preferences affect their behaviour in the main testing game in the direction predicted by theory but their beliefs over their opponent's preferences do not affect behaviour, thereby establishing partial support for the role of ambiguity aversion in game theory. The data in the previous section does not, however, allow us to identify the reason for the failure of beliefs over opponent's preferences to affect behaviour. We hypothesize that subjects have a low level of confidence in their reported beliefs and are, therefore, essentially 'guessing' in the belief measuring games. An alternative explanation is that subjects are unable to hold a consistent mental model of their opponent, and therefore respond differently in different environments. The results presented in this section, which are remarkably strong, support the former explanation of behaviour in the testing game.

We propose a simple new treatment: instead of eliciting subjects' beliefs over their opponent's preferences, we simply inform the subject of the choices that their opponent made in the preference measuring games. If subjects use the signal of their opponent's preferences to inform their behaviour as the column player in the testing game then we can conclude that subjects are able to form a coherent mental model of their opponent, and that it was a lack of confidence in beliefs that was driving behaviour in the original treatment. The results are stark: when a subject has observed his opponent's preferences, the opponent's preferences are strongly correlated with the subject's behaviour as the column player in the main testing game.

For the additional treatment, 130 new subjects were recruited and bought to the ELVSE lab.[25] The sessions lasted between 30 and 40 minutes and average earnings

[25]There were actually 131 subjects, but one subject had already participated in the original treat-

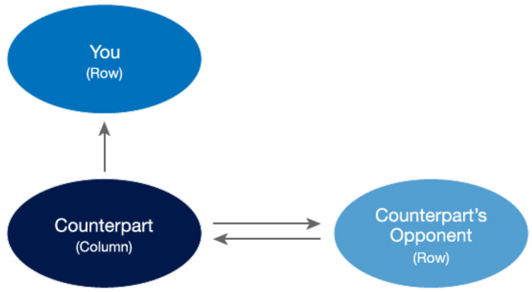

Figure 2.7: The matching protocol diagram, as shown to subjects. The subject was designated ”You”, and the subject's opponent was designated ”Counterpart”. Arrows are used to represent strategic interactions: the originator's choice influences the target's payoff. Note that the subjects were shown a transposed version of the game in figure 2.6, so the row and column labels are reversed.

were C$22.56. Each subject was matched with an opponent from the original experimental sessions, with the subject's opponent being selected in a pseudo-random fashion[26].

91 subjects were shown the choice made by their opponent in both the risk and ambiguity measuring games (i.e. their opponent's choice as the row player in the games in figures 2.2 and 2.5), while 39 subjects were shown only their opponent's choice in the ambiguity measuring game. Each subject was also given a summary of his opponent's performance on the comprehension tasks. The pool of opponents was heavily skewed towards opponents who had performed well on the comprehension tasks – the goal was not to investigate how a subject responded to his opponent's comprehension score, but rather to provide the subject with a credible signal that his opponent's choices represent reasoned decisions and not random behaviour generated by a poor understanding of the underlying game.

ment and was therefore excluded from the data. The subject had created two accounts in ORSEE, which wasn't noticed until after the session had finished. Given that there was no learning possible in the experiment, and no subject could observe any other subject's choices until after all decisions were finalized, we are still able to use the other subject's data from the offending session without fear of contamination.

[26]The pool of subjects was chosen in a manner that was designed to provide better statistical power than would have been provided using true random sampling. 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.

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 2.7: 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, and could earn up to a $1 bonus for completing the comprehension questions correctly (see appendix B.1.3 for details). As in the main treatment, we exclude subjects who performed poorly on the comprehension questions: 121 of the 130 subjects answered the comprehension questions correctly on their first attempt.[27]

The payoff structure that was used in the preference measuring games was designed for the original treatment in which an unobserved variable (the subject's true preferences) was expected to influence behaviour in both the preference measuring games and the testing game. As a consequence, we would expect the results in the original 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 2.5) it is clear that, in laymen's terms, strategies L and I are "dangerous", and strategy H is "safe". In the testing game (figure 2.6) 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.

The data for this new treatment is summarized in tables 2.11 and 2.12. Table 2.11 presents the data for subjects who observed both their opponent's risk and

[27] 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. See appendix B.3.3 for further details.

ambiguity preferences. Remarkably, every single subject who observes a Type 1 opponent plays X, which is their unique rationalizable strategy. There is also a large drop off in the proportion of subjects that play X, conditional on observing a highly risk averse opponent (Type 3 or Type 4) which is also in concordance with theory. Table 2.12 presents the data for subjects who observed only their opponent's ambiguity preferences.

Opponent type Observed preferences	Type 1 LRA, AN	Type 2 LRA, AA	Type 3 HRA, AN	Type 4 HRA, AA	Total
Pr(X)	1.00 (N=26)	0.92 (N=12)	0.56 (N=16)	0.30 (N=23)	0.69 (N=77)

Table 2.11: Percentage of subjects that play X in the game in figure 2.6, by opponents' preference type, restricted to subjects who passed the comprehension task, for subjects who observed both their opponents' risk and ambiguity preference. LRA denotes low risk aversion, HR denotes high risk aversion, AN denotes ambiguity neutral and AA denotes ambiguity averse.

Opponent type Observed preferences	Type 1 or 3 AN	Type 2 or 4 AA	Total
Pr(X)	0.95 (N=22)	0.69 (N=16)	0.84 (N=38)

Table 2.12: Percentage of subjects that play X in the game in figure 2.6, by opponents' preference type, restricted to subjects who passed the comprehension task, for subjects who only observed their opponents' ambiguity preference. AN denotes ambiguity neutral and AA denotes ambiguity averse.

Table 2.13 presents a more detailed analysis of the results. The average effects of observing a risk averse or ambiguity averse opponent are both large and significant although, again, they may be subject to omitted variable bias. In particular, the average effect of observing an ambiguity averse opponent is overstated. The conditional estimates demonstrate that the effect is working chiefly through the risk aversion, rather than ambiguity aversion, channel. The effect of observing a risk averse opponent is large and statistically significant whether we restrict the sample to only ambiguity neutral or ambiguity averse opponents. On the other hand, the effect of observing an ambiguity averse opponent, conditional on the opponent's risk aversion, is not significant. This may not be surprising given the difference in payoff structures between the risk and ambiguity measurement tasks: the risk task measures preferences over prospects with differences in expected values of a few dollars, while the ambiguity task measures preferences over prospects with differences in expected values of a few cents, implying that the signal of opponent's risk preferences is stronger than the signal of opponent's ambiguity preferences.

It could also be the case, however, that subjects are able to identify and reason about their opponent's risk preferences, but are unable to identify or understand the role of ambiguity preferences. Table 2.12 establishes that this is not the case: when subjects observe only their opponent's ambiguity preference they are more likely to play X when facing an ambiguity neutral than ambiguity averse opponent. The effect of ambiguity preferences is approximately half that of risk preferences, and a Fisher exact test with a null hypothesis of no effect of an opponent's ambiguity preferences on behaviour generates a two-sided p-value of 0.06 and a one-sided p-value of 0.04.

Effect of opponents	Restricted to opponents who are	$\Delta Pr(X)$
Ambiguity Aversion		-0.32**
		(0.003)
Risk Aversion		-0.56**
		(0.000)
Ambiguity Aversion	Low Risk Aversion	-0.08
		(0.316)F
Ambiguity Aversion	High Risk Aversion	-0.26
		(0.107)
Risk Aversion	Ambiguity Neutrality	-0.44**
		(0.000)F
Risk Aversion	Ambiguity Aversion	-0.61**
		(0.001)
Ambiguity Aversion & Risk Aversion		-0.70**
		(0.000)

Table 2.13: Change in proportion of subjects playing X in the testing game, as a function of opponent's preferences, restricted to subjects that passed the comprehension tests. $N = 77$. Pearson's χ^2 p-value shown in brackets. p-values denoted with an F are calculated using Fisher's exact test because of small or unbalanced samples. * indicates value is significantly different from 0 using a non-directional test at the 5% level and ** indicates significance at the 1% level.

It is obvious that the results for subjects with Type 1 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 Type 3 opponents also rationalizable by a Nash equilibrium?[28] 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 2.11) forms a Nash equilibrium of the game (where row players mix between A and C) when the row player has a CRRA utility function with risk aversion parameter $\beta = 2.81$. Given that the choice of C in the risk measurement game implies $\beta > 1.93$ the results are consistent with potential equilibrium play at reasonable levels of risk aversion.

Result 2.6. *When a subject observes a signal of her opponent's risk and ambiguity preferences, her behaviour as the column player in the game in figure 2.6 is affected by the signal in a fashion that is consistent with the theory of section 2.1.1.*

2.5 Discussion

The data presented in sections 2.3 and 2.4 provides a clear picture of the role of uncertainty aversion in games. A subject's preferences over risk and ambiguity are correlated with behaviour in our normal form testing game. There is no relationship, however, for subjects that passed our comprehension tests, between a subject's beliefs over her opponent's preferences and her behaviour in the testing game. In a follow up treatment, where subjects were shown their opponent's behaviour in the preference measuring tasks, a very strong relationship between the opponent's risk preferences and the subject's behaviour was observed.

Uncertainty aversion is not the only alternative theory that has been able to explain deviations from Nash equilibrium in previous work. Indeed, other models such as the level-k model, quantal response equilibrium and other regarding preferences can, to various extents, explain part of the data presented in this chapter. There is, however, a key piece of the data that these other models cannot explain: the relationship between ambiguity preferences in the classification game and behaviour in the testing game. One of the key motivations for the experimental design implemented here was to isolate and identify the role of preferences independently of these 'other' behavioural models. It may well be the case that a level-k (or

[28]Because of the very large range of mixed equilibrium probabilities for subjects with Type 2 or Type 4 opponents, there is not much point in repeating the analyses for those subjects.

quantal response, or other regarding preferences) model that is augmented with heterogeneous uncertainty preferences will provide a highly predictive model of strategic behaviour across a wide range of settings but these are not the focus of this chapter. However, for this chapter, we invoke the principle of parsimony and focus on isolating and identifying the role of strategic uncertainty in determining behaviour in strategic settings.

The results presented here are in line with the results found in other recent experimental papers that have, broadly speaking, taken an epistemic approach to behaviour in games. The experimental epistemic games literature seeks to identify the state of knowledge and modes of reasoning that are used by subjects in various experimental environments, with a goal of understanding the process of subject decision making via a convergence of theory and evidence.

Kneeland (2015) finds, in a paper that focuses on level-k reasoning (rather than ambiguity aversion), using a novel ring-game structure, that while most subjects (94%) are rational they do not form consistent conjectures regarding each other's behaviour. Kneeland also finds that a large proportion of subjects believe their opponents are rational (72%) and that 44% have second-order beliefs regarding the rationality of others.

The implications of the results in this chapter for subject rationality are broadly comparable with those found by Kneeland. Amongst ambiguity and risk neutral subjects, we find that 77% of subjects played an undominated strategy in the main testing game, and therefore exhibit rational behaviour. While this is lower than the 94% found by Kneeland, it was easier for Kneeland's subjects to identify dominated strategies because they were dominated by a pure strategy (rather than a mixed strategy as was the case in the game in figure 2.6). We find higher rates of equilibrium play than Kneeland[29], although the game in figure 2.6 is dominance solvable and therefore requires weaker epistemic assumptions to justify equilibrium behaviour than in Kneeland's game. Furthermore, the data in the treatment presented in section 2.4, where subjects were able to observe their opponent's behaviour in the classification games, can be interpreted as providing evidence that most subjects believe that their opponents are rational.

[29] The equilibrium analysis is not included in the final published version of Kneeland (2015), but can be found in earlier working paper versions.

Healy (2013) also provides empirical evidence regarding the epistemic status of laboratory subjects. One of his main finding is that subjects are generally poor at estimating their opponent's preferences over outcomes. Here, we find evidence that subjects are very much aware that they are poor at estimating their opponent's preferences: subjects who observed their opponent's behaviour in the classification games exhibited markedly different behaviour than subjects in the baseline treatment, and subjects responded to information about their opponent's preferences in a highly rational manner.

Esponda and Vespa (2014) find that, when subjects play a voting game against computer agents with known strategies, subjects can correctly extract information from the computers' strategies much more often after observing the computers' votes than they can extract information from a strategically equivalent simultaneous move game. Similarly, our subjects can extract information from their opponent's behaviour in the classification games if they have observed their opponent's behaviour. However, when a subject only predicts, but does not observe, the opponent's behaviour the subject does not behave as if her prediction of the opponent's preferences contains relevant information.

2.5.1 Conclusions

This chapter presents evidence that both risk and ambiguity preferences play a role in determining behaviour in normal form games, but that beliefs over an opponent's preferences do not influence behaviour in games. As demonstrated in a follow up treatment, subjects have a high level of uncertainty regarding their opponent's preferences; when informed of their opponent's preferences, subjects behaved in a fashion that is largely consistent with equilibrium play. The results suggest that both *strategic* uncertainty and uncertainty regarding an opponent's *preferences* play an important role in determining behaviour in normal form games. Additionally, this chapter presents evidence that ambiguity and risk preference are correlated in the context of games.

The results in this chapter, and the previous research discussed above (Kneeland (2015), Healy (2013) and Esponda and Vespa (2014)), provide some strong suggestions for empirically founded models of strategic behaviour.

Generally, assuming first order rationality of subjects appears to be a reasonable assumption (particularly when dominance relationships are strict). However, the standard game theoretic assumption of common knowledge of preferences is soundly rejected by the data: even mutual knowledge of preferences is too strong. Preference heterogeneity (across both risk and ambiguity dimensions) plays an important role in strategic behaviour, but subjects do not seem to have well calibrated beliefs regarding the preferences of others. On aggregate, the evidence suggests that preferences should be modeled as private information with diffuse or uncertain priors. However, the evidence from the follow-up treatment presented in section 2.4 suggests that when subjects are informed of their opponent's preferences (i.e. when beliefs over preferences become 'precise enough') then it is reasonable to assume that subjects will respond to this preference information in a rational fashion. Grant et al. (Forthcoming) provides an important step towards synthesizing some of these ideas in a general theoretical framework.

Furthermore, the results presented here suggest that studies that focus on either ambiguity or risk preferences in strategic environments may overstate the effects of risk or ambiguity preferences on behaviour via an omitted variable bias effect. Unpacking the relationship between preferences and strategic uncertainty requires a careful consideration of both risk and ambiguity aversion.

Chapter 3

Continuity, Inertia and Strategic Uncertainty: A Test of the Theory of Continuous Time Games.

3.1 Introduction

In game theoretic models, players usually make decisions in lock-step at a predetermined set of dates – a timing protocol we will call "Perfectly Discrete time." Most real world interaction, by contrast, unfolds asynchronously in unstructured continuous time, perhaps with some inertia delaying mutual response. Does this difference between typical modeling conventions and real-world interactions matter? Theoretical work on the effects of continuous time environments on behavior (developed especially in Simon and Stinchcombe (1989) and Bergin and MacLeod (1993)) focuses on what we will call "Perfectly Continuous time," a limiting case in which players can respond instantly (that is with zero *inertia*) to one another, and arrives at a surprising answer: Perfectly Discrete time and Perfectly Continuous time can often support fundamentally different equilibria, resulting in wide gaps in behavior between the two settings.

In this chapter we introduce new techniques that allow us to evaluate these theorized gaps in the laboratory directly and assess their relevance for understanding real world behavior. We pose two main questions. First, does the gulf between Perfectly Discrete and Perfectly Continuous time suggested by the theory describe real human behavior? Though equilibria exist that produce large differences in behavior (and authors such as Simon and Stinchcombe (1989) argue that these equilibria should be considered highly focal), multiplicity of equilibrium in Perfectly Continuous time means that the effect of continuous time is, ultimately, theoretically indeterminate. Second, we ask how empirically relevant these gaps are: can more realistic, imperfectly continuous time games (games with natural response delays that we call "Inertial Continuous time" games) generate Perfectly Continuous-like outcomes? Nash equilibrium suggests not but, as Simon and Stinchcombe (1989) and Bergin and MacLeod (1993) emphasize, even slight deviations from Nash equilibrium assumptions (á la ε-equilibrium) allow Perfectly Continuous-like behavior to survive as equilibria in the face of inertia, provided inertia is sufficiently small.

Recent experiments have begun to investigate the relationship between continuous and discrete time behavior in the lab (e.g. Friedman and Oprea (2012) and Bigoni et al. (2015))[1] but have not yet directly tested the theory motivating these questions for a simple reason: natural human reaction lags in continuous time settings generate inertia that prevents a direct implementation of the premises of the theory. These Inertial Continuous time settings are empirically important (and of independent interest) but are insufficient for a direct theory test because they generate very different equilibrium behavior from the Perfectly Continuous time environments that anchor the theory (a prediction we test and find strong though highly qualified support for in our data). In our experimental design, we introduce a new protocol ("freeze time") that eliminates inertia by pausing the game for several seconds after subjects make decisions, allowing them to respond "instantly" to actions

[1] Both Friedman and Oprea (2012) and Bigoni et al. (2015) report evidence from prisoner's dilemmas played with flow payoffs in Inertial Continuous time (i.e. subjects in these experiments suffer natural reaction lags that prevent instant response to the actions of others). While the Friedman and Oprea (2012) design varies the continuity of the environment (discrete vs. continuous time interaction) in deterministic horizon games, the Bigoni et al. (2015) design centers on varying the stochasticity of the horizon (deterministic vs. stochastic horizon) in continuous time games. Other more distantly related continuous time papers include experimental work on multi-player centipede games (Murphy et al., 2006), public-goods games (Oprea et al., 2014), network games (Berninghaus et al., 2006), minimum-effort games (Deck and Nikiforakis, 2012), hawk-dove games (Oprea et al., 2011), bargaining games (Embrey et al., 2015) and the effects of public signals (Evdokimov and Rahman, 2014).

of others (i.e. with no lag in game time) and thus allowing us to test Perfectly Continuous predictions. By systematically comparing behavior in this Perfectly Continuous setting to both Perfectly Discrete time and Inertial Continuous time settings we are able to pose and answer our motivating questions.

We apply this new methodology to a simple timing game similar to one discussed in Simon and Stinchcombe (1989) that is ideally suited for a careful test of the theory.[2] In this game, each of two agents decides independently when to enter a market. Joint delay is mutually beneficial up to a point, but agents benefit from preempting their counterparts (and suffer by being preempted). In Perfectly Discrete time, agents will enter the market at the very first opportunity, sacrificing significant potential profits in subgame perfect equilibrium. By contrast, in Perfectly Continuous time, agents can, in equilibrium, delay entry until 40% of the game has elapsed, thereby maximizing joint profits.[3] (Simon and Stinchcombe (1989) emphasize this equilibrium and point out that it uniquely survives iterated elimination of weakly dominated strategies, but many other equilibria – including inefficient immediate-entry equilibrium – exist in Perfectly Continuous time.) Importantly, Inertial Continuous time of the sort studied in previous experiments leads not to Perfectly Continuous time-like multiplicity in equilibrium but only to the inefficient instant entry predicted for Perfectly Discrete time – as Bergin and MacLeod (1993) point out even a small amount of inertia theoretically erases all of the efficiency enhancing potential of continuous time in Nash equilibrium.[4]

In the first part of our experimental design we pose our main question by comparing Perfectly Discrete and Perfectly Continuous time using a baseline set of

[2]Compared to, for instance, the continuously repeated prisoner's dilemma, our timing game has several advantages for a diagnostic test. First, the joint profit maximizing outcome predicted by Simon and Stinchcombe (1989) is interior, meaning simple heuristics like "cooperate until the end of the game" cannot be confused with equilibrium play. Second, the strategy space is considerably simpler than the prisoner's dilemma, making measurement of decisions and inferences about strategies crisper. Finally, the prisoner's dilemma frames the contrast between cooperation and defection somewhat starkly and may therefore trigger social behaviors that have little to do with the forces we designed our experiment to study – we speculated in designing the experiment that our timing game would be somewhat cleaner from this perspective.

[3]More precisely, the agents are maximizing joint profits subject to playing non-strictly dominated strategies in every subgame that is reached on the path of play.

[4]We note that all equilibria of the game we study in this chapter are subgame perfect and that some of our proofs rely on backwards induction. In the remainder of the chapter, for readibility, we omit the modifier "subgame perfect" when discussing equilibria of our game.

parameters and 60 second runs of the game. In the Perfectly Discrete time protocol, we divide the 60 second game into 16 discrete grid points and allow subjects to simultaneously choose at each grid point whether to enter the market. In the Perfectly Continuous time protocol we instead allow subjects to enter at any moment but, crucially, eliminate natural human inertia by freezing the game after any player enters, allowing her counterpart to enter "immediately" from a game-time perspective if she enters during the ample window of the freeze. We find evidence of large and extremely consistent differences in behavior across these two protocols. Virtually all subjects in the Perfectly Discrete time treatment suboptimally enter at the first possible moment while virtually all subjects in Perfectly Continuous time enter 40% of the way into the period, forming a tight mode around the joint profit maximizing entry time. The results thus support the conjecture of a large - indeed, from a payoff perspective, maximally large – gap between Perfectly Continuous and Discrete time behaviors.

In the second part of the design, we study how introducing realistic inertia into continuous time interaction changes the nature of the results observed in our Perfectly Continuous time treatment. Though Nash equilibrium predicts that even a tiny amount of inertia will force behavior back to Perfectly Discrete-like immediate entry times, alternatives such as ε-equilibrium suggest that Perfectly Continuous-like results may survive as equilibria at low levels of inertia. In Inertial Continuous time treatments we replicate our Perfectly Continuous time treatment but remove the freeze time protocol, thereby allowing natural human reaction lags to produce a natural source of inertia. We systematically vary the severity of this inertia by varying the speed of the game relative to subjects' natural reaction lags and find that when inertia is highest, entry times collapse to zero in continuous time as predicted by Nash equilibrium. However, when we lower inertia to sufficiently small levels, we observe large entry delays that are nearly as efficient as those observed in Perfectly Continuous time. Thus, realistic Inertial Continuous Time behavior is well approximated by the extreme of Perfectly Discrete time when inertia is large and better approximated by the extreme of Perfectly Continuous time when inertia is small. While these patterns are inconsistent with Nash equilibrium, they are, as both Simon and Stinchcombe (1989) and Bergin and MacLeod (1993) stress, broadly consistent with ε-equilibrium.

Though ε-equilibrium is consistent with our data, it also generates imprecise predictions. In the final part of the chapter, we consider, *ex post*, explanations that can more sharply organize our data in order to better understand how inertia and continuity influence behavior. Recent experimental work on dynamic strategic interaction (e.g. Dal Bo and Frechette (2011), Embrey et al. (2016), Vespa and Wilson (2016),Dal Bo and Frechette (2016)) has emphasized the crucial role strategic uncertainty (assumed away in Nash equilibrium) plays in predicting both equilibrium selection and non-Nash equilibrium behavior, and has focused especially on the predictive power of the basin of attraction of defection relative to cooperative alternatives. We show that the basin of attraction becomes more hospitable to continued cooperation at each moment in time as inertia falls towards the continuous limit and that measures of risk dominance starkly organize our data, tying our results directly in with this recent work.

We then consider more explicit ways of modeling subjects' responses to strategic uncertainty by studying simple (and parsimonious) heuristic rules (drawn from Milnor, 1954) discussed in the literature on highly uncertain environments: maximin ambiguity aversion (MAA), minimax regret avoidance (MRA) and Laplacian subjective expected utility (LEU). Of these, we find that the MRA decision rule dramatically outperforms Nash equilibrium, making extremely accurate point predictions for our game. We show that MRA also organizes data in an additional pair of diagnostic treatments that generate comparative statics unanticipated by Nash equilibrium. Finally, we consider the relevance of this result to other common continuous time games. In the online appendix we show theoretically that MRA predicts a smooth approach to Pareto optimal, Perfectly Continuous-like play as inertia falls to zero in an important, broad class of continuous time games under empirically sensible restrictions on beliefs. We directly test this claim by showing that MRA predictions almost exactly track data on behavior in continuous time prisoner's dilemma from previous work. Our analysis thus suggests that heuristic responses to strategic uncertainty like MRA may be a productive way of organizing and interpreting data in a wide range of dynamic strategic settings.

The results of our experiment suggest a role for Perfectly Continuous time theoretical benchmarks in predicting and interpreting real-world behavior, even if the world is never perfectly continuous. Changes in technology have recently narrowed

– and continue to narrow – the gap between many types of human interactions and the Perfectly Continuous setting described in the theory. Constant mobile access to markets and social networks, the proliferation of applications that speed up search and the advent of automated agents deployed for trade and search have the effect of reducing inertia in human interactions. Our results suggest, contra Nash equilibrium, that such movements towards continuity may generate some of the dramatic effects on behavior predicted for (and observed in) Perfectly Continuous time even if inertia never falls quite to zero. Guided by these results, we conjecture that the share of interactions that are better understood through the theoretical lens of Perfectly Continuous time than that of Perfectly Discrete time will grow as social and economic activity continues to be transformed by this sort of technological change.

The remainder of the chapter is organized as follows. Section 3.2 gives an overview of the main relevant theoretical results that form hypotheses for our experiment and section 3.3 describes the experimental design. Section 3.4 presents our results, Section 3.5 interprets the results in light of metrics and models of strategic uncertainty, and Section 3.6 concludes the chapter. Appendices collect theoretical proofs and the instructions to subjects.

3.2 Theoretical Background and Hypotheses

In section 3.2.1 we introduce our timing game and in section 3.2.2 we state and discuss a set of propositions characterizing Nash equilibrium and providing us with our main hypotheses. In section 3.2.3 we consider alternative hypotheses motivated by ε-equilibrium.

3.2.1 A diagnostic timing game

Consider the following simple timing game, adapted from one described in Simon and Stinchcombe (1989). Two firms, $i \in \{a,b\}$, each choose a time $t_i \in [0,1]$ at which to enter a market, perhaps conditioning this choice on the history of the game.[5] Payoffs depend on the order of entry according to the following symmetric

[5] To conserve notation, we normalize the length of the game to be 1 for the theoretical analysis. In our experiment, we sometimes vary the length of the game (and with it the severity of inertia and predicted time of entry) across treatments.

function:

$$U_a(t_a,t_b) = \begin{cases} \frac{1-t_b}{2}\left[\Pi_D + (t_b - t_a)(1+\frac{2}{1-t_b})\Pi_F\right] - c(1-t_a)^2 & \text{if } t_a < t_b \\ \frac{1-t_a}{2}\Pi_D - c(1-t_a)^2 & \text{if } t_a = t_b \\ \frac{1-t_a}{2}\left[\Pi_D - (t_a - t_b)\Pi_S\right] - c(1-t_a)^2 & \text{if } t_b < t_a \end{cases} \quad (3.1)$$

with parameters assumed to satisfy $0 < 2c < \Pi_S \leq \Pi_D < 4c$ and $\frac{4c}{3} \leq \Pi_F \leq 4c$. Though the applied setting modeled by this sort of game matters little for our relatively abstract experiment, we can interpret the model as one in which firms face quadratic costs for time spent in the market (parameterized by c), earn a duopoly flow profit rate of Π_D while sharing the market, earn a greater flow profit Π_F while a monopolist and suffer a permanent reduced earnings rate (parameterized by Π_S) proportional to the time one's counterpart has spent as a monopolist.

Several characteristics of this game are particularly important for what follows. First, firms earn identical profits if they enter at the same time and this simultaneous entry payoff is strictly concave in entry time, reaching a maximum at a time $t^* = 1 - \frac{\Pi_D}{4c} \in (0, \frac{1}{2})$. Second, if one of the firms instead enters earlier than the other (at time t'), she earns a higher payoff and her counterpart a lower payoff than had they entered simultaneously at time t'. The firms thus maximize joint earnings by delaying entry until an interior time t^* but at each moment each firm has a motivation to preempt its counterpart and to avoid being preempted.

3.2.2 Discrete, inertial and perfectly continuous time predictions

What entry times can be supported as equilibria in this game? The key observation motivating both the theory and our experiment is that the answer depends on how time operates in the game. In this subsection we characterize equilibrium under three distinct protocols: Perfectly Discrete time, Perfectly Continuous time and Inertial Continuous time (here we only sketch the main conceptual issues, deferring technical discussion to Appendix A).

We begin with Perfectly Discrete time, the simplest and most familiar case. Here, time is divided into $G+1$ evenly spaced grid points (starting always at $t = 0$)

on $[0,1]$ and players make simultaneous decisions at each of these points. More precisely, each player chooses a time $t \in \{0, 1/G, ..., (G-1)/G, 1\}$ at which to enter, possibly conditioning this choice on the history of the game, H_t at each grid point. Earnings are given by equation 3.1 applied to the dates on the grid at which entry occurred.[6] As in familiar dynamic discrete time games like the centipede game and the finitely repeated prisoner's dilemma there is a tension here between efficiency (which requires mutual delay until at least the grid point immediately prior to t^*) and individual sequential rationality (which encourages a player to preempt her counterpart). Applying the logic of backwards induction, strategies that delay entry past the first grid point unravel, leaving immediate entry at the first grid point, $t = 0$, as the unique subgame perfect equilibrium, regardless of G.

Proposition 3.1. *In Perfectly Discrete time, the unique subgame perfect equilibrium is for both firms to enter at time 0, regardless of the fineness of the grid, G.*

Proof. : See Appendix C.2.1.

At the opposite extreme, in Perfectly Continuous time players are not confined to a grid of entry times but can instead enter at any moment $t_i \in [0,1]$ (again, possibly conditioning on the history of the game at each t, H_t). Simon and Stinchcombe (1989) emphasize the relationship between the two extremes, modeling Perfectly Continuous Time as the limit of a Perfectly Discrete time game as G approaches infinity. In this limit, players can respond instantly to entry choices made by others: if an agent enters the market at time t her counterpart can respond by also entering at t, moving in response to her counterpart but at identical dates. Since, in our game, delaying entry after a counterpart enters is strictly payoff decreasing, no player can expect to succeed in preempting her counterpart (or have reason to fear being preempted). This elimination of preemption motives also protects efficient delayed entry from unravelling and thus makes it possible to support any entry time $t \in [0, t^*]$ as an equilibrium.[7]

[6] For example, if firm a entered at the third grid point, and firm b entered at the fifth grid point, the payoff for firm a is given by $U_a(\frac{2}{G}, \frac{4}{G})$.

[7] Entry times greater than t^* cannot be supported in equilibrium because they are always payoff dominated by t^*. Notice that despite the symmetry of the (joint entry) payoff function around t^*, the

Proposition 3.2. *In Perfectly Continuous time, any entry time $t \in [0, t^*]$ can be supported as a subgame perfect equilibrium outcome.*

Proof. We provide three proofs of this proposition. Appendix C.2.1 includes a self-contained heuristic proof, a more formal proof that draws directly from Simon and Stinchcombe (1989) and an alternative proof that instead follows the modeling approach of Bergin and MacLeod (1993).

Though it is possible for Perfectly Discrete and Perfectly Continuous behaviors to radically differ in equilibrium, this is hardly guaranteed. Because of multiplicity, Perfectly Continuous behavior may be quite different or quite similar to Perfectly Discrete behavior in equilibrium ($t = 0$ and t^* are both supportable in equilibrium in Perfectly Continuous time) depending on the principle of equilibrium selection at work. This multiplicity is in fact a central motivation for studying these environments in the laboratory. Simon and Stinchcombe (1989) emphasize that t^* is the unique entry time to survive iterated elimination of weakly dominated strategies in our game and they argue that this refinement is natural in the context of Perfectly Continuous time games. Evaluating the organizing power of this refinement is another central motivation for our study.

Remark 3.1. In Perfectly Continuous time, joint entry at $t^* = 1 - \frac{\Pi_D}{4c}$ is the only outcome that survives iterated elimination of weakly dominated strategies.

Proof. A heuristic proof is provided in appendix C.2.1. For further details, see Simon and Stinchcombe (1989).

Finally, Inertial Continuous time lies between the extremes of Perfectly Discrete and Perfectly Continuous time, featuring characteristics of each. Here, as in Perfectly Continuous time, players can make asynchronous decisions and are not confined to entering at a predetermined grid of times. However, as in Perfectly Discrete time, players are unable to respond instantly to entry decisions by their counterparts. In Inertial Continuous time, inability to instantly respond is due to what Bergin and MacLeod (1993) call *inertia* (here, simply response lags of ex-

equilibrium entry set is not symmetric around t^* because of the temporal nature of the game.

ogenous size δ).[8, 9] With inertial reaction lags, the logic of unravelling returns as players once again have motives to preempt one another. As a result, the efficient delayed entry supported in equilibrium in Perfectly Continuous time evaporates with even an arbitrarily small amount of inertia. Theoretically then, even a tiny amount of inertia pushes continuous time behavior to Perfectly Discrete levels.

Proposition 3.3. *In Inertial Continuous time, only entry at time 0 can be supported as a subgame perfect equilibrium regardless of the size of inertia,* $\delta > 0$*.*

Proof. See Appendix C.2.1.

Instead of modeling Perfectly Continuous time as a limit of Perfectly Discrete time as the grid becomes arbitrarily fine as Simon and Stinchcombe (1989) do, Bergin and MacLeod (1993) model it as the limit of Inertial Continuous time as inertia approaches zero. This alternative method for defining Perfectly Continuous time leads to an identical equilibrium set to the one described by Simon and Stinchcombe (1989) for our game.

3.2.3 Alternative hypothesis: Inertia and ε-equilibrium

Continuous time can fundamentally change Nash equilibrium behavior but this effect is extremely fragile: even a slight amount of inertia will eliminate any pro-cooperative effects of continuous time interaction in games like ours. Since inertia is realistic, this frailty in turn calls into question the usefulness of the theory for predicting and interpreting behavior in the real-world. Perhaps for this reason both Simon and Stinchcombe (1989) and Bergin and MacLeod (1993) motivate the theory of continuous time explicitly with reference to the more forgiving alternative of ε-equilibrium, emphasizing that any Perfectly Continuous time Nash equilibrium is arbitrarily close to some ε-equilibrium of a continuous time game with inertia (and vice versa).[10] If agents are willing to tolerate even very small deviations from

[8] Though, in the context of our experiment, inertia simply refers to natural human reaction lags, Bergin and MacLeod (1993) point out that more general types of inertia are possible.

[9] Throughout the chapter we define inertia δ as the ratio of an agent's reaction lag, δ_0, to the total length of the game, T (i.e. $\delta \equiv \delta_0/T$). Inertia is thus the fraction of the game that elapses before an agent can respond to her counterpart.

[10] More precisely, Simon and Stinchcombe (1989) make this same point with respect to synchronous, discrete time games with very fine time grids.

best response, they can support Perfectly Continuous-like outcomes as equilibria even in the face of inertia, provided inertia is sufficiently small. Indeed, in our game, when inertia is large, ε-equilibrium coincides perfectly with Nash equilibrium, supporting only immediate entry and mirroring Perfectly Discrete time Nash equilibrium. However when inertia falls below a threshold level (determined by ε) the equilibrium set expands to support any entry time $t \in [0, t^*]$, instead mirroring Perfectly Continuous time Nash equilibrium. We formalize this in the following proposition:

Proposition 3.4. *Consider a game in inertial continuous time. For any $\varepsilon > 0$, there exists $\hat{\delta} > 0$ such that for all levels of inertia $\delta < \hat{\delta}$ any entry time in $[0, t^*]$ can be sustained in a subgame perfect ε-equilibrium.*

For any $0 < \varepsilon < \Pi_F - c$ there exists $\underline{\delta}$ such that for all levels of inertia $\delta > \underline{\delta}$ immediate entry is the unique subgame perfect ε-equilibrium.

Proof. See Appendix C.3.2.[11]

This result is useful because it emphasizes that even very small deviations from the assumptions underlying Nash equilibrium – for instance small amounts of noise in beliefs or imprecision in payoffs specified in the game – can make either Perfectly Discrete or Perfectly Continuous time benchmarks more predictive, depending on the severity of inertia.[12] For this reason (and because of the important role ε-equilibrium plays in the theory), we built our experimental design in part with this alternative prediction in mind as an *ex ante* alternative hypothesis to Nash equilibrium. In section 3.5, we consider, *ex post*, more specific interpretations for the sort of non-Nash equilibrium behavior ε-equilibrium is capable of sustaining, focusing on the role of strategic uncertainty in supporting non-Nash equilibrium outcomes. By doing so we are able to more sharply organize our results and tie our findings to important themes explored in recent, closely related literature on dynamic strategic interaction.

[11] In Appendix C.3 we prove a set of propositions fully characterizing the ε-equilibrium sets for our protocols.

[12] Generically, $\hat{\delta} < \underline{\delta}$. However, when $\Pi_F = 4c$, as in our main treatments, $\hat{\delta} = \underline{\delta}$ which implies a discontinuity in the equilibrium set. See proposition C.4 for further details regarding the continuity (or otherwise) of the equilibrium set.

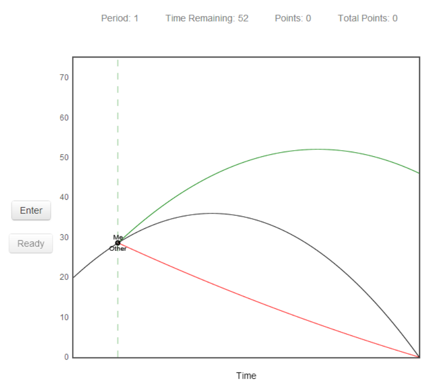

Figure 3.1: Screen shot from Perfectly Continuous and Inertial Continuous time treatments (under the low-temptation parametrization).

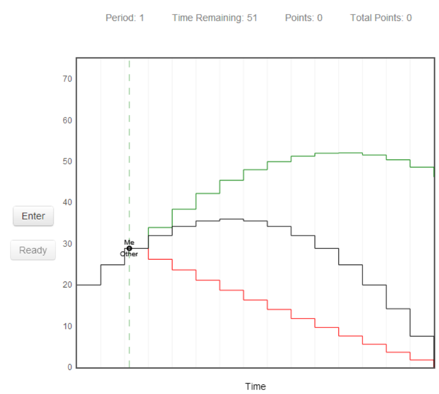

Figure 3.2: Screen shot from Perfectly Discrete time treatments (under the low-temptation parametrization).

3.3 Design and Implementation

In section 3.3.1 we discuss our strategy for implementing our three timing protocols in the lab and present the experimental software we built to carry out this strategy. In section 3.3.2 we present our treatment design.

3.3.1 Timing protocols and experimental software

We ran our experiment using a custom piece of software programmed in Redwood (Pettit et al. (2015)). Figures 3.1 and 3.2 show screenshots. Using this software, we implemented the three timing protocols described in Section 3.2 as follows:

Inertial Continuous time. Figure 3.1 shows an Inertial Continuous time screenshot. As time elapses during the period, the payoff dots (labeled "Me" and "Other") move along the joint payoff line (black center line) from the left to the right of the screen. (In most treatments periods last 60 seconds meaning it takes 60 seconds for the payoff dot to reach the right hand side of the screen.) When a subject is the first player to enter the market, her payoff dot shifts from the black to the green line (the top line), while her counterpart's payoff dot (the dot of the second mover) shifts to the red line (the bottom line).[13] When the second player enters the market, period

[13]Because of the nature of the payoff function, the green and red line change throughout the period

payoffs for both players are determined by the vertical location of each player's dot at the moment of second entry. Once both players have entered, they wait until the remaining time in the period has elapsed before the next period begins (see the instructions in Appendix B for more detail). Because subjects, on average, take roughly 0.5 seconds to respond to actions by others, subjects have natural inertia in their decision making that should theoretically generate Inertial Continuous time equilibrium behavior.

Perfectly Continuous time. The Perfectly Continuous time implementation is identical to the Inertial Continuous time implementation (as shown in Figure 3.1) except that when either subject presses the space bar to enter the game freezes (we call this the "Freeze Time" protocol) and the payoff dots stop moving from left to right across the screen for five seconds. Subjects observe a countdown on the screen and the first mover's counterpart is allowed to enter during this time. If the counterpart enters during this window, her response is treated as simultaneous to her counterpart's entry time and both players earn the amount given by the current vertical location of their payoff dot. Otherwise, the game continues as in Inertial Continuous time once the window has expired. Regardless, subjects must wait until the remaining time in the period has elapsed before the next period begins. The length of the pause was calibrated to be roughly 10 times longer than the median reaction lag measured in Inertial Continuous time, giving subjects ample time to respond, driving inertia to 0 and thus satisfying the premises of Perfectly Continuous time models.[14]

Perfectly Discrete time. Figure 3.2 shows a screen shot for the Perfectly Discrete treatments, which is very similar to the continuous time screen but for a few

prior to entry and stabilize once one player has entered.

[14]An alternative way of formally implementing Perfectly Continuous time is to allow subjects to pre-specify an entry time and a stationary response delay (possibly set to zero). We opted to use the Freeze Time protocol instead of this sort of strategy method for three reasons. First, employing the strategy method would force us to substantially constrain subjects' strategy space, eliminating or limiting the dependence of strategies on histories. Second, using the "Freeze Time" protocol allows us to directly compare entry decisions to inertia generated by naturally occurring reaction lags in Inertial Continuous time – a central goal of the experiment that would be impossible using the strategy method. Finally, for realism, we wanted subjects to actually see the unfolding of payoffs and behavior in real time. Nonetheless, see Duffy and Ochs (2012) for evidence from a distantly related entry game that simultaneous choice and dynamic implementations can generate very similar results.

changes. First, periods are divided into $G = 15$ subperiods, which begin at grid-points $t = \{0, 4, 8, ..., 56\}$ (measured in seconds)[15], each marked by a vertical gray line on the subject's screen. Instead of moving smoothly through time, as in the continuous time treatments, the payoff dots follow step functions and "jump" to the next step on the payoff functions at the end of each subperiod. Actions are shrouded during a subperiod, so payoff dots will only move from the black to the green (or red) payoff lines after the subperiod in which a subject chose to enter has ended. Payoffs are determined according to equation (3.1), calculated at the grid point that began the subperiod in which the subject entered.[16]

3.3.2 Treatment design and implementation

Our experimental design has three parts. In the first we implement 60 second timing games using the parameter vector $(c, \Pi_D, \Pi_F, \Pi_S) = (1, 2.4, 4, 2.16)$ under the extremes of Perfectly Discrete and Perfectly Continuous time.[17] We call these Baseline treatments PD (Perfectly Discrete) and PC (Perfectly Continuous).

Second, we examine the effects of inertia on continuous-time decisions by running a series of Inertial Continuous time treatments using the same Baseline parameters. In the IC_{60} treatment we run periods lasting 60 seconds each (just as in the PC and PD treatments). In the IC_{10} and IC_{280} treatments we repeat the IC_{60} treatment but speed up or slow down the clock so that periods finish in 10 or 280 seconds (respectively). By speeding up the game clock so that the game lasts only 10 seconds (the IC_{10} treatment) we dramatically increase the severity of inertia; by slowing down the game so that it takes 280 seconds to finish (the IC_{280} treatment),

[15] We used the convention that any subjects who were yet to enter at the $t = 56$ subperiod will be forced to enter at the $t = 60$ subperiod, which would result in a payoff of 0. In practice, however, no subjects came close to waiting this long to enter.

[16] For example, if in our Baseline treatment a subject entered in the first subperiod and her counterpart entered in the third subperiod, payoffs would be given by $U(0, \frac{2}{15})$ for the subject and $U(\frac{2}{15}, 0)$ for his counterpart.

[17] To allow the entire payoff space to be shown on a single reasonably scaled plot we truncated the maximum payment to be 75 points per period (for context $U(t^*, t^*)$ was normalized to be 36 points). This truncation (which subjects can clearly see on their screen) only affects the payoff of the first mover under the unusual circumstance that her opponent delays entry for a significant amount of time. Regardless, this design choice only affects payoffs that are well off the equilibrium path and does not affect any of the equilibrium sets discussed in the chapter.

we substantially reduce the severity of inertia.[18]

Finally, in the Low Temptation treatments (discussed in Section 3.5.2), we examine the robustness of explanations for our main results by changing the payoff functions in Perfectly Continuous and Discrete time. In the L-PD (Low temptation - Perfectly Discrete) and L-PC (Low temptation - Perfectly Continuous) treatments we replicate the PD and PC treatments but lower the premium from preempting one's counterpart, Π_F, from 4 to 1.4. Changing Π_F has no effect on Nash equilibrium in either case but can have substantial effects in Perfectly Discrete time under some alternative theories.

All treatments are parameterized such that t^* occurs 40% of the way into the period (the 7th subperiod in Perfectly Discrete time treatments).[19]

We ran the PD, IC, PC, L-PD and L-PC treatments using a completely between-subjects design. In each case we ran 4 sessions with between 8 and 12 subjects participating. Each session was divided into 30 periods, each a complete run of the 60 second game, and subjects were randomly and anonymously matched and rematched into new pairs at the beginning of each period. We ran the IC_{10} and IC_{280} treatments using a within-subject design consisting of 3 blocks each composed of 3 IC_{280} periods followed by 7 IC_{10} periods, for a total of 30 periods.[20] Once again, subjects were randomly and anonymously rematched into new pairs each period.

We conducted all sessions at the University of British Columbia in the Vancouver School of Economics' ELVSE lab between March and May 2014. We randomly invited undergraduate subjects to the lab via ORSEE (Greiner (2004)), assigned them to seats, read instructions (reproduced in Appendix B) out loud and gave a brief demonstration of the software. In total 274 subjects participated, were paid based on their accumulated earnings and, on average, earned \$26.68 (includ-

[18] Recall that we define inertia δ as the ratio of an agent's reaction lag, δ_0, to the total length of the game, T (i.e. $\delta \equiv \delta_0/T$).

[19] In 60 second period treatments this occurs 24 seconds into the period while in the IC_{10} and IC_{280} treatments this occurs after 4 or 112 seconds respectively.

[20] We used a within design for these treatments mostly because we were concerned that the extreme duration of IC_{280} periods would cause boredom in subjects if repeated a number of times. By interspersing these with fast-paced IC_{10} we were able to reduce this concern. During the experiment, we revealed the next period's treatment (IC_{280} or IC_{10}) only after the conclusion of the previous period.

ing a $5 show up payment).[21] Sessions (including instructions, demonstrations and payments) lasted between 60 and 90 minutes.

3.4 Results

In section 3.4.1 we report results from the Baseline treatments, comparing Perfectly Continuous and Perfectly Discrete time behaviors under identical parameters. The data strongly supports Simon and Stinchcombe (1989)'s conjecture of a large gap between Perfectly Continuous and Discrete time: PD subjects nearly always inefficiently enter immediately while PC subjects nearly always delay entry until t^*, the joint profit maximizing entry time. In section 3.4.2, we study the relationship between the relatively realistic setting of Inertial Continuous time and the extremes of Perfectly Continuous and Discrete time, varying the severity of subjects' inertia by varying the speed of an Inertial Continuous time game under Baseline parameters. We find that at high levels of inertia, behavior follows Nash equilibrium predictions, collapsing to Perfectly Discrete levels. As inertia drops towards zero, however, entry times approach Perfectly Continuous time levels, a result inconsistent with Nash equilibrium.

Most of the distinctive predictions and comparative statics discussed in Section 3.2 concern the timing of *first entry* and documenting first entry times will be our primary focus in the data analysis. Before turning to this data, however, it is useful to briefly document second mover behavior across treatments. Focusing attention on behavior after the first 10% of periods (after subjects have had a few periods to become comfortable with the interface), we find quite uniform and sensible behavior across treatments: subjects almost universally enter as soon as possible (given inertia or discretization) following a counterpart's entry, meaning that subjects constrain themselves to playing *admissible strategies* (strategies that are weakly undominated, see Brandenburger et al. (2008)). In both Perfectly Continuous and Perfectly Discrete time, over 95% of second movers enter at the first possible opportunity after their first-moving counterparts (immediately in PC and no later than the very next sub-period in PD).[22] In Inertial treatments we measure

[21] Funds for subject payments were provided by a research grant from the Faculty of Arts at the University of British Columbia.

[22] About 5% of subjects in the PC protocol entered with a delay of exactly 0.1 seconds, which we

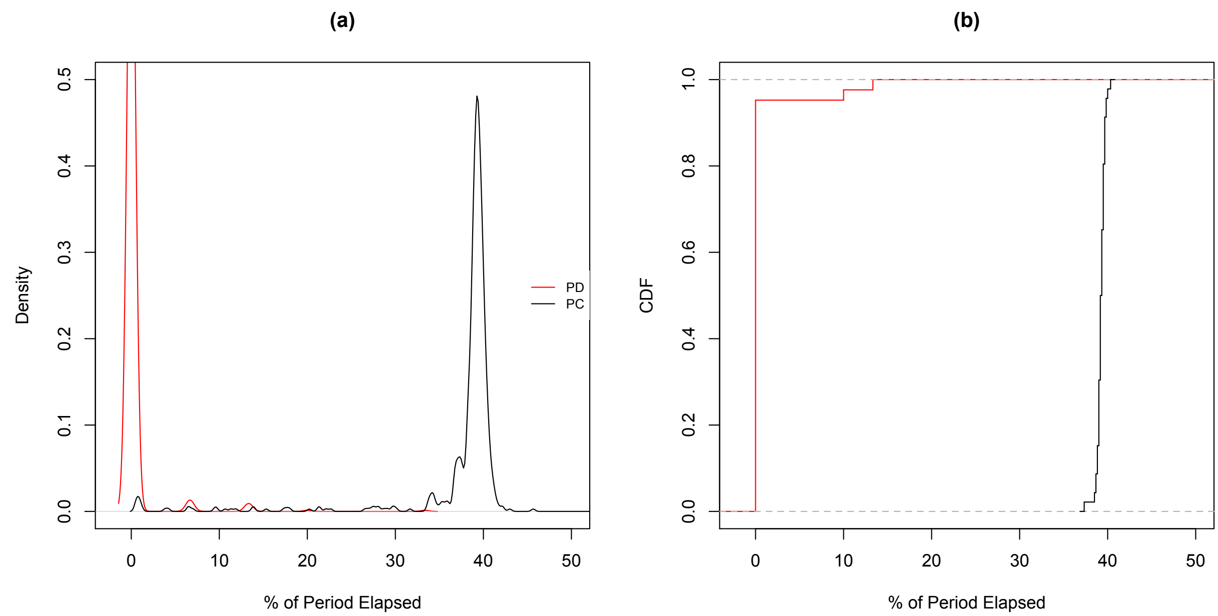

Figure 3.3: (a) The left hand panel shows kernel density estimates of entry times (normalized as fraction of the period elapsed) in the PD and PC treatments. For both treatments t^*, which generates the maximal symmetric payoff, lies 40% of the way into the period. (b) The right hand panel shows CDFs of subject-wise medians from product limit estimates (Kaplan and Meier (1958)) of intended entry times.

the median subject's reaction lag, δ_0, at 0.5 seconds, closely matching reaction lags documented in previous research (e.g. Friedman and Oprea, 2012). Given the incentives in our game, these rapid responses strongly suggest that subjects understood the structure and incentives of the game across treatments (as delay in response is strictly dominated in each treatment in the experiment). Unless otherwise noted, remaining references to entry times will refer to the timing of first entry.

3.4.1 Perfectly continuous and discrete time

Figure 3.3 (a) plots kernel density estimates of observed entry times for our PD (in red) and PC treatments (in black). Figure 3.3 (b) complements the kernel density estimates by plotting CDFs of subject-wise median entry times using product limit estimation intended to minimize the potential downward bias introduced by first movers preempting – and therefore censoring – the intended entry times of second movers.[23]

The results are striking. In the PD treatment, virtually all subjects choose to enter immediately as the theory predicts, generating highly inefficient outcomes. The PC treatment, by contrast, induces radically different[24] behavior: entry times are tightly clustered near t^*, with subjects maximizing joint earnings by delaying entry until about 40% of the period has elapsed. Recall that though t^* is only one of a continuum of equilibria in PC, it is the outcome uniquely selected by elimination of weakly dominated strategy and is advanced as a focal prediction by Simon and Stinchcombe (1989). The tightly clustered behavior in the PC treatments supports this conjectured focality and suggests that equilibrium selection is very uniform in Perfectly Continuous time. This pattern of behavior thus strongly supports the conjecture that Perfectly Discrete and Perfectly Continuous time induce fundamentally different behaviors in otherwise identical games.

Result 3.1. *Under Baseline parameters, Perfectly Continuous interaction induces fundamentally different behavior from Perfectly Discrete interaction. While subjects virtually always enter immediately in the PD treatment, they virtually always delay entry until t^* in the PC treatment.*

---

believe is due to a rounding error by the software and which we treat as a zero second lag in this calculation.

[23]Specifically we use techniques introduced by Kaplan and Meier (1958) to calculate non-parametric, maximum likelihood estimates of each subject's distribution of intended entry times in the face of censoring bias introduced by counterpart preemption. The procedure uses observed entry times to partially correct censoring bias introduced in periods in which the subject is preempted by her counterpart. For each subject we estimate these distribution functions and then take the median. Figure 3.3 (b) plots distributions (across subjects) of these medians.

[24]Mann-Whitney tests on session-wise median product-limit estimates of entry times allows us to reject the hypothesis that PC and PD distributions are the same at the five percent level.

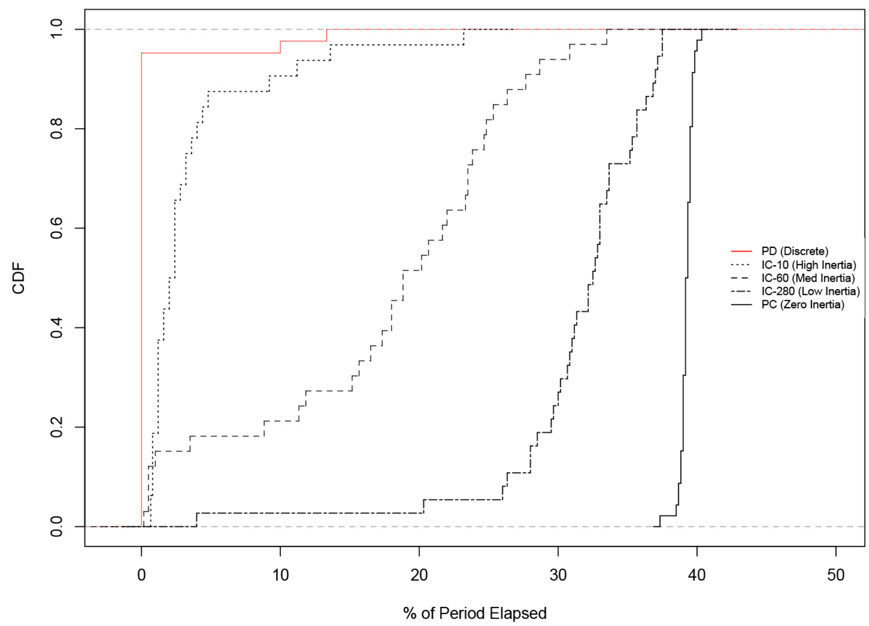

Figure 3.4: CDFs of subject-wise product limit estimates of entry times in each of the main treatments of the experiment.

3.4.2 Inertia and continuous time

Perfectly Continuous time generates a dramatic change in behavior, but environments with zero inertia are probably rare in the real world. How robust are these extreme results to a re-introduction of inertia into the game? In order to study this question we ran a series of Inertial Continuous time (IC) treatments, varying the severity of inertia from very high to very low. In the IC_{60} treatment we duplicated the PC treatment but eliminated the freeze time protocol, allowing subjects' reaction lags to generate natural inertia in the game. In the IC_{10} and IC_{280} treatments, run within-subject, we sped up (IC_{10}) or slowed down (IC_{280}) the game clock relative to the 60 second IC_{60} periods, generating periods that lasted 10 or 280 seconds respectively. Speeding up the game dramatically increases the magnitude of inertia (defined, recall, as the ratio of reaction lags to game length) while slowing down the game reduces inertia substantially.

Figure 3.4 shows the results, plotting CDFs of subject-wise median product-limit estimates of entry times for the IC_{10} (high inertia), IC_{60} (moderate inertia), IC_{280} (low inertia) and the *PC* (zero inertia) treatments (for reference we also plot the *PD* treatment in red). The results reveal dramatic and quite systematic effects of inertia on continuous time behavior as inertia drops towards zero. First, the tight optimal entry delays observed in the *PC* treatment almost completely collapse in the high inertia case, generating Perfectly Discrete-like near-immediate entry as predicted by Nash equilibrium. However, when we reduce the severity of inertia, CDFs shift progressively to the right, with median entry times rising to $t = 0.2$ at medium inertia and $t = 0.3$ (where subjects earn 95% of earnings available at t^*) at low inertia and finally reach $t = t^* = 0.4$ when inertia reaches zero.[25] The results thus show that entry times rise smoothly towards Perfectly Continuous levels as inertia falls towards zero, providing us with a next result:

Result 3.2. *High levels of inertia cause entry delay to completely collapse as Nash equilibrium predicts. However as inertia falls towards zero, entry times approach Perfectly Continuous levels.*

The survival of high levels of cooperative delay in the face of small amounts of inertia is starkly inconsistent with Nash equilibrium (which predicts a complete collapse in cooperative delay with any inertia) but broadly consistent with ε-equilibrium. Though the results are perfectly consistent with ε-equilibrium, the smooth path of convergence is not explained by ε-equilibrium due to multiplicity (once inertia falls enough to allow entry times later than $t = 0$, *any* entry time in $[0, t^*]$ is supportable in ε-equilibrium). In the next section we consider our results in light of recent findings in the literature on dynamic strategic interaction and develop a more satisfying, structured and precise explanation for these patterns.

3.5 Discussion: Strategic Uncertainty and Continuity

Why does inertia have a "smooth" effect on entry instead of causing the immediate collapse in delay predicted by Nash? ε-equilibrium is broadly consistent with this

[25] An exact Jonckheere-Terpstra test allows us to reject the hypothesis that distributions of session-wise median product limit estimates of entry times are invariant to inertia against the alternative hypothesis that they are (weakly)monotonically ordered by inertia ($p < 0.001$).

pattern but provides little insight into either its source or (due to multiplicity) its precise shape. One appealing answer is that the rich dynamic environment of a continuous time game makes it difficult to arrive at the sort of common knowledge required to support Nash equilibrium, forcing subjects to grapple with unresolved strategic uncertainty when making their decisions. Indeed, strategic uncertainty has emerged as a central explanatory variable for cooperation in dynamic games and both finitely and infinitely repeated prisoner's dilemmas in prominent recent work. To measure the strategic risk of cooperating, the literature typically restricts attention to the strategies Always Defect and Grim Trigger and calculates the basin of attraction of defection (hereafter, the BOA) – the minimal probability one must assign to one's counterpart playing Grim in order for Grim to be a best response. Intuitively, the greater the basin of attraction, the more risky it is to attempt to cooperate: when the BOA is greater than 0.5 it becomes risk dominant to always defect. Both prospective experiments (e.g. Dal Bo and Frechette (2011), Embrey et al. (2016), Vespa and Wilson (2016)) and wide ranging retrospective meta-analyses (Dal Bo and Frechette (2016), Embrey et al. (2016)) reveal that this simple short-hand measure and corresponding notions of risk dominance have startlingly strong predictive power for cooperation rates in both infinitely repeated games where cooperation is an equilibrium and – importantly for our application – finitely repeated games where it is not.

Importantly, this simple measure of strategic uncertainty, when adapted to our game, also crisply organizes the large treatment effects of inertia we observe in our data. To adapt this measure to our setting, we consider the strategies "enter now at time t" and "wait to enter at t^*"– the closest analogues to Always Defect and Grim Trigger for our game – and calculate the "enter now" BOA for each t in $[0,t^*]$, allowing us study how the strategic risk of entering immediately changes as the game progresses.[26] Figure 3.5 plots the BOA at each t for each of our

[26] Again, here and in the remainder of this section, we restrict attention to *admissible strategies* (a strategy is admissible if it is not weakly dominated: see Brandenburger et al. (2008) for an extensive discussion of the role of admissibility in games.). The main implication is that when a player chooses "enter now" and the other "wait until t^*," the second entrant responds to her counterpart by entering as soon as possible (given the timing protocol). The strategy "wait until t^*" should therefore be read as "enter at t^* or as soon as possible after my opponent enters, whichever is earliest." As emphasized above, such admissible strategies are virtually universally employed in the data.

continuous time treatments and plots portions of the BOA at which “enter now” is risk dominant in red. Under high inertia (IC_{10}) the BOA is always 1 and it is always risk dominant to enter now[27] while under zero inertia (*PC*) the BOA is always 0 and delay is risk dominant until t^*. In the two intermediate treatments IC_{60} and IC_{280}, the BOA changes over time, with immediate entry becoming risk dominant at a different intermediate time in each case.

We make two observations. First, except where the measure reaches its boundary of 1, the basin of attraction is smaller, at each t, in treatments with larger inertia, suggesting that cooperation is indeed more strategically risky (relative to immediate entry) at higher levels of inertia. Second, the times at which immediate entry becomes risk dominant almost perfectly corresponds to median entry times in all of our treatments: BOA reaches 0.5 at times (0, 0.198, 0.308 and 0.4)[28] in treatments (IC_{10},IC_{60}, IC_{280} and *PC*) and median product limit entry times track closely at (0.022,0.188,0.325, 0.392). Thus, subjects enter in each treatment precisely when immediate entry becomes risk dominant, suggesting that strategic uncertainty has a strong role in shaping our treatment effects.

3.5.1 Three decision rules

The basin of attraction provides a convenient and easily interpretable measure of strategic uncertainty and suggests a strong link between strategic uncertainty and behavior in our data. It is, however, built on fundamental simplifications that make it better suited to benchmarking levels of strategic uncertainty in our game than to modeling exactly how subjects make decisions in the face of strategic uncertainty. Can we put our findings on a firmer footing by considering specific heuristic responses to strategic uncertainty that generate point predictions against which we can compare our data? Because we are conducting this exercise *ex post*, our aim is to focus on parsimonious decision rules that put minimal structure on beliefs subjects hold about their counterparts' strategies, and are therefore difficult to adjust to fit the data *ex post*. To achieve this parsimony we consider models that replace Nash equilibrium's extreme assumption that agents perfectly know one another's

[27]The BOA is also always 1 in the PD treatment.

[28]Strictly speaking, the BOA never rises above 0 in the PC treatment but we describe the separatrix as 0.4 to highlight the fact that even at times arbitrarily close to 0.4, entry is not risk dominant.

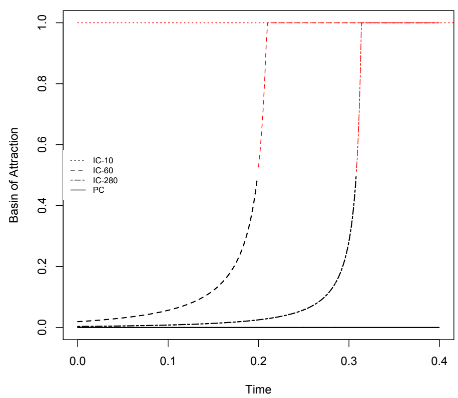

Figure 3.5: Basin of attraction of immediate entry, calculated at each time t. Red coloring denotes points at which immediate entry is risk dominant. Horizontal lines at 1 (for IC-10) and 0 (for PC) signify treatments in which the basin of attraction does not change over time.

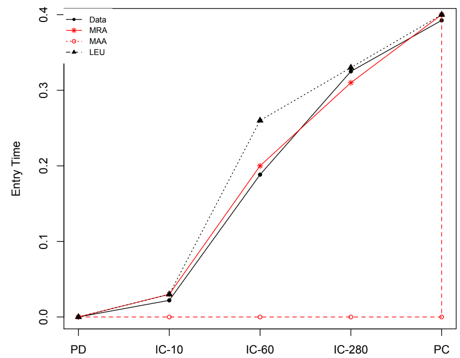

Figure 3.6: Median product limit entry times and predictions from the Minimax Regret Avoidance (MRA), Maximin Ambiguity Aversion (MMA) and Laplacian Expected Utility (LEU) models.

strategies with the opposite extreme assumption that agents are maximally uncertain, weighting all counterpart strategies symmetrically *ex ante*.[29, 30]

In a classic paper, Milnor (1954) considers three non-parametric (and therefore highly parsimonious) decision rules for uncertain environments like this that we can apply to our strategy set and compare to our data. The Laplacian[31] Ex-

[29] See Stoye (2011a) and Milnor (1954) for descriptions of the symmetry axiom we have in mind. See Arrow and Hurwicz (1972) for an argument that this type of symmetry is appropriate in models of fundamental uncertainty.

[30] As in previous sections, we place only two restrictions on the set of priors: (i) we restrict to admissible strategies and (ii) entry occurs in $[0, t^*]$. Both of these characteristics of strategies are virtually universally observed in the data. In considering models of strategic uncertainty, we restrict attention to simple decision rules and beliefs that are not disciplined by equilibrium, though several authors have proposed equilibrium extensions of these rules (e.g. Renou and Schlag (2010), Halpern and Pass (2012), Lo (2009)). While we avoid adopting the stronger assumptions and greater structure required of these equilibrium concepts for this exercise, we note that many of these equilibrium concepts generate identical predictions to those discussed below.

[31] So named for Laplace's (1824) argument that uniform beliefs should be applied to unknown

pected Utility (LEU) rule is the subjective expected utility response to this type of strong uncertainty and models agents as simply choosing actions that constitute best responses to a uniform (or "Laplacian") distribution of entry times by their counterparts on $[0, t^*]$. The Minimax Regret Avoidance (MRA) rule, first proposed in Savage (1951) and axiomatized by Milnor (1954) and Stoye (2011a), results from relaxing the independence of irrelevant alternatives axiom and specifies that agents choose the action that minimizes the largest *ex post* regret (the difference between earnings actually generated by a strategy choice and the earnings a different strategy choice might have generated given counterparts' strategies) over all strategies. Finally, the Maximin Ambiguity Avoidance[32] (MAA) rule, proposed by Wald (1950) and axiomatized by Milnor (1954), Gilboa and Schmeidler (1989) and Stoye (2011b), relaxes the independence axiom in expected utility theory and specifies that agents choose a strategy that yields the largest minimum payment (over other subjects' strategies) an agent could achieve.[33]

As we show in Online Appendix C.4.1, these decision rules make very different predictions for the way inertia shapes decisions in continuous time games like ours. The MAA rule predicts exactly what Nash equilibrium predicts: immediate entry at $t = 0$ for any inertia greater than zero. The MRA and LEU rules, by contrast, each predict a smooth pattern of progressively later entry as inertia falls towards zero, terminating at t^* when inertia is zero, though the rate of convergence is different in each case. We calculate predictions for each of these rules for each treatment, and present the results in Figure 3.6 along with medians of subject-wise product limit estimates. Of the three heuristics, MAA does the worst by predicting exactly what Nash equilibrium does for each treatment. Both MRA and LEU, by contrast, do an excellent job of tracking the data but the MRA heuristic fits point estimates from

events due to the principle of insufficient reason (see e.g. Morris and Shin, 2003). Laplacian beliefs have an important role in the literature on global games (see Morris and Shin (2003)).

[32]We deliberately avoid the more common MEU acronym for this decision rule to emphasize a subtle difference in interpretation between the standard MEU model and our application. In the standard MEU model, as in Gilboa and Schmeidler (1989), the set of priors is treated as endogenous to the agent's preferences. By contrast, we interpret MAA as a decision rule that is applied to an exogenous set of uninformative beliefs (as in Stoye (2011b)).

[33]Milnor refers to these as the Laplace, Savage and Wald rules respectively. He also discusses a fourth rule that he calls Hurwicz (commonly today called α-maxmin) which we reject because it has a free parameter and therefore can be "tuned" to the data in an *ex post* exercise like this one.

the data almost perfectly and is the most accurate of the three models. We report this as our next result:

Result 3.3. *Median entry times across treatment are almost perfectly organized by predictions made by the MRA decision rule, suggesting that reactions to strategic uncertainty are an important driver of behavior.*

Our results suggest that assuming subjects are highly strategically uncertain about their counterparts' behavior can generate significantly better predictions than making the opposite assumption that strategic uncertainty is eliminated in equilibrium. Interestingly evidence for such strategic uncertainty doesn't seem to ease much as subjects acquire experience in our data: median entry times in the final period of play track MRA predictions across levels of inertia just as well as product limit estimates do using data from the whole session (as visualized in Figure 3.6). There is, moreover, no evidence of movement towards Nash equilibrium over time in any of our continuous time treatments (except in IC_{10} where MRA actually predicts Nash-like outcomes), suggesting that strategic uncertainty continues to play an essential role in determining behavior even after dozens of periods of play. As it turns out, MRA predictions are surprisingly robust to the type of feedback subjects acquire in dynamic settings like ours because most uncensored feedback subjects receive via repeated play concerns whether counterparts tend to enter early in the game. Since regret is primarily shaped by the possibility that one's counterpart will enter later in the game, MRA predictions change little when subjects learn that early entry events by counterparts are unlikely. Consequently, MRA predictions tend to be fairly durable in the face of experience, just as later entry times are in our data.[34]

[34] Consider a subject who repeatedly enters at or near the MRA predicted entry time, t_{MRA}. She will learn the approximate distribution of entry times over the interval $[0, t_{MRA}]$, but because of censoring she will not learn anything about the distribution of (intended) entry times over the interval $[t_{MRA}, t^*]$. For all entry times $t < t_{MRA}$ the subject faces her maximal regret when her opponent enters at t^* and, because of censoring, she cannot rule out that her opponent may intend to enter at t^*. For entry times $t > t_{MRA}$ she faces her maximal regret when her opponent slightly pre-empts her, which again she cannot rule out. In other words, our MRA prediction is robust to a subject learning that some, or even all, early entry times are never used. This is because early entry times do not cause large regret: regret is maximized when opponent's are either fully cooperative (enter at t^*) or enter immediately prior to the subject's intended entry time. Given that the MRA decision rule seeks to minimize maximal regret, it is uncertainty over these later entry times that causes MRA deviations from Nash equilibrium. See the proof in Online Appendix C.4.2 for more detail on the mapping between the set

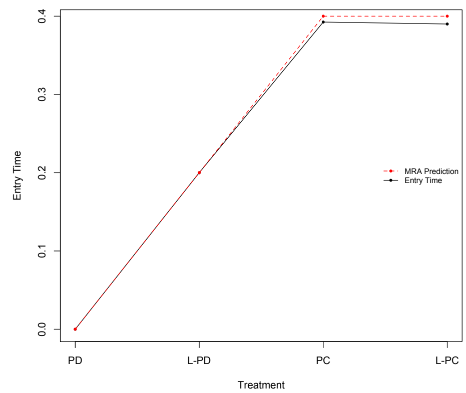

Figure 3.7: Median observed entry times (first column) and MRA predictions from the diagnostic L-PD and L-PC treatments. PD and PC treatments from the main design are also included for comparison.

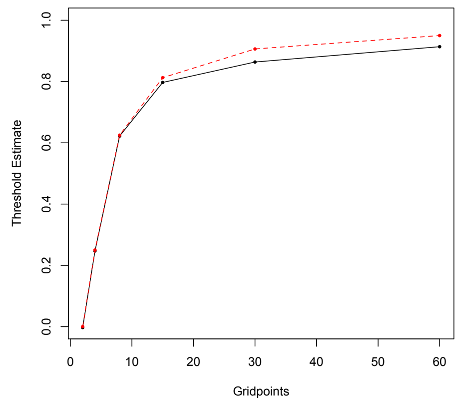

Figure 3.8: Median cooperation rates from the Grid-n treatment from Friedman and Oprea (2012) compared to MRA predictions.

3.5.2 Validation using alternative comparative statics

We designed and ran two additional treatments to study whether our explanation for the comparative static effect of inertia can also explain other, distinct comparative statics. In the L-PD and L-PC treatments we replicate the Perfectly Discrete and Perfectly Continuous treatments but dramatically lower the preemption temptation parameter Π_F from 4 to 1.4. In the PC treatment, lowering this parameter has no effect on strategic risk as measured by the basin of attraction and does not change the MRA point prediction under Perfectly Continuous time protocols (the prediction is $t = 0.4$ in either case). That is, under the class of explanations we've considered thus far, the PC and L-PC treatments should generate identical behaviors.

By contrast, in Perfectly Discrete time protocols, strategic uncertainty changes a great deal when we lower Π_F in the L-PD treatment: while immediate entry is always risk dominant in the PD treatment (under the simple BOA measures discussed above) it becomes risk dominant only at an interior point in the L-PD treatment, as

of believed entry times and MRA predictions.

in the IC_{60} and IC_{280} treatments, suggesting lowering Π_F may generate a later entry time in discrete time. Most importantly for our purposes, the MRA prediction rises from 0 to 0.2 when we lower Π_F in the L-PD treatment.[35]

Figure 3.7 shows median subject-wise product limit estimates for the PD, L-PD, PC and L-PC treatments and MRA point predictions. The results nearly perfectly track the point predictions provided by the MRA heuristic. Entry times rise from 0 to 0.2 when we lower Π_F in the L-PD treatments but remain constant at about 0.4 when we make the same parameter change in the L-PC treatments, just as the MRA rule suggests.[36]

Result 3.4. *Results from additional diagnostic treatments varying payoff parameters in discrete and continuous time are well organized by measures of strategic uncertainty and point estimates are virtually identical to point predictions generated by the MRA rule.*

3.5.3 Validation using other continuous time games

The MRA decision rule organizes behavior in our game remarkably well, predicting, in particular, the smooth approach to Perfectly Continuous-like cooperative benchmarks we observe as inertia falls to zero. How relevant are these sorts of results for understanding behavior in other continuous time games? To find out, we test the MRA rule against data from the continuous prisoner's dilemma, the simplest game in a broad and empirically important class of games in which efficient outcomes are in tension with individual incentives. The Grid-n treatment in Friedman and Oprea (2012) studies 60 second prisoner's dilemmas that are divided up into 4, 8, 16, 32 and 60 Perfectly Discrete time subperiods, within subject. This time protocol creates the equivalent of exogenous reaction lags in continuous time lasting 50%, 25%, 12.5%, 6.6%, 3.3% and 1.6% of the game, respectively, generating a similar effect to inertia in our Inertial Continuous time games.

[35] We originally designed these additional treatments to validate ε-equilibrium comparative statics which are similar to and consistent with MRA predictions but are less precise, just as with treatments from the main design.

[36] A Mann-Whitney allows us to reject the hypothesis that session-wise median product-limit estimates of entry times in the PD and L-PD treatments are from the same distribution ($p = 0.017$); the same test does not allow us to reject the same hypothesis regarding the PC and L-PC treatments ($p = 0.183$).

In Figure 3.8 we plot median final mutual cooperation times (measured as a fraction of the period) as a function of the number of grid points.[37] Over this we overlay MRA predictions[38] for the earliest time at which mutual cooperation can evaporate as a function of the number of grid points. (Any date after the times plotted can be supported under the MRA). Strikingly, these earliest MRA defection times nearly perfectly match median final cooperation times, converging towards the Perfectly Continuous time limit of 1 (cooperation until the very end of the period) as the number of grid points grow large and the forced reaction lag grows small. The results thus provide strong out-of sample confirmation that the MRA heuristic organizes convergence paths to Perfectly Continuous time benchmarks.[39]

Result 3.5. *The earliest MRA defection time generates accurate point predictions for convergence to continuity in the continuous time prisoner's dilemma.*

The prisoner's dilemma is the simplest in a broad set of strategically similar games that include important applications like Bertrand pricing, Cournot quantity choice, public goods, and team production problems. Though, of course, we cannot test every game in this class, the fact that MRA-predicted results like ours extend to the continuous prisoner's dilemma is strongly suggestive that such effects are relevant for a much larger set of strategically similar games. In Online Appendix C.4.2 we provide some support for this intuition by showing that, under plausible (and particularly empirically relevant) specifications of beliefs over the strategy space, the MRA rule predicts similar convergence results for this broad class of dilemma-like games. Combined, our results suggest that rules like the

[37] As in the other analyses in this chapter we use the full dataset in making these measurements. Restricting attention to the final 2/3 of the session as Friedman and Oprea (2012) do generates similar results.

[38] Calculated under restrictions on beliefs discussed in Online Appendix C.4.2.

[39] MRA also predicts the high rates of cooperation and low variation over parameters Friedman and Oprea (2012) observe in their (Inertial) Continuous time treatments. These treatments study 60 second continuous time prisoner's dilemmas – prisoner's dilemmas with flow payoffs realized in continuous time (though with subject-generated inertia). Friedman and Oprea's (2012) design sets mutual cooperation payoffs of 10 and "suckers" payoffs of 0 and varies the temptation payoff (x) and defection payoff (y) cyclically over 32 periods over four parameterizations: Hard (x=18, y=8), Easy (x=14, y=4), Mix-a (x=18, y=4), Mix-b (x=14, y=8). The median final time of mutual cooperation (in 60 second periods) are 59.6, 58.4, 58.44, 57.6 in the Easy, Mix-a, Mix-b and Hard treatments, which are very tightly clumped near the Perfectly Continuous benchmark time of 60. This nearly perfect cooperation and minimal variation over parameters is explained by MRA, which predicts earliest collapse of cooperation at 58.9, 58.5, 56.5 and 55.5 seconds in these four treatments.

MRA, founded in strategic uncertainty, provide an empirically plausible mechanism by which we might expect nearly-Perfectly Continuous levels of cooperation to emerge and persist even in the presence of realistic inertia.

3.6 Conclusion

Perfectly Continuous and Perfectly Discrete time are both idealizations but they are illuminating ones, functioning as strategic analogues to vacuums in the physical sciences. Like vacuums, they are environments in which theoretical forces are cast in particularly high relief and results can be crisply interpreted in the light of theory. Although Perfectly Discrete time behavior has been exhaustively studied in thousands of experimental investigations, Perfectly Continuous time has never been studied before and for a very simple reason: natural frictions in human interaction that loom especially large in the relatively fast pace setting of a laboratory experiment push strategic environments meaningfully away from the Perfectly Continuous time setting described in the theory. Our chapter introduces a methodological innovation that eliminates these frictions, allowing us to observe, for the first time, Perfectly Continuous behavior. By observing and comparing behavior across these two "pure" environments and by comparing both to more naturalistic protocols in between we learn some fundamental things about dynamic strategic behavior.

Results from our initial baseline parameters are nearly perfectly organized by benchmarks proposed in the literature. Though our game suffers from multiple (indeed, a continuum of) equilibria in Perfectly Continuous Time, we observe entry times tightly clustered at the interior joint profit maximizing entry time under this timing protocol. This decisive equilibrium selection strongly supports a weak dominance refinement argued for by Simon and Stinchcombe (1989) for Perfectly Continuous games. By contrast, under the exact same parameters, in Perfectly Discrete Time we observe almost universal, highly inefficient first-period entry that is perfectly in line with backwards induction. Thus our baseline results show strong evidence of a large and economically significant gulf between Perfectly Discrete and Perfectly Continuous time behaviors.

How do results from these artificial settings relate to more realistic strategic interactions? Most real human decisions are made neither perfectly synchronously (as in Perfectly Discrete time) nor with instant response (as in Perfectly Continuous time). More realistic are real time, asynchronous settings in which there is some delay in mutual responses, even if small. Nash equilibrium predicts that even a tiny amount of such inertia will be sufficient to erase all of the cooperative equilibria generated by Perfectly Continuous time. However ε-equilibrium suggests that the correspondence between Inertial Continuous time behavior and the benchmarks of Perfectly Discrete and Perfectly Continuous time depends crucially on the size of inertia. While very high levels of inertia can cause ε-equilibrium sets to coincide with Perfectly Discrete behavior (as suggested by Nash), very low levels of inertia can push the ε-equilibrium set to coincide with the Perfectly Continuous equilibrium set. We study such settings in our Inertial Continuous time treatments in which subjects interact (under Baseline parameters) in continuous time but with natural human reaction lags (clocked at roughly 0.5 seconds in our subjects). By varying the speed of the game clock we are able to alter the severity of naturally occurring inertia in subjects' decision making and study the robustness of Perfectly Continuous time behavior to multiple levels of inertia. Our results show that Nash equilibrium-like collapses to Perfectly Discrete-like benchmarks occur in continuous time when inertia is very high. But at low levels of inertia, subject entry delays approach the efficient levels generated in (and predicted for) the Perfectly Continuous treatment.

We close the chapter by considering sharper (psychologically) and crisper (predictively) explanations for our results than ε-equilibrium can provide. Recent research on dynamic strategic interactions have amassed a great deal of evidence that strategic uncertainty faced by attempting cooperation (as measured by the basin of attraction for defection) and related notions of risk dominance have strong predictive power for cooperation even in games in which cooperation is unsustainable as a Nash equilibrium. Applying similar measures to our game, we find that strategic uncertainty subjects face when attempting to cooperate rises sharply with inertia, supporting a conjecture that strategic uncertainty shapes behavior in continuous time games in a way that Nash equilibrium cannot capture.

Inspired by the organizing power of this measure in our (and other) data, we consider a series of simple heuristics that replace Nash equilibrium's extreme assumption that subjects perfectly know their counterparts' strategies with the opposite extreme that subjects know very little about counterparts' strategies. We show that several such decision rules strongly outperform Nash equilibrium and that one (the Maximin Regret Avoidance model) almost perfectly matches cross-treatment point estimates from our data. To strengthen our analysis we expose the MRA rule to an additional test using a pair of additional treatment and find similarly strong evidence. Finally, we show that heuristics like the MRA generate increasingly cooperative behavior as games grow more continuous in an important class of empirically relevant games. We test this prediction on continuous time prisoner's dilemmas from past work and show, once again, that the MRA heuristic nearly perfectly organizes both point predictions and treatment effects over which Nash equilibrium makes starkly counterfactual predictions. The results suggest that benchmarks that assume no knowledge of others' strategies and impose little structure on subjects' beliefs do a better job of anticipating behavior (even behavior of experienced subjects) than benchmarks that assume perfect knowledge. These models also explain why, as in our data, Perfectly Continuous-like behavior persists even in the presence of inertia.

The results from our experiment – and supporting theoretical benchmarks – suggest an appealing framework for understanding the relationship between the abstractions of Perfectly Discrete and Perfectly Continuous time and real world behavior. Perfectly Discrete and Perfectly Continuous time predictions can be thought of as polar outcomes that each approximate realistic (Inertial Continuous time) behavior when inertia is either very high or very low, respectively. Indeed, we can easily push real time (Inertial) behavior close to either Perfectly Discrete or Perfectly Continuous time behavior simply by varying the severity of inertia. Concretely, these sorts of results suggest that Perfectly Continuous time benchmarks can, in some cases, be more empirically relevant than Discrete time benchmarks, even if agents face frictions that should be sufficient to short circuit Perfectly Continuous time equilibria under standard theory. The rise of thick online global markets, always-accessible mobile technology, friction reducing applications and automated online agents have made strategic interactions more asynchronous and lags in response less severe. These trends, which seem likely to intensify in the coming years, have the effect of pushing many interactions closer to the setting of Perfectly

Continuous time. Though these technological changes may never drive inertia entirely to the Perfectly Continuous limit of zero, our results suggest that behavior can nonetheless come close to Perfectly Continuous levels as inertia falls. This deviation from standard theory in turn suggests that we might expect Perfectly Continuous time predictions to become an increasingly relevant way of understanding economic behavior relative to the Perfectly Discrete predictions most often used in economic models.

Chapter 4

Mental Equilibrium and Mixed Strategies for Ambiguity Averse Agents

4.1 Introduction

Consider a game between two agents that is mediated by a game theorist. The agents report their strategies to the game theorist, who then resolves the outcome of the game and pays the agents their winnings (or collects their losses). The game theorist may allow mixed strategy reports from the agents, and resolve the mixed strategy herself, or she may require that the agents report a pure strategy. If the game theorist requires pure strategy reports, as is the case in Chapter 2, then the game theorist should be aware that the agents may still be using a mixed strategy that they are resolving privately before reporting. Under the standard formulation, where agents have expected utility preferences, the set of equilibrium under the two reporting requirements will be indistinguishable.

However, when agents have ambiguity averse preferences then the different reporting requirements may induce different games with different equilibrium. Equilibrium concepts such as Lo (2009), Dow and Werlang (1994) and Eichberger and Kelsey (2000) enforce pure strategy reporting and generate larger equilibrium sets than Lo (1996) or Klibanoff (1996) which allow for mixed strategy reports. The difference arises because ambiguity averse preferences are non-linear and agents may have a strict preference for mixed strategies. The equilibrium concepts that enforce pure strategy reporting assume that only pure strategies are available to agents: they implicitly rule out private pre-play mixing. By its very nature, however, private pre-play mixing will be unobservable to the game theorist and cannot readily be prevented.

How, then, does allowing for private pre-play mixing affect the equilibrium in papers such as Lo (2009), Dow and Werlang (1994) and Eichberger and Kelsey (2000)? The answer, provided for the first time in this chapter, is that allowing for private pre-play mixing has no effect on the equilibrium set for agents with preferences that lie in the intersection of Choquet Expected Utility (Schmeidler, 1989) and Maxmin Expected Utility (Gilboa and Schmeidler, 1989). While this result is of independent interest, it is particularly relevant for experimental tests of ambiguity averse equilibrium concepts. Recent experiments[1] have used equilibrium concepts that restrict themselves to pure strategies, and thereby implicitly assume that their subjects are using only pure strategies and are not engaging in pre-play mixing. Given that it is not possible to actively prevent subjects from pre-play mixing, the results in this Chapter are essential for a direct interpretation of the data in the previous experimental literature using pure strategy solution concepts.

We establish our key result using an application of results from Gilboa and Schmeidler (1994) for agents in individual decision making problems. The application of Gilboa and Schmeidler (1994) to games yields two additional insights. First, the technique allows for an extension of the static equilibrium concepts discussed above into dynamically consistent extensive form equilibrium concepts. Second, because the technique linearizes non-linear ambiguity averse preferences, it also allows us to formulate the equilibrium set of a game with ambiguity averse agents as the solution to a linear complementarity problem – thereby providing a simple mechanical technique for solving games with ambiguity averse agents.

The rest of this chapter is organized as follows. Section 4.2 introduces the mathematical tools from Gilboa and Schmeidler (1994) that are used in the rest

[1] See Section B.1.5 for an overview of the literature, that also includes Chapter 2 of this dissertation.

of the chapter. Section 4.3 introduces the notion of a game as a set of interacting decision problems and presents the key idea of a "mental" state space. Sections 4.4 and 4.5 apply the structure of a "mental" state space to the equilibrium concepts contained in Lo (2009) and Dow and Werlang (1994), respectively, and establishes the key result of this Chapter. Section 4.6 demonstrates the extension of a "mental" equilibrium to an extensive form game in a dynamically consistent fashion, using the equilibrium concept from Lo (2009) as an example. Section 4.7 demonstrates how the "mental" formulation of a game gives rise to a linear complementarity program that can be easily solved to find the equilibrium of a normal form game. Section 4.8 concludes, and proofs are relegated to Appendix D.

4.2 Preliminaries

Suppose that there exists a finite[2] set of *states* of the world, $\omega \in \Omega$, and that an *act*, f, maps each state to an *outcome* in $\mathbb{R}$; that is $f : \Omega \mapsto \mathbb{R}$. Choquet Expected Utility (CEU), first introduced by Schmeidler (1989), generalizes Subjective Expected Utility by allowing a decision maker to hold non-additive beliefs which are represented by a capacity, ν, defined over the set of events $\Sigma = 2^{\Omega}$. Suppose, without loss of generality, that for a given act, f, the set of states can be ordered so that $f(\omega_1) \geq f(\omega_2) \geq \ldots \geq f(\omega_n)$. A CEU agent calculates her utility of an act by evaluating the (discrete) Choquet integral:

$$\int f \, d\nu = \sum_{i=1}^{n} f(\omega_i) \left[\nu(\cup_{m=1}^{i} \omega_m) - \nu(\cup_{m=1}^{i-1} \omega_m) \right] \tag{4.1}$$

We shall assume throughout that the capacities, ν_i, are *belief functions* and we use V to denote the space of all such capacities. That is, we assume that $\nu_i(\Sigma_i) = 1$ and that ν_i is totally monotone.[3] Under these assumptions it is also possible to represent the agent's preferences using Maxmin Expected Utility (Gilboa and Schmeidler, 1989, 1994).

Throughout this chapter, we shall rely on two basic mathematical results that are demonstrated in Gilboa and Schmeidler (1994). Firstly, the non-additive mea-

[2]Although the finiteness of the state space can be relaxed, we restrict attention to finite sets here to mirror our later restriction to finite games.

[3]A totally monotone capacity is convex, but the converse need not hold.

sure ν can be spanned by an additive measure over an appropriately defined (larger) state space. Secondly, we can represent an agent with CEU preferences over Ω as, equivalently, having SEU preferences over the larger state space with an appropriately transformed set of acts.

Result 4.1 (Adapted from Gilboa and Schmeidler (1994))**.** *For $T, A \in \Sigma' = \Sigma \backslash \{\emptyset\}$, define*

$$e_T(A) = \begin{cases} 1 & T \subseteq A \\ 0 & \textit{otherwise} \end{cases}$$

Then the set $\{e_T\}_{T \in \Sigma'}$ forms a linear basis for V. The unique coefficients $\{\alpha_T^\nu\}$ satisfying

$$\nu = \sum_{T \in \Sigma'} \alpha_T^\nu e_T$$

are given by

$$\alpha_T^\nu = \sum_{S \subseteq T} (-1)^{|T|-|S|} \nu(S) \tag{4.2}$$

Furthermore, if ν is totally monotone then $\alpha_T^\nu \geq 0$ for all $T \in \Sigma'$ and if ν is normalized then $\sum_{T \in \Sigma'} \alpha_T^\nu = 1$.

Result 4.1 provides the key building block for this chapter: any non-additive measure over a state space can be spanned by an appropriately formed set of states constructed from the power set of the original state space. Furthermore, when the non-additive measure is a belief function then the spanning coefficients can be interpreted as probabilities over the newly constructed state space. Note the relationship between Result 4.1 and the proof of the representation theorem for CEU in Schmeidler (1989). In Result 4.1, we begin with a non-additive measure and 'restore' additivity by extending the state space. In Schmeidler (1989), the primitive is a SEU representation with respect to an additive measure, which is then extended to generate a CEU representation with respect to a non-additive measure. This tight relationship between Choquet Expected Utility and Subjective Expected Utility is formalized in the next result.

Recall that when ν is a belief function the core of ν is simply the set of probability measures, p, such that $p(A) \geq \nu(A)$ for all $A \in \Sigma$.[4] We now state result 4.2.

Result 4.2 (Corollary 4.4 from Gilboa and Schmeidler (1994))**.** *Suppose that ν is a belief function. Then for every $f \in F$*

$$\int f\, d\nu = \sum_{T \in \Sigma'} \alpha_T^\nu \left[\min_{\omega \in T} f(\omega) \right] \tag{4.3}$$

$$= \min_{p \in Core(\nu)} \sum_{\omega \in \Omega} p(\{\omega\}) f(\omega). \tag{4.4}$$

Result 4.2 demonstrates that an agent with CEU preferences with respect to a belief function can have their preferences represented via either MEU or SEU preferences. While this relation between MEU and CEU preferences is both straightforward and well known, the representation with SEU preferences requires the formation of a new set of acts over the set $\Sigma = 2^\Omega$, with the outcome associated with each new act defined by the min function in equation 4.3. We shall call these acts "mental" acts, and will sometimes refer to the event space as the "mental" state space. This terminology reflects that the event space may not be observable to an external observer and may, therefore, represent the mental accounting of the agent.[5]

Definition 4.1. *A mental act, f', is an extension of an act, f, defined over the event space, Σ', such that $f' : \Sigma' \mapsto \mathbb{R}$ with $f'(T) = \min_{\omega \in T} f(\omega)$ for all $T \in \Sigma'$.*

It follows from result 4.2 that the preferences of an agent, with CEU preferences with respect to a belief function, can be written in the expected utility form with respect to mental acts over the mental state space. Billot and Vergopoulos (2014) generalize this result to a broader set of preferences than just CEU with

[4] Note that, because each p in the core is additive, the core can be equivalently defined to be the restriction of this set of probability measures to an equivalent set of probability measures defined over Ω. In some places we use core (ν) to denote this restriction, but this should be clear from the context.

[5] The interpretation of the event space as a mental "state" space implies that the agent, at least in his mental accounting, is not well calibrated about the nature of the world. In a sense, this is precisely the trade-off that allows us to move from non-linear preferences in the observable state space to linear preferences in the mental state space.

respect to a belief function, and explore the implications of result 4.2 for problems with a single decision maker in some detail. It is possible to generalize the results in this chapter to cover the broader class of preferences covered in Billot and Vergopoulos (2014).

Before continuing with the substantive heart of the chapter, we first review why the introduction of mixed strategies may change the equilibrium set in games with ambiguity averse agents relative to a game where only pure strategies are available.

4.2.1 Mixing: Anscombe-Aumann vs. Savage interpretations of uncertainty

Suppose that an agent, who may choose from a finite set of acts, obtains a mixing device (a random number generator, say) and that they may condition their choice of act on the output of the mixing device. How should we define a mix over two acts, and how should the agent evaluate the payoffs that arise from these mixed acts? It is well established in the literature (Seo, 2009; Eichberger et al., 2016) that, for an ambiguity averse agent, the answer to these questions depends crucially on the way that we model the mixing device. The two most common modeling strategies originate with Anscombe and Aumann (1963) and Savage (1954). While both Anscombe and Aumann (1963) and Savage (1954) developed an axiomatization of additive expected utility preferences, our interest in these papers lies in the modeling frameworks rather than the representation theorems. Each of these frameworks has been extended and adapted to formulate representations of non-additive Choquet Expected Utility and Maxmin Expected Utility preferences; in what follows, we illustrate the role of mixing with reference to various axiomatizations of CEU preferences.

The Anscombe and Aumann (1963) model of decision making is a two stage model.[6] In the first stage the (uncertain) state of the world is resolved, and the outcome of an act in a state is a (risky) lottery; this timing structure is sometimes referred to as a horse-roulette lottery. The second stage is the resolution

[6] In the original formulation the Anscombe and Aumann (1963) model had three stages. Here, we refer to the cleaner two-stage adaptation of Anscombe and Aumann (1963) that was developed in Fishburn (1970) and used in later works including Schmeidler (1989). Seo (2009) also uses the three-stage formulation and provides a brief discussion.

of the roulette lottery and, in the CEU model in an Anscombe-Aumann state space (Schmeidler (1989)), the agent is assumed to have expected utility preferences over roulette lotteries. The randomization device is therefore appropriately modeled as acting state-by-state on the roulette lotteries: consider two acts, f and g and their mixture $\beta f + (1-\beta)g$. For each state s, the mixed act is defined such that $(\beta f + (1-\beta)g)(s) = \beta f(s) + (1-\beta)g(s)$.

The Savage (1954) model of decision making is a one stage model. All uncertainty, whether it is ambiguous or unambiguous, is resolved simultaneously. For example, in the Sarin and Wakker (1992) axiomatization of CEU in the Savage model there is a distinction between ambiguous and unambiguous states, while in the Gilboa (1987) axiomatization no such distinction is made. In this framework there is no clear meaning to the idea of mixing two acts; the object $\beta f + (1-\beta)g$ is simply not defined. The appropriate way to incorporate such an act in a Savage framework is to enlarge the state space in such a way that the new act can be defined in an appropriate fashion across each state. For example, each state, s, could be split into a state s_f and a state s_g and then we would define $(\beta f + (1-\beta)g)(s_f) = f(s)$ and $(\beta f + (1-\beta)g)(s_g) = g(s)$, and if the agent understands that the states s_f and s_g were realized with subjective probabilities in accordance with some mixing device then we would expect their preferences to reflect this knowledge.[7] As discussed below this approach is not compatible with our interpretation, in Section 4.3, of a game as a set of interacting decision problems.

As an example, consider the following setup. There are two states $s = \{A, B\}$ and two acts f and g, with payoffs given in table 4.1. For simplicity, assume that the agent is risk neutral.

	$s=A$	$s=B$
f	20	10
g	10	20

Figure 4.1: An example with two acts and two states.

Suppose that the agent has CEU preferences with $\nu(A) = 0.4 = \nu(B)$. The agent therefore values both f and g such that $U(f) = U(g) = 14$.

[7] See Klibanoff et al. (2005) for a formal exposition.

Now, suppose that we are working in an Anscombe-Aumann framework, and consider the prospect $h = \frac{1}{2}f + \frac{1}{2}g$. In the state $s = A$, h will generate a fifty-fifty lottery of 20 and 10. In the state $s = B$, h will generate a fifty-fifty lottery of 10 and 20. Therefore, in each state, h generates a lottery which the agent will value at 15.[8] h produces a constant lottery in all states, and should therefore be valued at the value of that lottery. Notice that this implies that, for our ambiguity averse agent, $h \succ f$ and $h \succ g$: the agent strictly prefers the mixed prospect h over the prospects that form the support of h.

In the Savage framework there is only one resolution of uncertainty so we cannot assign a lottery as an outcome; the outcome that occurs when a state occurs after an act has been chosen must be real valued. So, what value should we assign h in the state $s = A$? It is tempting to claim that this value should be 15, as in the Anscombe-Aumann example. However, recall that our agent has CEU preferences that represent pessimism. It seems reasonable that our agent may be pessimistic and believe that they will hold g more often when $s = A$ and hold f more often when $s = B$; in essence, our state space is not rich enough to fully capture the agent's preferences.

The standard approach in the Savage framework is to expand the state space by forming a new state space that is the cross product of the original state space and the set of possible outcomes of the available randomization device. This is precisely the approach outlined by Eichberger and Kelsey (1996), wherein they establish an indifference to randomizations when the state space is expanded in this fashion for the class of preferences considered in this chapter. We do not pursue the Eichberger and Kelsey (1996) approach in this chapter because it is difficult to reconcile this expanded state space with the interpretation of a game as a set of interacting decision problems with pure strategies as the choice objects, as we seek to do in Section 4.3.

As an alternative, we can use the results presented in section 4.2 to translate the agent's preferences into SEU preferences over the event space. To see this, we introduce a new state $s = \{A, B\}$ and form two new prospects f' and g' (figure 4.2). We call this the mental state space method, as the new state is a mental construct

[8] Recall that, in the CEU framework, the agent has standard SEU preferences over objective lotteries.

of the agent.

	$s=A$	$s=B$	$s=\{A,B\}$
f	20	10	$\cdot$
g	10	20	$\cdot$
f'	20	10	10
g'	10	20	10

Figure 4.2: An example with two acts and two states extended into the event space.

It is easy to see from Result 4.1 that $\alpha_A^v = \alpha_B^v = 0.4$ and $\alpha_{\{A,B\}}^v = 0.2$. Then Result 4.2 allows us to write the agents preferences in an expected utility form, and we conclude that we can form the prospect $h' = \frac{1}{2}f' + \frac{1}{2}g'$ and that the agent must have preferences such that $U(h') = \frac{1}{2}U(f') + \frac{1}{2}U(g') = 14$. Furthermore, because preferences are preserved in the mapping from the original acts to mental acts by construction, we conclude that $U(h) = 14$.

It is also important to identify the timing of randomizations in each of the models. In the Anscombe-Aumann model it is explicitly the case that the objective mixing occurs after the subjective state is realized. In the case of a game, this would be consistent with the case where agents submit mixed strategies to the game theorist, the game theorist announces the mixed strategies and then resolves the mixed strategies to reveal the outcome of the game. Alternatively, the game theorist could simply announce the (pure) outcome of the game as a function of the submitted mixed strategies. In the Savage model mixing, and by extension, mixed strategies are not well defined. In this case, the only reasonable interpretation is that agents use and report pure strategies to the game theorist.

The mental state space interpretation of mixed acts is that the mixing device is resolved (weakly) *before* the resolution of the state. In this interpretation we may model the resolution of uncertainty in two stages. Ex-ante, the agent has SEU preferences over the mental state space and mental acts and, by extension, any available mixing devices. At the interim stage, the mixing device has been resolved and the agent is holding one of the original acts in the original state space (over which there may be subjective uncertainty). Ex-post, all uncertainty has been resolved and the agent receives the outcome associated with the realized state and act selected by the randomization device. In a game setting, this interpretation is consistent with the agent using a private mixing device before reporting the pure strategy realization of the mixing device to the game theorist.

The example presented above highlights the two main concerns for mixed strategies in games with ambiguity averse agents (and also suggests a resolution). If the game theorist models the game with an Anscombe and Aumann (1963) state space then agents will have a strict preference for mixed strategies which, in many cases, causes the behaviour of ambiguity averse agents to be indistinguishable from standard Nash agents (Lo, 1996; Klibanoff, 1996). Alternatively, if the game theorist models the game with a Savage (1954) state space then the payoff for mixed strategies is not well defined and the game theorist does best by restricting agents to pure strategies. Each of these approaches is considered in the previous literature, which is summarized in the appendix in Table A.1.

The resolution proposed in this chapter is to expand the game into the mental state space, where mixed strategies are well defined provided that the mixing device is realized before strategies are revealed.

The discussion so far is consistent with, and could also have been motivated by, more recent literature on the timing of randomization in individual decision problems. Seo (2009) demonstrates that, for agents who violate the reduction of compound lotteries, the above argument hinges critically on whether the objective randomization occurs before or after uncertainty is resolved. Furthermore, Halevy (2007) provides experimental evidence that subjects who exhibit ambiguity aversion overwhelmingly do not reduce compound lotteries. Therefore we conclude, for ambiguity averse agents, that the order of realization of randomizations is important. Additionally, Eichberger et al. (2016) demonstrate that, for a dynamically consistent agent, the agent will be indifferent between the randomization and the initial prospects if the randomization occurs before the uncertainty is resolved , but that a preference for randomization may exist if the randomization occurs after the uncertainty is resolved.

4.3 Games as Interacting Decision Problems

It is possible to reformulate a standard game as a set of interacting decision problems.[9] Recall that a normal-form game is fully defined by the set of players, I, a set of strategies, $\{S_i\}$, for each player and a set of utility functions, $\{u_i\}$. For each agent we can formulate their decision problem as choosing from a set of acts over states where the states are given by $\Omega_i = \{\times S_{j\neq i}\}$ and $f_i(\omega) = s_i(\omega) = u_i(s_i, \omega) = u_i(s_i, s_{j\neq i})$.

4.3.1 Translating a game into the mental space

Any finite normal form game can be translated into a game played in the mental state space via an application of result 4.2 and definition 4.1. Figures 4.3 and 4.4 provide an example for the 3×2 game studied in Chapter 2.[10]

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

Figure 4.3: An example game

	X	Y	$\{X,Y\}$
A	25,20	14,12	14,$\cdot$
B	14,20	25,12	14,$\cdot$
C	18,12	18,22	18,$\cdot$
$\{A,B\}$	$\cdot$,20	$\cdot$,12	
$\{A,C\}$	$\cdot$,12	$\cdot$,12	
$\{B,C\}$	$\cdot$,12	$\cdot$,12	
$\{A,B,C\}$	$\cdot$,12	$\cdot$,12	

Figure 4.4: An example game in the mental state space with payoffs

The interpretation of figure 4.4 is that the row player perceives there to be an

[9] Aumann (1987) is the classic example.

[10] Groes et al. (1998) use a similar visual representation of extending a game into the event space, although they allow for both optimism and pessimism when extending preferences into the extra states, and exclude explicit randomizations. In contrast, the purpose of expanding the state space in this Chapter is precisely to allow for the inclusion of mixed strategies.

extra possible state, the state $\{X,Y\}$. In this state, the row player is uncertain which strategy the column player will play and evaluates their payoffs pessimistically. If the row player forms beliefs that can be represented by the belief function ν_R in the game in figure 4.3 then the row players beliefs over states are given by the $\alpha_T^{\nu_R}$ given by definition 4.1.

4.3.2 Mixed strategies

The interpretation of mixed strategies in a game with ambiguity averse agents requires some care. Typically, the previous literature has taken one of two approaches: either use an Anscombe and Aumann (1963) interpretation of uncertainty in which case mixed strategies are well defined but their inclusion renders the behaviour of ambiguity averse agents as indistinguishable from SEU agents, or use a Savage (1954) interpretation under which mixed strategies are not well defined and therefore are excluded from the feasible set of strategies (mixed equilibria are then interpreted using either a population interpretation or as an equilibrium in beliefs). See table A.1 for a summary of the previous literature and their treatment of mixed strategies.

Recently, the decision theory and experimental design literature has drawn a distinction between the case where subjective uncertainty is resolved before objective risk and the case where risk is resolved before uncertainty (Seo (2009), Azrieli et al. (2014), Baillon et al. (2014) and Eichberger et al. (2016)). The interpretation of mixed strategies proposed here is closely related to this distinction: we propose thinking of mixed strategies as a randomization over pure strategies that occurs in an agent's mental state space *before* the game is played. The implications for experimental implementations of games with ambiguity averse subjects are two-fold. Firstly, if subjects must choose pure strategies (i.e. they must enter a pure strategy choice into a computer which then matches the subject's response with another subject) then we are free to interpret the subject as playing an explicitly mixed strategy (where the subjects may resolve the randomization internally). Given that, as demonstrated below, the introduction of mixed strategies does not affect the equilibrium set this is purely a matter of interpretation. However, if providing subject's with the ability to choose a mixed strategy (perhaps the subject is given the ability to choose a point in a simplex to represent their strategy) then it is critically important to resolve the mixing device before the subject views their opponent's strategy choice.[11]

4.4 Lo-Nash Equilibrium

This section introduces Lo-Nash equilibrium, following Lo (2009) closely. Define a set of players $N = \{1, \ldots, n\}$, let each player $i \in N$ have a finite set of actions A_i, and define $A = \times_{i \in N} A_i$ and $A_{-i} = \times_{j \neq i \in N} A_j$. We shall endow each agent with a Von Neumann-Morgenstern utility function $u : A \mapsto \mathbb{R}$. Suppose that an agent has uncertainty regarding the strategy choices of their opponents, A_{-i}. Then we can regard a strategy, a_i, as an act over the state space A_{-i} generating a payoff $u_i(a_i, a_{-i})$ when the state a_{-i} is realized.

In a manner consistent with Gilboa and Schmeidler (1989)'s MEU formulation, we suppose that an agent's beliefs regarding their opponents strategies are a closed and convex set of probability measures $\Phi_i \subseteq \Delta(A_{-i})$. Given Φ_i an agents preferences are represented by

$$\min_{\phi \in \Phi} \sum_{a_{-i} \in A_{-i}} u_i(a_i, a_{-i}) \phi_i(a_{-i}).$$

Furthermore, we use σ to denote a probability measure on A. We define $\sigma^{A_i}(a_i) = \sum_{a_{-i} \in A_{-i}} \sigma(a_i, a_{-i})$ as the marginal distribution of σ on A_i and $\sigma^{A_{-i}}(a_{-i}) = \sum_{a_i \in A_i} \sigma(a_i, a_{-i})$ as the marginal distribution of σ on A_{-i}. Then, in the usual fashion we write $\sigma(a_{-i}|a_i) = \frac{\sigma(a_i, a_{-i})}{\sigma^{A_i}(a_i)}$.

Finally, we write $\operatorname{supp}\sigma$ to denote the support of the probability distribution σ, and define $\operatorname{supp}\Phi_i$ to be the union of the supports of the elements of Φ_i, and write Φ to denote the profile $(\Phi_i)_{i \in N}$. We are now ready to define a Lo-Nash equilibrium.

Definition 4.2 (Lo-Nash equilibrium)**.** *A pair $< \sigma, \Phi >$ forms a Lo-Nash equilib-*

[11] Saito (2015) provides an alternative framework for understanding preferences for randomization under uncertainty. In the Saito framework, a preference (or not) for randomization is an endogenous feature of an agent's preferences rather than simply a function only of the timing of realizations. In Saito's notation an agent with preference parameter $\delta = 0$ is indifferent to randomization while an agent with $\delta = 1$ behaves as if a randomization device completely eliminates the effects of uncertainty, and an agent may hold any $\delta \in [0, 1]$.

rium if it satisfies

$$\sigma(\cdot|a_i) \in \Phi_i \;\; \forall a_i \in supp\, \sigma^{A_i}, \forall i \in N \tag{4.5}$$

$$supp\, \Phi_i = \times_{j \neq i} supp\, \sigma^{A_j} \;\; \forall i \in N \tag{4.6}$$

and

$$a_i \in \arg\max_{\hat{a}_i \in A_i} \min_{\phi_i \in \Phi_i} \sum_{a_{-i} \in A_{-i}} u_i(\hat{a}_i, a_{-i}) \phi_i(a_{-i}) \;\; \forall a_i \in supp\, \sigma^{A_i}, \forall i \in N \tag{4.7}$$

Equation 4.7 requires that all strategies that are played in an equilibrium are best responses, with preferences defined as MEU preferences with respect to the equilibrium conjectures Φ. Equations 4.5 and 4.6 are the consistency requirements: equation 4.6 ensures that a strategy is played with a positive probability iff it is expected to be played with a positive probability, and equation 4.5 forces actual strategies to be contained in the belief sets. Note that 4.5 allows for conditioning of σ on a_i - this allows for strategies to be correlated, but the realized strategy must lie within player i's belief set for all a_i. Nash equilibrium is a special case of Lo-Nash equilibrium, thereby ensuring existence of Lo-Nash equilibrium for all finite normal form games. Section A.1.1 presents a detailed example of Lo-Nash equilibrium.

4.4.1 Formulating Lo-Nash equilibrium as a mental equilibrium

To introduce a mental version of Lo-Nash equilibrium we first need to extend the state space into the event space and define preferences over the new state space. In a Lo-Nash equilibrium, each agent i faces a state space A_{-i}: their opponent's strategy space is their state space. We denote the extended state space for agent i to be $\Sigma_i = \{2^{A_{-i}}/\emptyset\}$. Then we define an agents' extended utility function, for all $T \in \Sigma_i$, as:

$$u'(a_i, T) = \min_{a_{-i} \in T} u(a_i, a_{-i}).$$

Notice that when $T \in A_{-i}$ then $u'(a_i, T) = u(a_i, a_{-i})$ so that the extended utility function is consistent with the original utility function. We also define the natural extension of σ over Σ_i: $\sigma(a_i, T) = \sum_{a_{-i} \in T} \sigma(a_i, a_{-i})$ for all $T \in \Sigma_i$. Finally we introduce the probability measure $\alpha_i \in \Delta(\Sigma_i)$ which can be interpreted as agent i's belief over the event space Σ_i. We write α to denote the profile $(\alpha_i)_{i\in N}$.

We are now ready to define our Mental Lo-Nash equilibrium. In essence Mental Lo-Nash equilibrium is simply an application of Result 4.1 to Lo-Nash equilibrium, as we establish in Theorem 4.1 below.

Definition 4.3 (Mental Lo-Nash equilibrium)**.** *A pair $<\sigma,\alpha>$ form a Mental Lo-Nash equilibrium if*

$$\sigma(T|a_i) \geq \sum_{\tau\subseteq T} \alpha_i(\tau) \;\; \forall T \in \Sigma_i, \forall a_i \in supp\, \sigma^{A_i}, \forall i \in N \tag{4.8}$$

$$\sum_{\{T: a_{-i}\notin T\}} \alpha_i(T) = 1 \Leftrightarrow \prod_{j\neq i} \sigma^{A_j}(a_j) = 0 \;\; \forall a_{-i} \in A_{-i}, \forall i \in N \tag{4.9}$$

and

$$a_i \in \arg\max_{\hat{a}_i\in A_i} \sum_{T\in\Sigma_i} u_i'(\hat{a}_i, T)\alpha_i(T) \;\; \forall a_i \in \; supp\; \sigma^{A_i}, \forall i \in N \tag{4.10}$$

Remark 4.1. If we restrict α such that $\alpha_i(T) = 0$ for all $T \notin A_{-i}$ for all i then Equation 4.8 reduces to

$$\sigma^{A_{-i}}(T) = \alpha_i(T) \;\; \forall T \in A_{-i}, \forall i \in N \tag{4.11}$$

and Equation 4.9 is redundant.

Equation 4.10 reduces to the standard Nash equilibrium best response condition because $u_i(a_i,T) = u_i'(a_i,T)$ for all $T \in A_{-i}$.

Evidence. When $\alpha_i(T) = 0$ for all $T \notin A_{-i}$ then $\sum_{T\in A_{-i}} \alpha_i(T) = 1$. Given that $\sum_{T\in A_{-i}} \sigma(T|a_i) = 1$ because σ is a probability measure, Equation 4.8 implies that $\sigma(T|a_i) = \alpha_i(T)$ for $T \in A_{-i}$ for every a_i. This ensures that the conditional distributions of σ are equal, and that they must also be equal to the marginal distribution – σ is a product measure. Equation 4.9 is then trivially satisfied. □

Remark 4.1 implies that when beliefs are additive over the strategy space, then Mental Lo-Nash equilibrium reduces to Nash equilibrium. We therefore do not need to independently establish the existence of Mental Lo-Nash equilibrium: the existence of Nash equilibrium guarantees the existence of Mental Lo-Nash equilibrium. It is not, however, necessarily the case that there exists a Mental Lo-Nash equilibrium that is not a Nash equilibrium (the prisoner's dilemma provides a simple counterexample).

Theorem 4.1. *If $<\sigma,\alpha>$ is a mental Lo-Nash equilibrium then there exists a Φ such that $<\sigma,\Phi>$ a Lo-Nash equilibrium. Conversely, if $<\sigma,\Phi>$ is a Lo-Nash equilibrium and Φ_i is the core of a belief function for all i then there exists an α such that $<\sigma,\alpha>$ is a mental Lo-Nash equilibrium.*

The interpretation of a mental Lo-Nash equilibrium is straightforward when σ is a product measure (i.e. strategies are uncorrelated). In this case, each agent uses an independent pre-play mixing device yet may place a positive weight on beliefs that imply correlation between their opponents, or even that their opponents strategies may be correlated with their own strategy. Such beliefs may be justified in cases where an agent believes that their opponents may be able to communicate, or where they fear that their opponents may be able to see the agent's strategy choice (or a signal correlated with their strategy choice) before making a decision. It is feasible that some subjects in a standard game theory experiment may not trust the experimenter and may hold such concerns. It would also be possible for an experimenter to cultivate such concerns (either deliberately or otherwise) via the instructions given to subjects.[12]

When σ is not a product measure, the interpretation is more complicated. As a potential example, imagine a Professor of business strategy who suggests that their students use the positioning of a second hand on a clock face as a randomization device – when students of the same Professor play against each other in a simultaneous game, this will be enough to induce correlated randomizations. It is, therefore, the case that there exists mental Lo-Nash (and, also, Lo-Nash) equilibria that involve correlated outcomes. However, as discussed by Lo (2009), Lo-Nash equilibrium is distinct from Correlated Equilibrium.[13]

4.5 Equilibrium Under Uncertainty

Dow and Werlang (1994) introduce equilibrium under uncertainty, a belief based

[12] Any experiment that intended to investigate such concerns would need to be carefully designed to avoid deceiving the subjects.

[13] Lo actually names his equilibrium concept "Correlated Nash equilibrium" to emphasize the connection.

equilibrium concept for two-player games that identifies only the set of (pure) strategies that an agent believes their opponent will use. While equilibrium under uncertainty is one of the earliest ambiguity averse equilibrium concepts (more modern concepts allow for more sophisticated preferences and belief structures)[14] we focus on it here as it allows for an interesting interpretation of equilibrium when it is translated into the mental equilibrium framework outlined above. Specifically, the equilibrium presented in Dow and Werlang's lone theorem can be interpreted, through the lens of mental equilibrium, as an equilibrium of a normal form game with pre-play communication where pessimistic agents believe that their opponent may deviate from their pre-agreed strategy. Alternatively, it may be modeled as a correlated equilibrium where pessimistic agents believe their opponent's may disregard their signal with some positive probability.

4.5.1 Equilibrium under uncertainty

We focus on the special case where the belief functions take a particular form: $\nu_i(T) = (1-c_i)\pi(T)$ for all $T \in \Sigma'$ such that $T \neq \Omega_i$ where π is an additive probability distribution.[15] Dow and Werlang (1994) also consider a more general class of capacities, but focus on and prove existence for the special case considered here. Belief functions of this form exhibit constant ambiguity aversion, with the level of ambiguity aversion parameterized by the variable c_i, which we shall treat as an exogenous parameter of the game.

Definition 4.4 (Adapted from Dow and Werlang (1994))**.** *Consider a two player game with $N = \{i, j\}$. Suppose that the agents have CEU preferences with respect to a belief function, and that the belief functions are such that $\nu_i(T) = (1-c_i)\pi_j(T)$ for all $T \in \Sigma_i'/\Omega_i$ and $\nu_j(T) = (1-c_j)\pi_i(T)$ for all $T \in \Sigma_j'/\Omega_j$ and $\nu_i(\Omega_i) = \nu_j(\Omega_j) = 1$.*[16]

[14] In particular Eichberger and Kelsey (2000) generalizes Dow and Werlang (1994) to allow for $n > 2$ players.

[15] To be explicit, $\nu(\Omega) = \pi(\Omega) = 1$. This formulation generates a simple set of α^ν: $\alpha_T^\nu = (1-c_i)\pi(T)$ for all singleton $T \in \Sigma'$, $\alpha_\Omega^\nu = c_i$ and $\alpha_T^\nu = 0$ otherwise. These belief functions are the special case of neo-additive capacities introduces in Chateauneuf et al. (2007) restricted to ambiguity aversion.

[16] Note the switch in subscripts between the belief function ν_i (denoting agent i's beliefs) and π_j which, in the sequel, will be interpreted as agent j's mixed strategy.

Then $\{\nu_i, \nu_j\}$ *form an equilibrium under uncertainty iff* $s_i \in \arg\max_{s'_i \in S_i} \int s'_i d\nu_i$ *for all* $s_i \in supp(\pi_j)$ *and* $s_j \in \arg\max_{s'_j \in S_j} \int s'_j d\nu_j$ *for all* $s_j \in supp(\pi_i)$.

Notice that this definition of equilibrium is strictly an equilibrium in beliefs; the equilibrium places no restrictions on strategies, but only restricts the beliefs that each player may hold about their opponent's strategies: beliefs must be mutually consistent in the sense that every strategy that I believe my opponent might play must be optimal for my opponent given their beliefs. An alternative definition of equilibrium under uncertainty, using sub-additive probabilities instead of capacities, is included, along with an example, in Section A.1.2. It is possible, however, as formulated in section 4.5.2, to reformulate and reinterpret an equilibrium under uncertainty as a mixed strategy equilibrium in the mental state space.

4.5.2 Mental equilibrium under uncertainty

A mental equilibrium under uncertainty is a pair of strategies and beliefs over the mental state space such that the strategies are best responses given the beliefs, and the beliefs are consistent with the strategies in the sense that no observable event occurs less often than the beliefs suggest it should occur.

Definition 4.5 (Mental Equilibrium under uncertainty)**.** *A pair* $< \sigma, \alpha >$ *form a Mental equilibrium under uncertainty if*

$$\sigma^{A_j}(T) \geq \sum_{\tau \subseteq T} \alpha_i(\tau) \quad \forall T \in \Sigma_i, \forall i \in N \tag{4.12}$$

and

$$\sigma^{A_i} \in \arg\max_{\hat{\sigma} \in \Delta(A_i)} \sum_{T \in \Sigma_i} u'_i(\hat{\sigma}, T)\alpha_i(T) \quad \forall i \in N \tag{4.13}$$

The following theorem confirms that an equilibrium under uncertainty can be rewritten as a mental equilibrium under uncertainty. We restrict attention to the case of constant ambiguity aversion because of the interesting interpretation of such equilibrium afforded by Definition 4.6 below. Dow and Werlang (1994) also argue that constant ambiguity is the most interesting case because it allows the game to be parameterized by the degree of ambiguity exhibited by the agents. Theorem 4.2 can be generalized to more general belief functions, given an appropriate rewriting of Definition 4.4.

Theorem 4.2. *Suppose that* $\nu_i = (1-c_i)\sigma^{A_j}$ *and* $\nu_j = (1-c_j)\sigma^{A_i}$. *Define* α_i *such that*

$$\alpha_i(T) = \begin{cases} (1-c_i)\sigma^{A_j}(T) & \text{if } T \in A_j \\ c_i & \text{if } T = A_j \\ 0 & \text{otherwise} \end{cases}$$

If $< \nu_i, \nu_j >$ *form an equilibrium under uncertainty, then* $< \sigma, \alpha >$ *form a mental equilibrium under uncertainty. Furthermore, for any* σ'^{A_j} *such that* $\sigma'^{A_j}(a_j) \geq \alpha_i(a_j)$ *for all* $a_j \in A_j$ *and supp* $\sigma'^{A_j} =$ *supp* σ^{A_j}, $< \sigma', \alpha >$ *is also a mental equilibrium under uncertainty.*

If $< \sigma, \alpha >$ *form a mental equilibrium under uncertainty, then* $< \nu_i, \nu_j >$ *form an equilibrium under uncertainty.*

A mental equilibrium under uncertainty can also be interpreted in the same way that we can interpret a mental Lo-Nash equilibrium: agents have access to a pre-play mixing device that they consult before playing a pure strategy, and the construction of the mixing device and the beliefs regarding the opponent's mixing device must satisfy the requirements of the relevant equilibrium. There is also an alternative interpretation of a mental equilibrium under uncertainty.

Consider the following story that supports a mental equilibrium under uncertainty. There are two agents, who are about to play a normal form game with each other. Before the game commences the agents are allowed to communicate with each other, and reach an agreement on which (potentially mixed) strategy they will play.[17] After the pre-play communication each agent retires to their own room to decide their strategy in private. Once in their room, the agents may question the sincerity of their opponents commitment; these doubts may be objectively justified (as in Aumann (1990)) or may be purely subjective. The agent concludes that they can trust their opponent to follow the agreement with probability $1-c$ and that their opponent will deviate from the agreement with probability c.[18] Furthermore,

[17] In the case of a mixed strategy, the agents are agreeing on the mixing device that they will consult prior to selecting their pure strategy.

[18] We treat the parameter c as exogenous, but a richer model may seek to endogenize it.

the agent is pessimistic in the sense of Maxmin or Choquet Expected Utility.

Some natural questions arise: Under what conditions will the agent be willing to play their agreed strategy? Which pre-play commitments are self-enforcing in this environment? The answers are straightforward: if the pre-play commitments form a mental equilibrium under uncertainty then they will be self-enforcing and the agents will be willing to play their agreed strategy.

Definition 4.6. *If* $<\sigma,\alpha>$ *form a mental equilibrium under uncertainty and*

$$\alpha_i(T) = \begin{cases} (1-c_i)\sigma^{A_j}(T) & \text{if } T \in A_j \\ c_i & \text{if } T = A_j \\ 0 & \text{otherwise} \end{cases}$$

then $<\sigma,\alpha>$ *form a* c-robust agreement.

When $c = 0$, we recover a common justification for Nash equilibrium: if agents believe that their opponent will follow their pre-play commitment with probability 1, then the strategies are self enforcing iff the strategies form a Nash equilibrium. We illustrate c-robust agreements by presenting an example.

	X	Y
A	3,3	0,2
B	2,0	1,1

Figure 4.5: A normalized stag hunt game.

Consider the stag hunt game presented in figure 4.5 and the associated symmetric c-robust agreements:

- If $c > \frac{1}{2}$ there is a unique robust agreement: $\{B,Y\}$.
- If $c \leq \frac{1}{2}$ there are three robust agreements: $\{B,Y\}$, $\{A,X\}$ and the mixed robust agreement $\sigma(X) = \sigma(A) = \frac{1}{2(1-c)}$.

Notice that in the mixed robust agreement the weight placed on the 'cooperative' strategy increases as the parameter c increases. Intuitively, in order for an agent to be willing to mix (i.e. be indifferent between their strategies) against a relatively untrustworthy opponent the agreement must be weighted more heavily towards ‘cooperation’. As the level of trust decreases, we find a point ($c = \frac{1}{2}$) where cooperation is no longer sustainable.

4.6 Extensive Form Games

Modeling extensive form games with ambiguity averse agents is a difficult task given that ambiguity averse agents do not, in general, have dynamically consistent preferences. One key implication of this is that Kuhn’s theorem (establishing the equivalence of behavioral and mixed strategies) breaks down in games with ambiguity averse agents, thereby forcing the modeler to make difficult choices regarding the structure of allowable strategies (Aryal and Stauber, 2014). The natural structure of a mental equilibrium provides a possible solution: because preferences have a SEU representation in the mental state space Bayesian updating will produce dynamically consistent preferences *in the mental state space*. The only remaining challenge is to determine when dynamically consistent preferences in the mental state space map into dynamically consistent preferences in the observable game.

We begin by stating the formal requirements for dynamically consistent observable preferences, and then discuss the intuition behind the result.

Consider the standard formulation for an extensive form game $< I, A, X, n, H, \iota, u >$, where I is the finite set of players, X is the (partially ordered) finite set of nodes (with x_0 denoting the initial node), $A(x)$ denotes the actions available at each node, $n(x,a)$ denotes the successor to node x when action a is played, $\iota(x)$ denotes the player to move at node x, $H(x)$ is an information partition such that if $x' \in H(x)$ then $\iota(x) = \iota(x')$, $A(x) = A(x')$ and $H(x') = H(x)$, and u is a set of preferences over outcomes. We write A_i, H_i and u_i to denote the actions, histories and (Bernoulli) utility function for player i. X_i denotes the nodes at which player i has the move.

We map the extensive form game into a mental extensive form game as follows. As is standard, $s_i : H_i \mapsto A_i$ denotes a strategy for player i. The set of player i strategies is S_i, $S = \times_{i \in I} S_i$ and $S_{-i} = \times_{j \neq i} S_j$. Player i’s mental state space is given by $\Sigma_i = 2^{S_{-i}}/\emptyset$.[19] Beliefs over the mental state space are denoted by α_i that are formulated from beliefs over the observable game according to Result 4.1. Result 4.2 therefore implies that preferences have a SEU representation over the mental

[19] As for a normal form game, the mental state space is equivalent to the observable event space.

state space. By assumption, beliefs over the mental state space are updated via Bayesian updating.

Given X, n and H_i we write $\psi_i : X_i \bigcup x_0 \mapsto S_{-i}$ to denote the set of opponents strategies that are consistent with play reaching node $x \in X_i$. Note that $\psi_i(x_0) = S_{-i}$ for all i because the initial node is consistent with any feasible opponent strategy.

Assuming perfect recall, it is possible to build an ordered filtration of S_{-i} using ψ_i. First, note that the successor operator n defines a partial order over all nodes; call this partial order $\succeq_n$. Define the total order $\succeq_{\psi_i}$ to be an arbitrarily chosen total order such that $x \succeq_n x'$ implies $x \succeq_{\psi_i} x'$ for all $x, x' \in \{X_i \bigcup x_0\}$. The total order $\succeq_{\psi_i}$ allows us to order the set $\{X_i \bigcup x_0\}$ as $\{x_0, x_1, \ldots x_n\}$. We are now ready to define the filtration Ψ_i as follows:

Definition 4.7. *Define $\Psi_i = \{\Psi_{x_0}, \Psi_{x_1}, \ldots, \Psi_{x_n}\}$ to be a set of partitions. The first partition is defined to be $\Psi_{x_0} = \Sigma_i$. Subsequent partitions are defined recursively. Suppose that the m-th partition has k elements, so that it may be written as $\Psi_{x_m} = \{\Psi^1_{x_m}, ..., \Psi^k_{x_m}\}$. Then the $m+1$-th partition is given by $\Psi_{x_{m+1}} = \{\Psi^1_{x_m}/\psi(x_{m+1}), ..., \Psi^k_{x_m}/\psi(x_{m+1}), \psi(x_{m+1})\}$, after discarding any $\Psi^j_{x_m}/\psi(x_{m+1})$ that are empty. Finally, if $\Psi_x = \Psi_{x'}$ for any $x, x' \in \{X_i \bigcup x_0\}$ then delete either Ψ_x or $\Psi_{x'}$ from the set.*[20]

Definition 4.8 (Filtration)**.** $\Psi_i = \{\Psi_{x_0}, \Psi_{x_1}, \ldots, \Psi_{x_n}\}$ *is a filtration if each Ψ_{x_m} is a partition of A_{-i} and for each $\Psi^k_{x_m} \in \Psi_{x_m}$ there exists a $\Psi^j_{x'_m} \in \Psi_{x'_m}$ such that $\Psi^k_{x_m} \subseteq \Psi^j_{x'_m}$ for all m, m' with $m > m'$.* [21]

Lemma 4.3. *If the extensive form game satisfies perfect recall then Ψ_i is a filtration for all i.*

We are now ready to state the main result of this section, which provides conditions on the agents' beliefs such that the agent will exhibit dynamically consistent preferences. The proof is borrowed directly from Eichberger et al. (2005). The novelty of the result rests on using Result 4.1 as the vehicle that allows us

[20]Note that, by construction, $\Psi_x = \Psi_{x'}$ whenever $x' \in h(x)$.

[21]Note that the term filtration is used in slightly different forms in different literature. It is possible to recover the most common definition (typically used in the mathematics literature) by completing each Ψ_{x_j} into a σ-algebra.

to translate the Eichberger et al. (2005) result from an individual decision making framework into a game theoretic setting.

Theorem 4.4. *Consider an extensive form game in the mental state space. For each player, write the final stage of their filtration as* $\Psi^i_{x_n} = \{\Psi^i_1, ..., \Psi^i_k\}$. *If*

$$\sum_{\kappa=1}^{k} \sum_{s \subseteq \Psi^i_\kappa} \alpha_i(s) = 1 \tag{4.14}$$

and the agent uses Bayesian updating to update α *then agent i has dynamically consistent preferences.*

Furthermore, if $\Psi^i_{x_n}$ *is non-trivial (i.e.* $|\Psi^i_\kappa| \geq 2$ *for all* κ*) and the agent believes all nodes are reached with positive probability then*

$$\sum_{\kappa=1}^{k} \sum_{s \subseteq \Psi^i_\kappa} \alpha_i(s) = 1$$

if and only if agent i has dynamically consistent preferences.

Proof. See Theroem 2.1 in Eichberger et al. (2005). Note that the condition $\sum_{\kappa=1}^{k} \sum_{s \subseteq \Psi^i_\kappa} \alpha_i(s) = 1$ is equivalent to their condition $\sum_{E \in \mathscr{E}} \nu(E) = 1$. □

The practical implication of equation 4.14 is that the agent may only be uncertain (i.e. have non-additive beliefs) regarding actions occur after the agents last decision node on that branch, or regarding which node they are at within an information set when there are no future nodes at which the agent has the move where he learns which node he was previously at within the information set.[22] This is a reasonably restrictive condition, although as the theorem states this condition is tight in a fully mixed equilibrium for dynamically consistent agents who are never the last agent to act.

In a general framework, dynamic consistency will imply ambiguity neutrality (via the tight connection between dynamic consistency and Savage's P2 postulate). That is, if we wish to impose dynamic consistency across all possible extensive form games *at the same time* then the agents must necessarily be expected utility

[22] We conjecture that this condition implies that the agent's beliefs must satisfy the rectangularity condition of Epstein and Schneider (2003).

maximizers. Here, however, we take a weaker stance of enforcing dynamic consistency for a single, predefined, game tree. Any mental equilibrium of the normal form game derived from the game tree such that equilibrium beliefs satisfy Equation 4.14 will be dynamically consistent. Checking whether equation 4.14 holds for a given equilibrium is straightforward: the mental state space is already specified as part of the equilibrium so all that remains is to find the finest partition of the filtration for each subject, which can usually be read directly off the game tree.

It is also straightforward to define a sub-game perfect refinement of mental Lo-Nash equilibrium.

Definition 4.9 (Sub-game perfect mental Lo-Nash equilibrium)**.** *A pair $<\sigma,\alpha>$ is a sub-game perfect mental Lo-Nash equilibrium if it forms a Lo-Nash equilibrium on every sub-game when restrictions of both σ and α to any sub-game are found via Bayesian updating.*

4.7 Computability

One natural concern with the use of ambiguity averse equilibrium concepts is that they are difficult to calculate. In this section we demonstrate that the computational complexity of the mental equilibrium concepts presented above are similar to the computational complexity of finding a Nash equilibrium. The fact that mental equilibrium consist of linear utility functions is the key feature that makes this possible. Importantly, the mental equilibrium approach provides an algorithmic methodology for solving for the equilibrium of either Dow and Werlang (1994) or Lo (2009): first, find the mental state space and associated payoffs and then apply the same techniques that are used for computing Nash equilibrium.

The textbook approach, which we present here following Shoham and Leyton-Brown (2009), for computationally searching for Nash equilibrium involves formulating the problem as a linear complementarity problem (LCP). The LCP can then be solved as a feasibility program. The LCP for a two player game with finite action sets is given by:

$$\sum_{a_i \in A_i} u_j(a_j, a_i)\sigma(a_i) + r(a_j) = U_j \qquad \forall a_j \in A_j \tag{4.15}$$

$$\sum_{a_j \in A_j} u_i(a_i, a_j)\sigma(a_j) + r(a_i) = U_i \qquad \forall a_i \in A_i \tag{4.16}$$

$$\sum_{a_i \in A_i} \sigma(a_i) = 1, \sum_{a_j \in A_j} \sigma(a_i) = 1 \tag{4.17}$$

$$\sigma(a_i) \geq 0, \sigma(a_j) \geq 0 \qquad \forall a_i \in A_i, a_j \in A_j \tag{4.18}$$

$$r(a_i) \geq 0, r(a_j) \geq 0 \qquad \forall a_i \in A_i, a_j \in A_j \tag{4.19}$$

$$r(a_i)\sigma(a_i) = 0, r(a_j)\sigma(a_j) = 0 \qquad \forall a_i \in A_i, a_j \in A_j \tag{4.20}$$

A solution to the LCP consists of a set of σ, r and U such that the σ form a Nash equilibrium, the U denote the equilibrium utility for each player and the r a set of slack parameters that are set to 0 for any strategy that is used in the equilibrium. The LCP can be solved using, for example, the Lemke-Howson algorithm.

We can also write a mental equilibrium under uncertainty as a LCP. We present here the case considered in Theorem 4.2:

$$\sum_{a_j \in A_j} u'_i(a_i, a_j)\alpha_i(a_j) + u'_i(a_i, A_j)\alpha(A_j) + r(a_i) = U_i \qquad \forall a_i \in A_i \tag{4.21}$$

$$\sum_{a_i \in A_i} u'_j(a_j, a_i)\alpha_j(a_i) + u'_j(a_j, A_i)\alpha(A_i) + r(a_j) = U_j \qquad \forall a_j \in A_j \tag{4.22}$$

$$\sum_{T \in A_j} \alpha_i(T) = c_i, \sum_{T \in A_i} \alpha_j(T) = c_j \tag{4.23}$$

$$\alpha_i(A_j) = 1 - c_i, \alpha_j(A_i) = 1 - c_j \tag{4.24}$$

$$\alpha_i(T_j) \geq 0, \alpha_j(T_i) \geq 0 \qquad \forall T_j \in A_j, \forall T_i \in A_i \tag{4.25}$$

$$r(a_i) \geq 0, r(a_j) \geq 0 \qquad \forall a_i \in A_i, a_j \in A_j \tag{4.26}$$

$$\alpha(a_i)r(a_i) = 0, \alpha(a_j)r(a_j) = 0 \qquad \forall a_i \in A_i, a_j \in A_j \tag{4.27}$$

This LCP solves for the beliefs of the agents (the α), but does not directly provide us with the strategies (the σ). However, the α bound a convex hull of feasible σ and it is possible to write down a feasible σ directly from the output of the LCP.

When we view the values (c_i, c_j) as fixed parameters of the game, then the (c_i, c_j) parameters and equation 4.24 drop out and the mental equilibrium under uncertainty LCP is exactly as computationally difficult as the Nash equilibrium LCP. If, on the other hand, we treat (c_i, c_j) as variables then we have 1 additional parameter, and no additional equations, to calculate for each agent. Given that increasing the strategy set for an agent in the Nash equilibrium LCP introduces both an additional parameter and an additional equation, we conclude that solving a $k \times k$ mental equilibrium under uncertainty is harder than solving a $k \times k$ Nash equilibrium but easier than solving a $(k+1) \times (k+1)$ Nash equilibrium.

Mental Lo-Nash equilibrium is more computationally intensive because of the stricter requirements on the supports of the strategies and the larger mental state space. The mental Lo-Nash LCP is given by:

$$\sum_{T \in \Sigma_i} u'_i(a_i, T)\alpha_i(T) + r(a_i) = U_i \qquad \forall a_i \in A_i \tag{4.28}$$

$$\sum_{T \in \Sigma_j} u'_j(a_j, T)\alpha_j(T) + r(a_j) = U_j \qquad \forall a_j \in A_j \tag{4.29}$$

$$\sum_{T \in \Sigma_i} \alpha_i(T) = 1, \sum_{T \in \Sigma_j} \alpha_j(T) = 1 \tag{4.30}$$

$$\sum_{T: a_j \in T} \alpha_i(T) = I(a_j) \qquad \forall a_j \in A_j \tag{4.31}$$

$$\sum_{T: a_i \in T} \alpha_j(T) = I(a_i) \qquad \forall a_i \in A_i \tag{4.32}$$

$$I(a_j) \leq 1 \qquad \forall a_j \in A_j \tag{4.33}$$

$$I(a_i) \leq 1 \qquad \forall a_i \in A_i \tag{4.34}$$

$$\alpha_i(T_j) \geq 0, \alpha_j(T_i) \geq 0 \qquad \forall T_j \in \Sigma_i, \forall T_i \in \Sigma_j \tag{4.35}$$

$$r(a_i) \geq 0, r(a_j) \geq 0 \qquad \forall a_i \in A_i, a_j \in A_j \tag{4.36}$$

$$(1 - I(a_i))r(a_i) = 0, (1 - I(a_j))r(a_j) = 0 \qquad \forall a_i \in A_i, a_j \in A_j \tag{4.37}$$

The first two equations contain 2^k α parameters compared to k parameters in the Nash case, and we have also introduced the I parameters which impose a sub-constraint on the α parameters in addition to the standard requirements that the α form a probability distribution. While the problem has substantially more parameters, it is not is any sense fundamentally different from the Nash equilibrium LCP.

4.8 Conclusions

This chapter presents a methodology that extends static ambiguity averse equilibrium concepts that only allows for pure strategies to allow for mixed strategies, as well as be extended into a dynamically consistent ambiguity averse equilibrium concept for extensive form games. The technique is restricted to preference structures which allow only ambiguity aversion.[23] Furthermore, an application of the technique to Dow and Werlang's (1994) equilibrium under uncertainty produces a model of partially robust agreements.

The theoretical interest in the structure and interpretation of mixed strategy equilibrium for agents with uncertainty averse preferences is readily apparent. Unlike standard SEU agents, ambiguity averse agents are not necessarily indifferent between a mixed strategy and the pure strategy supports of the mixed strategy. Considering this, some of the standard interpretations of mixed strategies (via purification arguments or population interpretations of mixed strategies) may not be appropriate in the context of ambiguity averse agents; indeed, a majority of ambiguity averse equilibrium concepts explicitly restrict their analyses to pure strategies. In this context, the methodology introduced here can be viewed as an equilibrium-preserving interpretation of mixed strategies: when agents are ambiguity averse and have access to a mixing device that resolves before the strategic interaction occurs then the agents will, in equilibrium, be indifferent between using the mixing device or not. In this fashion the population interpretation of mixing can be restored without the need to explicitly assume away the existence of mixing devices.

A commonly held piece of folk wisdom is that people play pure strategies, and do not use mixed strategies. Experimental evidence appears to contradict this,

[23] Preference structures that allow only for ambiguity seeking can also easily be accommodated by defining the dual of $e_T(A)$ in theorem 4.1. Dealing with preferences structures that allow for either ambiguity aversion or ambiguity seeking is a significantly more difficult proposition, although it may be tractable under some (possibly unpalatable) restrictions.

however, and indicates that a majority of subjects make use of randomization devices when available in zero-sum games (see Schachat (2002) for example). Given this evidence that mixed strategies are used when available, extending an equilibrium concept to allow for mixed strategies while preserving other features of the equilibrium represents an improvement in the descriptive power of the equilibrium concept.

The ubiquitous rise of the internet and digital communications has lead to a proliferation of pre-programmed bots that are increasingly interacting in strategic situations.[24] In some cases bots that play pure strategies are exploitable, thereby necessitating the use of randomization. In cases where the programmer has expected utility preferences over the outcomes produced by their bot there is no conceptual problem with evaluating the (ex-ante) expected performance of differently specified programs. However, in cases where the programmer has ambiguity averse preferences the evaluation of such programs is a conceptually challenging task. The theory presented in this chapter presents one potential way forward. Given that current ambiguity averse equilibrium concepts that allow for mixed strategies are observationally equivalent to Nash equilibrium for 2-player games, a theory that allows for mixed strategy equilibrium that is distinct from Nash equilibrium is necessary for ambiguity averse preferences to be a positive descriptor of behaviour in games.

[24] Online pricing algorithms for concerts and sporting events interacting with automated purchasing algorithms owned by scalpers are but one example.

Chapter 5

Conclusion

In this dissertation, I have contributed to the literature on strategic uncertainty in games with non-Subjective Expected Utility agents.

Chapter 2 provides the most thorough and convincing evidence to date that ambiguity preferences influence behaviour in normal form games. It also demonstrates that the usual assumption of common knowledge of preferences does not hold among a typical experimental population; even the much weaker mutual knowledge of preferences fails to obtain. However, in an experimental intervention that restored mutual knowledge of preferences, solution concepts that allow for ambiguity aversion perform remarkably well.

Chapter 2 introduces three experimental innovations. First, it introduces a methodology for measuring risk and ambiguity preferences, and beliefs over opponent's preferences, using normal form games. Second, it identifies and separates the effects of risk aversion and ambiguity aversion in a normal form game. Third, it separates the effect of preferences on behaviour from the effects of beliefs over opponent's preferences in a normal form game.

Chapter 3 investigates the role of naturally occurring reaction lags in determining how subjects perceive and respond to strategic uncertainty. One key aspect of the results is that when reaction lags are either very small or very large then behaviour approximates predictions that are formed from models without strategic uncertainty. This implies that in some settings with very large reaction lags discrete time Nash equilibrium is a good approximation of behaviour and in other settings with very small reaction lags continuous time Nash equilibrium is a good approximation of behaviour. We speculate that as technology improves and communication and reaction lags shrink, continuous time models of behaviour will become increasingly important in applications.

The original contributions of Chapter 3 have two main dimensions. First, we introduce a new experimental technique that allows us to implement the premises of the theory of continuous time games directly in the laboratory. Second, we identify that theories of strategic uncertainty, particularly minimax regret theory, explain the data collected in our experiment remarkably well. We further demonstrate that minimax regret theory can also explain data from similar previous experiments, and that a theoretical application of minimax regret theory generates plausible predictions across a broad class of dilemma-like games in inertial continuous time.

Chapter 4 provides a theoretical motivation for interpreting experiments with ambiguity averse agents using solution concepts that restrict the action space to pure strategies. The Chapter establishes that, for two key equilibrium concepts, allowing pre-play mixing is equivalent to restricting to pure strategies. The theoretical results of Chapter 4 are essential for interpreting the experimental results of Chapter 2. The key insight of the Chapter lies in identifying that the main result of Gilboa and Schmeidler (1994) can be profitably applied to games with ambiguity averse agents: straightforward applications of Gilboa and Schmeidler (1994) generate novel results relating to mixed strategies, dynamic consistency and computability in games with ambiguity averse agents.

Bibliography

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. → pages 28

F.J. Anscombe and R.J. Aumann. A definition of subjective probability. *The Annals of Mathematical Statistics*, 34(1):199–205, 1963. → pages 78, 82, 84

K.J Arrow and L. Hurwicz. An optimality criterion for decision-making under ignorance. In J.L. Ford C.F. Carter, editor, *Uncertainty and Expectations in Economics: Essays in Honour of G.L.S. Shackle*. Basil Blackwell, Oxford, 1972. → pages 63, 175

Gaurab Aryal and Ronald Stauber. A note on Kuhn's theorem with ambiguity averse players. *Economics Letters*, 125(1):110–114, 2014. → pages 93

Robert J. Aumann. Correlated equilibrium as an expression of bayesian rationality. *Econometrica*, 55(1):1–18, 1987. → pages 83

Robert J. Aumann. *Economic Decision-Making: Games, Econometrics and Optimisation*, chapter Nash Equilibria are not Self-Enforcing, pages 201–206. North Holland, 1990. → pages 91

Yaron Azrieli and Roee Teper. Uncertainty aversion and equilibrium existence in games with incomplete information. *Games and Economic Behavior*, 73: 310–317, 2011. → pages 114

Yaron Azrieli, Christopher P. Chambers, and P. J. Healy. Incentives in experiments: A theoretical analysis. Working Paper, 2014. → pages 15, 84, 129

Sophie Bade. Ambiguous act equilibria. *Games and Economic Behavior*, 71: 246–260, 2011. → pages 113, 114

Aurelien Baillon, Yoram Halevy, and Li Chen. Experimental elicitation of ambiguity attitude using the random incentive system. Mimeo, 2014. → pages 15, 84, 129

P. Battigalli, S. Cerreia-Vioglio, F. Maccheroni, and M. Marinacci. A note on comparative ambiguity aversion and justifiability. IGIER Working Paper Series, September 2015. → pages 12

Dirk Bergemann and Karl H. Schlag. Pricing without priors. *Journal of the European Economic Association*, 6(2/3):560–569, 2008. → pages 177

Dirk Bergemann and Karl H. Schlag. Robust monopoly pricing. *Journal of Economic Theory*, 146:2527–2543, 2011. → pages 177

James Bergin and W. Bentley MacLeod. Continuous time repeated games. *International Economic Review*, 34(1):21–37, 1993. → pages 2, 41, 42, 43, 44, 49, 50, 159, 160, 164, 165, 166, 167

B. Douglas Bernheim. Rationalizable strategic behavior. *Econometrica*, 52(4): 1007–1028, 1984. → pages 11

S. Berninghaus, K.-M. Ehrhart, and M. Ott. A network experiment in continuous time: The influence of link costs. *Experimental Economics*, 9:237–251, 2006. → pages 42

Maria Bigoni, Marco Casari, Andrzej Skrzypacz, and Giancarlo Spagnolo. Time horizon and cooperation in continuous time. *Econometrica*, 83(2):587–616, 2015. → pages 42

Antoine Billot and Vassili Vergopoulos. Expected utility without parsimony. CES Working Paper, March 2014. → pages 77, 78

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. → pages 28

Adam Brandenburger, Amanda Friedenberg, and H. Jerome Keisler. Admissibility in games. *Econometrica*, 76(2):307–352, 2008. → pages 56, 61

Colin F. Camerer and Risto Karjalainen. *Models and Experiments in Risk and Rationality*, chapter Ambiguity Aversion and Non-additive Beliefs in Non-cooperative Games: Experimental evidence, pages 325–358. Springer Netherlands, 1994. → pages 9, 130

H Carlsson and Eric van Damme. Global games and equilibrium selection. *Econometrica*, 61:989–1018, 1993. → pages 3

Alain Chateauneuf, Jurgen Eichberger, and Simon Grant. Choice under uncertainty with the best and worst in mind: Neo-additive capacities. *Journal of Economic Theory*, 137(1):538–576, 2007. → pages 89

Soo Hong Chew, Mark Ratchford, and Jacob S. Sagi. You need to recognize ambiguity to avoid it. 2013. → pages 15, 21, 30, 123, 132, 133

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. → pages 28

Pedro Dal Bo and Guillaume Frechette. The evolution of cooperation in infinitely repeated games: Experimental evidence. *American Economic Review*, 101: 411–429, February 2011. → pages 45, 61

Pedro Dal Bo and Guillaume Frechette. On the determinants of cooperation in infinitely repeated games: A survey. Mimeo, March 2016. → pages 45, 61

Mark Dean and Pietro Ortoleva. Is it all connected? A testing ground for unified theories of behavioral economics phenomena. Mimeo, March 2014. → pages 28

C. Deck and N. Nikiforakis. Perfect and imperfect real-time monitoring in a minimum-effort game. *Experimental Economics*, 15:71–88, 2012. → pages 42

Carmela Di Mauro and Massimo F. Castro. Kindness, confusion, or …ambiguity? *Experimental Economics*, 14:611–633, 2011. → pages 130

Carmela Di Mauro and Anna Maffioletti. Attitudes to risk and attitudes to uncertainty: experimental evidence. *Applied Economics*, 36(4):357–372, 2004. → pages 28

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. → pages 3, 5, 10, 12, 14, 73, 74, 75, 88, 89, 90, 96, 99, 113, 114, 119, 120

John Duffy and Jack Ochs. Equilibrium selection in static and dynamic entry games. *Games and Economic Behavior*, 76:97–116, 2012. → pages 53

Jurgen Eichberger and Da Kelsey. Uncertainty aversion and preference for randomization. *Journal of Economic Theory*, 71:31–43, 1996. → pages 80

Jurgen Eichberger and David Kelsey. Non-additive beliefs and strategic equilibria. *Games and Economic Behavior*, 30:183–215, 2000. → pages 5, 11, 14, 73, 74, 89, 114

Jurgen Eichberger and David Kelsey. Are the treasures of game theory ambiguous? *Economic Theory*, 48:313–339, 2011. → pages 130, 131

Jurgen Eichberger and David Kelsey. Optimism and pessimism in games. *International Economic Review*, 55(2):483–505, 2014. → pages 114

Jurgen Eichberger, Simon Grant, and David Kelsey. CEU preferences and dynamic consistency. *Mathematical Sociel Sciences*, 49:143–151, 2005. → pages 94, 95

Jurgen Eichberger, David Kelsey, and Burkhard C. Schipper. Granny versus game theorist: Ambiguity in experimental games. *Theory and Decision*, 64:333–362, 2008. → pages 9, 130, 131

Jurgen Eichberger, Simon Grant, and David Kelsey. Randomization and dynamic consistency. Economic Theory (Forthcoming), 2016. → pages 78, 82, 84, 128, 129

Daniel Ellsberg. Risk, ambiguity and the Savage axioms. *Quarterly Journal of Economics*, 75(4):643–669, 1961. → pages 9

Matthew Embrey, Guillaume R. Frechette, and Steven F. Lehrer. Bargaining and reputation: An experiment on bargaining in the presence of behavioural types. *Review of Economic Studies*, 82:608–631, 2015. → pages 42

Matthew Embrey, Guillaume R. Frechette, and Sevgi Yuksel. Cooperation in the finitely repeated prisoner's dilemma. Mimeo, 2016. → pages 45, 61

Larry G. Epstein. Preference, rationalizability and equilibrium. *Journal of Economic Theory*, 73:1–29, 1997. → pages 5, 11, 12, 114, 134

Larry G. Epstein and Yoram Halevy. No two experiments are identical. Mimeo, December 2014. → pages 16, 124

Larry G. Epstein and Martin Schneider. Recursive multiple-priors. *Journal of Economic Theory*, 113:1–31, 2003. → pages 95

Ignacio Esponda and Emanuel Vespa. Hypothetical thinking and information extraction in the laboratory. *American Economic Journal: Microeconomics*, 6 (4):180–202, 2014. → pages 39

Piotr Evdokimov and David Rahman. Cooperative institutions. Mimeo, August 2014. → pages 42

Peter C. Fishburn. *Utility Theory for Decision Making*. John Wiley & Sons, 1970. → pages 78

Daniel Friedman and Ryan Oprea. A continuous dilemma. *American Economic Review*, 102(1):337–363, 2012. → pages xi, 42, 178, 182

Itzhak Gilboa. Expected utility with purely subjective non-additive probabilities. *Journal of Mathematical Economics*, 16:65–88, 1987. → pages 79

Itzhak Gilboa and David Schmeidler. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18:141–153, 1989. → pages 7, 11, 29, 64, 74, 75, 85, 115, 176

Itzhak Gilboa and David Schmeidler. Additive representations of non-additive measures and the Choquet integral. *Annals of Operations Research*, 52:43–65, 1994. → pages 74, 75, 76, 77, 102, 119, 192

Jacob K. Goeree and Charles A. Holt. Ten little treasures of game theory and ten intuitive contradictions. *American Economic Review*, 95(5):1402–1422, 2001. → pages 131

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. → pages 29

Simon Grant, Idione Meneghel, and Rabee Tourky. Savage games. *Theoretical Economics*, Forthcoming. → pages 13, 40, 114

Ben Greiner. An Online Recruitment System for Economic Experiments. MPRA Paper 13513, University Library of Munich, Germany, 2004. URL http://ideas.repec.org/p/pra/mprapa/13513.html. → pages 55

Ben Greiner. Subject pool recruitment procedures: Organizing experiments with orsee. *Journal of the Economic Science Association*, 1(1):114–125, July 2015. → pages 20

Ebbe Groes, Hans Jorgen Jacobson, Birgitte Sloth, and Torben Tranaes. Nash equilibrium with lower probabilities. *Theory and Decision*, 44:37–66, 1998. → pages 83, 114

Yoram Halevy. Ellsberg revisited: An experimental study. *Econometrica*, 75(2): 503–536, 2007. → pages 82, 131

Yoram Halevy and Vincent Feltkamp. A Bayesian approach to uncertainty aversion. *Review of Economic Studies*, 72:449–466, 2005. → pages 29

Joseph Y. Halpern and Rafael Pass. Iterated regret minimiziation: A new solution concept. *Games and Economic Behavior*, 74:184–207, 2012. → pages 63, 175, 176, 177, 178

P. J. Healy. Epistemic foundations for the failure of Nash Equilibrium. Mimeo, 2013. → pages 38, 39

Chip Heath and Amos Tversky. Preference and belief: Ambiguity and competence in choice under uncertainty. *Journal of Risk and Uncertainty*, 4: 5–28, 1991. → pages 124

Frank Heinemannn, Rosemarie Nagel, and Peter Ockenfels. Measuring strategic uncertainty in coordination games. *Review of Economic Studies*, 76:181–221, 2009. → pages 10, 130

Charles A. Holt and Susan K. Laury. *Handbook of the Economics of Risk and Uncertainty*, chapter Assessment and Estimation of Risk Preferences, pages 135–201. Elsevier, 2014. → pages 125

Asen Ivanov. Attitudes to ambiguity in one-shot normal-form games: An experimental study. *Games and Economic Behavior*, 71:366–394, 2011. → pages 10, 130, 131, 132

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. → pages 129

E.L. Kaplan and Paul Meier. Nonparametric estimation from incomplete observations. *Journal of the American Statistical Association*, 53(282): 457–481, June 1958. → pages 57, 58

David Kelsey and Sara le Roux. An experimental study on the effect of ambiguity in a coordination game. *Theory and Decision*, Forthcoming, 2015a. → pages 9, 29, 124, 130, 132, 133

David Kelsey and Sara le Roux. Dragon slaying with ambiguity: Theory and experiments. Mimeo, 2015b. → pages 130, 133

John M. Keynes. *A Treatise on Probability*. MacMillan & Co., 1921. → pages 9

Peter Klibanoff. Uncertainty, decision and normal form games. Mimeo, July 1996. → pages 74, 82, 113, 114

Peter Klibanoff, Massimo Marinacci, and Sujoy Mukerji. A smooth model of decision making under ambiguity. *Econometrica*, 73(6):1849–1892, 2005. → pages 12, 79

Terri Kneeland. Identifying higher order rationality. *Econometrica*, 83(5): 2605–2079, 2015. → pages 38, 39

Frank H. Knight. *Risk, Uncertainty and Profit*. Hart, Schaffner & Marx; Boston: Houghton Mifflin Company, The Riverside Press, Cambridge, 1921. → pages 9

David M. Kreps. *Notes on the Theory of Choice*. Westview Press, 1988. → pages 128

Ehud Lehrer. Partially-specified probabilities: Decisions and games. Mimeo, 2011. → pages 113, 114

Peter B. Linhart and Roy Radner. Minimax-regret strategies for bargaining over several variables. *Journal of Economic Theory*, 48:152–178, 1989. → pages 177

Kin Chung Lo. Equilibrium in beliefs under uncertainty. *Journal of Economic Theory*, 71:443–484, 1996. → pages 3, 74, 82, 113, 114

Kin Chung Lo. Correlated Nash Equilibrium. *Journal of Economic Theory*, 144: 722–743, 2009. → pages 5, 10, 12, 14, 63, 73, 74, 75, 85, 88, 96, 113, 114, 115, 176, 177

Massimo Marinacci. Ambiguous games. *Games and Economic Behavior*, 31: 191–219, 2000. → pages 114

John W. Milnor. Games against nature. In Robert M. Thrall, Clyde H. Coombs, and R. L. Davis, editors, *Decision Processes*. Wiley, 1954. → pages 45, 64, 175, 176, 177

Stephen Morris and Huyn Song Shin. *Advances in Economics and Econometrics: Theory and Applications*, chapter Global games: Theory and applications, pages 56–114. Cambridge University Press, 2003. → pages 3, 64

Ryan O. Murphy, Amnon Rapoport, and James E. Parco. The breakdwon of cooperation in iterative real-time trust dilemmas. *Experimental Economics*, 9 (2):147–166, 2006. → pages 42

Ryan Oprea, K. Henwood, and Daniel Friedman. Seperating the hawks from the doves: Evidence from continuous time laboratory games. *Journal of Economic Theory*, 146:2206–2225, 2011. → pages 42

Ryan Oprea, Gary Charness, and Daniel Friedman. Continuous time and communication in a public-goods experiment. *Journal of Economic Behaviour and Organization*, 108:212–223, 2014. → pages 42

David G. Pearce. Rationalizable strategic behavior and the problem of perfection. *Econometrica*, 52(4):1029–1050, 1984. → pages 11

J. Pettit, Jeremy Hewitt, and Ryan Oprea. Redwood: Software for graphical, browser-based experiments in discrete and continuous time. mimeo, 2015. → pages 20, 52

Briony D. Pulford and M Colman, Andrew. Ambiguous games: Evidence for strategic ambiguity aversion. *Quarterly Journal of Experimental Physiology*, 60(8):1083–1100, 2007. → pages 130, 131

Roy Radner. Collusive behavior in noncooperative epsilon-equilibria of oligopolies with long but finite lives. *Journal of Economic Theory*, 22: 136–154, 1980. → pages 168

Ludovic Renou and Karl H. Schlag. Minimax regret and strategic uncertainty. *Journal of Economics Theory*, 145:264–286, 2010. → pages 63, 177

Frank Riedel and Linda Sass. Ellsberg games. *Theory*, 76(4):469–509, 2014. → pages 114

Jorn Rothe. Uncertainty aversion and equilibrium in normal form games. In *Contributions to Game Theory and Management: The Third International Conference Game Theory and Management*, volume III, 2010. → pages 113, 114

Kota Saito. Preferences for flexibility and randomization under uncertainty. *American Economic Review*, 105(3):1246–1271, 2015. → pages 85

Rakesh Sarin and Peter P. Wakker. A simple axiomatization of nonadditive expected utility. *Econometrica*, 60(6):1255–1272, 1992. → pages 79

Leonard J. Savage. The theory of statistical decision. *Journal of the American Statistical Association*, 46(253):55–67, 1951. → pages 64, 177

Leonard J. Savage. *The Foundations of Statistics*. Wiley, New York, 1954. → pages ii, 78, 79, 82, 84

Jason M. Schachat. Mixed strategy play and the minimax hypothesis. *Journal of Economic Theory*, 104:189–226, 2002. → pages 100

David Schmeidler. Subjective probability and expected utility without additivity. *Econometrica*, 57(3):571–587, 1989. → pages 12, 29, 74, 75, 76, 78, 79

Kyoungwon Seo. Ambiguity and second-order belief. *Econometrica*, 77(5): 1575–1605, 2009. → pages 78, 82, 84

Yoav Shoham and Kevin Leyton-Brown. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press, 2009. → pages 96

Leo K. Simon and Maxwell B. Stinchcombe. Extensive form games in continuous time: Pure strategies. *Econometrica*, 57(5):1171–1214, 1989. → pages 2, 41, 42, 43, 44, 46, 48, 49, 50, 56, 58, 69, 159, 160, 161, 163, 164, 165, 173

Jorg Stoye. Axioms for minimax regret choice correspondences. *Journal of Economic Theory*, 146:2226–2251, 2011a. → pages 63, 64, 175, 177

Jorg Stoye. Statistical decisions under ambiguity. *Theory and Decision*, 70(2): 129–148, 2011b. → pages 64, 176

Emanuel Vespa and Alistair J. Wilson. Experimenting with equilibrium selection in dynamic games. Mimeo, 2016. → pages 45, 61

Abraham Wald. *Statistical Decision Functions*. Wiley, 1950. → pages 64, 177

Appendix A

A Brief Introduction to the Theory of Games With Ambiguity Averse Agents

This appendix provides a brief background on the theory of games with ambiguity averse agents; it is particularly relevant for readers of both Chapters 2 and 4. Section A.1 surveys the existing solution concepts for ambiguity averse agents. Section A.1.1 re-introduces Lo-Nash equilibrium (originally previewed in Chapter 4) and uses the game studied in Chapter 2 as an illustrative example. Section A.1.2 presents a slightly different formulation of equilibrium under uncertainty than that presented in Chapter 4. A comparison of the differences in the equilibrium sets between Lo-Nash equilibrium and equilibrium under uncertainty highlights the role that the different consistency conditions play across the two equilibria; in particular, (C,Y) is an equilibrium of the testing game from Chapter 2 for equilibrium under uncertainty but not for Lo-Nash equilibrium.

A.1 Theoretical Literature

There is a growing literature that models the extension of game theory to subjects with non-neutral ambiguity preference. There is, however, surprisingly little agreement on how to proceed. Should an ambiguity averse equilibrium concept maintain mutual knowledge of rationality? Should stochastic independence be preserved, or should we allow for correlation between agent's strategies? How much (if at all) should consistency of beliefs be relaxed? Each of these questions could be used to categorize the literature in various ways. We focus, instead, on the role of mixed strategies.

To reiterate, there are three ways for the game theorist to deal with mixed strategies. Either allow agents to hold strict preferences for mixed strategies, restrict the strategy space to contain only pure strategies (and thereby avoid the issue altogether), or to place restrictions on the model such that the agent does not have a strict preference for mixed strategies. Models that fit the third category require careful attention to be paid regarding the timing of realizations of randomization devices. Currently, the only two papers that have used the third approach have not been sufficiently careful (Rothe (2010) and Lehrer (2011)). Table A.1 provides a summary of the state of the literature.

As can be seen in table A.1 a majority of the literature simply excludes mixed strategies. The experimental design in Chapter 2 could be reinterpreted using any of the equilibrium concepts that exclude mixed strategies, and this appendix includes the equilibrium derivations for Dow and Werlang (1994) and Lo (2009). It is not, however, compatible with models that allow a strict preference for randomization, such as Lo (1996) or Bade (2011).[1] For example, propositions 4 and 5 in Lo (1996) imply that for two-player games with a unique strict pure strategy Nash equilibrium, such as the game in figure 2.6 here, the only ambiguity averse equilibrium is the Nash equilibrium – a conclusion that is rejected in our data. Bade (2011) gives agents access to an ambiguous randomization device (i.e. agents can condition their strategy on a draw from an Ellsberg urn, for example), but finds observational equivalence between Nash equilibria and ambiguity averse equilibria in two-player games.

Returning to the questions raised in the opening paragraph of this section, in the opinion of this author, the most convincing answers have been provided by Lo

[1] The experimental results presented in Chapter 2 *could* be interpreted as providing some evidence against models such as Lo (1996) or Klibanoff (1996). An alternative viewpoint would be to argue that a proper test of these papers would require an explicit mixing device be provided to subjects. Similarly, a proper test of Bade (2011) would require subjects to have access to an ambiguous mixing device.

Paper	Equilibrium concept	Treatment of mixed strategies
Paper	**Equilibrium concept**	**Treatment of mixed strategies**
Dow and Werlang (1994)	Nash equilibrium under uncertainty	Pure strategies only
Lo (1996)	Belief equilibrium	Strict preference for mixed strategies allowed
Klibanoff (1996)	Equilibrium with uncertainty aversion	Strict preference for mixed strategies allowed
Epstein (1997)	Rationalizability	Pure strategies only
Groes et al. (1998)	Nash equilibrium with lower probabilities	Pure strategies only
Marinacci (2000)	Belief equilibrium	Pure strategies only
Eichberger and Kelsey (2000)	Equilibrium under uncertainty	Pure strategies only
Lo (2009)	Correlated Nash equilibrium	Pure strategies only
Rothe (2010)	Choquet Nash equilibrium	Mixed strategies have payoffs that are assumed to be linear in utility
Bade (2011)	Ambiguous act equilibrium	Strict preference for mixed strategies allowed. Allows for ambiguous strategies (using an ambiguous mixing device)
Azrieli and Teper (2011)	Ex-ante J equilibrium (games with incomplete information)	Strict preference for mixed strategies allowed.
Lehrer (2011)	Partially specified equilibrium	Mixed strategies have payoffs that are assumed to be linear in utility
Grant et al. (Forthcoming)	Savage games	Pure strategies only
Eichberger and Kelsey (2014)	Equilibrium under ambiguity	Pure strategies only
Riedel and Sass (2014)	Ellsberg equilibrium	Strict preference for mixed strategies allowed. Allows for ambiguous strategies (using an ambiguous mixing device)

Table A.1: This table contains summary of the treatment of mixed strategies in the literature of ambiguity averse equilibrium concepts. Only papers that define and introduce equilibrium concepts are included, applications are not included. Note that the approach taken by Lehrer (2011) and Rothe (2010), of assuming that the payoffs of mixed strategies are linear in utility, raises some difficulties that are not addressed in either paper.

(2009). Lo posits an epistemic approach to ambiguity averse game theory; beginning with the epistemic conditions of Nash equilibrium, what are the consequences of relaxing rationality to allow for ambiguity averse preferences? The resultant equilibrium concept substantially relaxes stochastic independence[2], allows agents to hold set valued beliefs regarding their opponent's strategies and requires actual play to lie within the set of beliefs, and maintains mutual knowledge of rationality.

The only desiderata missing from Lo (2009) is that it does not allow for mixed strategies; the strategy space is restricted to pure strategies, and a population interpretation of mixing is invoked. Chapter 4 extends Lo (2009) to allow for mixed strategies.

A.1.1 Lo-Nash equilibria

This section introduces Lo-Nash equilibrium, following Lo (2009) closely. Define a set of players $N = \{1,\dots,n\}$, let each player $i \in N$ have a finite set of actions A_i, and define $A = \times_{i\in N} A_i$ and $A_{-i} = \times_{j\neq i\in N} A_j$. We shall endow each agent with a Bernoulli utility function $u : A \mapsto \mathbb{R}$. Suppose that an agent has uncertainty regarding the strategy choices of their opponents, A_{-i}. Then we can regard a strategy, a_i, as an act over the state space A_{-i} generating a payoff $u_i(a_i, a_{-i})$ when the state a_{-i} is realized.

In a manner consistent with Gilboa and Schmeidler (1989)'s MEU formulation, we suppose that an agent's beliefs regarding their opponents strategies are a closed and convex set of probability measures $\Phi_i \subseteq \Delta(A_{-i})$. Given Φ_i an agents preferences are represented by

$$\min_{\phi\in\Phi} \sum_{a_{-i}\in A_{-i}} u_i(a_i, a_{-i})\phi_i(a_{-i}).$$

Furthermore, we use σ to denote a probability measure on A. We define $\sigma^{A_i}(a_i) = \sum_{a_{-i}\in A_{-i}} \sigma(a_i, a_{-i})$ as the marginal distribution of σ on A_i and $\sigma^{A_{-i}}(a_{-i}) = \sum_{a_i\in A_i} \sigma(a_i, a_{-i})$ as the marginal distribution of σ on A_{-i}. Then, in the usual fashion we write $\sigma(a_{-i}|a_i) = \frac{\sigma(a_i,a_{-i})}{\sigma^{A_i}(a_i)}$.

[2]In fact, Lo refers to his equilibrium as Correlated Nash equilibrium to emphasize the similarities with Aumann's Correlated equilibrium. We avoid using this name, and instead refer to the equilibrium as Lo-Nash equilibrium, to avoid confusion between the two concepts.

Finally, we write supp σ to denote the support of the probability distribution σ, and define supp Φ to be the union of the supports of the elements of Φ. We are now ready to define a Lo-Nash equilibrium.

Definition A.1 (Lo-Nash equilibrium)**.** *A pair $<\sigma,\Phi>$ forms a Lo-Nash equilibrium if it satisfies*

$$\sigma(\cdot|a_i) \in \Phi_i \;\; \forall a_i \in supp\, \sigma^{A_i}, \forall i \in N \tag{A.1}$$

$$supp\, \Phi_i = \times_{j \neq i} supp\, \sigma^{A_j} \;\; \forall i \in N \tag{A.2}$$

and

$$a_i \in \arg\max_{\hat{a}_i \in A_i} \min_{\phi_i \in \Phi_i} \sum_{a_{-i} \in A_{-i}} u_i(\hat{a}_i, a_{-i})\phi_i(a_{-i}) \;\; \forall a_i \in supp\, \sigma^{A_i}, \forall i \in N \tag{A.3}$$

Equation A.3 requires that all strategies that are played in an equilibrium are best responses, with preferences defined as MEU preferences with respect to the equilibrium conjectures Φ. Equations A.1 and A.2 are the consistency requirements: equation A.2 ensures that a strategy is played with a positive probability iff it is expected to be played with a positive probability, and equation A.1 forces actual strategies to be contained in the belief sets. Note that equation A.1 allows for conditioning of σ on a_i - this allows for strategies to be correlated, but the realized strategy must lie within player i's belief set for all a_i.

Lo-Nash equilibrium: an example

Consider the normal form game presented in figure A.1. For this section, we shall follow the standard game theoretic approach and assume that the payoffs in figure A.1 are utility values, thereby abstracting from issues of risk aversion. Recall that the game has a unique Nash equilibrium (A,X).

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

Figure A.1: Testing game

In contrast, there are fully mixed Lo-Nash equilibria in this game. Intuitively, this is easy to see. If the row player is uncertain about the column player's strategy choice, and the row player is ambiguity averse, then there are set-valued beliefs for which C is a best response. There does not, however, exist a Lo-Nash equilibrium where the row player plays C with probability 1. The consistency requirements for Lo-Nash equilibria would then imply that the column player would know with certainty that the row player will play C, so the unique best response for the column player is Y. In this case, C is no longer the best response for the row player and the equilibrium breaks down. It is the case, however, that C can be played with a probability that is arbitrarily close to one. The full set of Lo-Nash equilibria for this game is large.

Lemma A.1. *The set of distributions, Σ, that can be supported as Lo-Nash equilibria for the game in figure A.1 is* $\Sigma = \{\sigma(A,X) = 1\} \bigcup \{\sigma : 0 < \sigma(A),\ 0 < \sigma(B),\ 0 < \sigma(C),\ \frac{4}{11} \leq \sigma(X|\cdot) \leq \frac{7}{11}\ \} \bigcup \{\sigma : \sigma(A) = 0,\ 0 < \sigma(C) < 1,\ \sigma(X|\cdot) \leq \frac{7}{11},\ \sigma(Y) < 1\ \} \bigcup \{\sigma : \sigma(B) = 0,\ 0 < \sigma(C) < 1,\ \frac{4}{11} \leq \sigma(X|\cdot),\ \sigma(Y) > 0\ \}$.

Proof of lemma A.1. We proceed case-by-case, considering each possible set of pure strategy supports for the row player in turn. It follows directly from the definition of Lo-Nash equilibrium that we can find the full set of equilibrium distributions for each case by first identifying the largest feasible equilibrium sets of beliefs.

It is straightforward to see that the only pure strategy Lo-Nash equilibrium is $\{A,X\}$.

Next we consider equilibria that are fully mixed for the row player. The row player will be indifferent between all three of their strategies only when $\Phi_R = \{\phi : \frac{4}{11} \leq \phi(X) \leq \frac{7}{11}\}$. This implies that the column player can play any mixed strategy within this range. The largest set of beliefs such that the column player is indifferent between both X and Y is beliefs of complete uncertainty. The consistency requirements on the row player are therefore very weak, such that they need to use each strategy with a strictly positive probability. Note that the equilibrium allows for correlation between strategies.

Next, we consider equilibria where the row player does not play A. In such cases, the largest feasible set of row player beliefs are given by $\Phi_R = \{\phi : \phi(X) \leq$

$\frac{7}{11}\}$. Consistency requirements restrict the column player's beliefs to be $\Phi_C = \{\phi : \phi(A) = 0, \phi(B) > 0, \phi(C) > 0\}$, which is also the largest feasible set of beliefs that ensures column player indifference. Therefore the row player can play any mix that places positive weight on both B and C. The correlation requirements allow for equilibria that are arbitrarily close to pure $\{C,Y\}$. Consider the equilibrium with $\sigma(B,X) = 7\varepsilon$, $\sigma(B,Y) = 4\varepsilon$ and $\sigma(C,Y) = 1-11\varepsilon$ with all other outcomes being assigned a probability of 0, and take ε to be arbitrarily small. The equilibria where the row player does not play B are built in a similar fashion.

There are no equilibria where the row player mixes between only A and B. In such a case, the column player must believe that $\phi(C) = 0$ and therefore they have a unique best response of X. □

This example illustrates a key feature of Lo-Nash equilibrium. The source of ambiguity is not from doubts over which Nash equilibrium is to be played, nor is it over which strategy will be realized in a mixed strategy Nash equilibrium. The source of ambiguity in this game is endogenous to the game structure and agent's preferences. The only requirement is that the row player has ambiguity averse preferences, so that C is a best response for some potential row player beliefs and thereby breaking the rationalizability chain required to produce the unique (A,X) equilibrium.

As is always the case, the unique Nash equilibrium is also a Lo-Nash equilibrium. Ambiguity averse preferences, on their own, need not necessarily affect the play of a game if the agents still form point-valued conjectures. If an agent believes, with certainty, that their opponent will play a particular (possibly mixed) strategy then there is no role for their ambiguity to play as, subjectively, the agent faces no ambiguity.

The other key feature of the game in figure A.1 is that equilibria other than (A,X) exist iff the row player is ambiguity averse (lemmas A.2 and A.3). If the row player has SEU preferences, then the column player's ambiguity preference has no effect on the equilibrium set. Using this result, we can differentiate the effect of a players own ambiguity preference (which should affect row player behaviour only) from the effect of a players beliefs about their opponent's ambiguity preference (which should affect column player behaviour only).

Definition A.2. *A Lo-Nash equilibrium is* proper for agent i *iff* Φ_i *is not a singleton. A Lo-Nash equilibrium is* proper *iff it is proper for at least one agent.*

Lemma A.2. *All proper Lo-Nash equilibrium of the game in figure A.1 are proper for the row player.*

Proof of lemma A.2. From lemma A.1, all proper Lo-Nash equilibria involve the row player mixing between C and at least one of A or B. If the row player has singleton beliefs, then their set of best responses can never include C. Therefore, any proper Lo-Nash equilibrium must be proper for the row player. □

Lemma A.3. *There exist proper Lo-Nash equilibrium of the game in figure A.1 that are not proper for the column player.*

Proof. Consider the equilibrium where $\sigma(A,X) = \sigma(A,Y) = \frac{1}{6}$, $\sigma(B,X) = \sigma(B,Y) = \frac{1}{9}$, $\sigma(C,X) = \sigma(C,Y) = \frac{2}{9}$, $\Phi_R = \{\phi : \frac{4}{11} \leq \phi(X) \leq \frac{7}{11}\}$ and $\Phi_C = \{\phi : \phi(C) = \frac{4}{9}, \phi(A) = \frac{1}{3}, \phi(B) = \frac{2}{9}\}$. □

A.1.2 Nash Equilibrium under uncertainty

In this section we describe Dow and Werlang's (1994) Nash Equilibrium under uncercainty. The analysis will follow Dow and Werlang (1994) very closely. We begin by introducing the notion of a sub-additive probability, P, being a probability which satisfies

$$P(A) + P(B) \leq P(A \cap B) + P(A \cup B).$$

The expected utility of an agent with respect to a sub-additive probability over a non-negative random variable, X, is given by the Choquet integral

$$E(X) = \int_{\mathbb{R}^+} P(X \geq x)dx.$$

In many cases this will be equivalent to calculating the Maxmin Expected Utility with respect to the core of the sub-additive probability (see Gilboa and Schmeidler (1994) for details).

Dow and Werlang define the support of a sub-additive probability to be an event, A, such that $P(A^c) = 0$ and $P(B^c) > 0$ for all $B \subset A$. We note that the support need not be unique, and that there are other reasonable definitions of the support of a sub-additive probability that are not used here.

A Nash equilibrium under uncertainty is then simply the requirement that all strategies that are in the support of an opponent's beliefs are best responses for an agent given the agent's beliefs. Label the set of available strategies for player i as A_i.

Definition A.3 (From Dow and Werlang (1994))**.** *A pair (P_1, P_2) of non-additive probabilities P_1 over A_1 and P_2 over A_2 is a Nash Equilibrium under Uncertainty if there exist a support of P_1 and a support of P_2 such that*

(i) for all a_1 in the support of P_1, a_1 maximizes the expected utility of player 1, given that P_2 represents player 1's beliefs about the strategies of player 2; and conversely,

(ii) for all a_2 in the support of P_2, a_2 maximizes the expected utility of player 2, given that P_1 represents player 2's beliefs about the strategies of player 1.

We refrain from giving the full set of equilibrium sub-additive probabilities, as many of them produce observationally equivalent outcomes, but instead establish two facts: that a fully mixed equilibrium can be supported, and that pure $\{C, Y\}$ can be supported, in a Nash Equilibrium under Uncertainty.

Lemma A.4. *The pure strategy equilibrium $\{C, Y\}$ can be supported by the equilibrium sub-additive probability P_C with $P_C(Y) = \frac{3}{11}$ and $P_C(X) = 0$ and additive probability P_R with $P_R(C) = 1$.*

Proof of lemma A.4. The support of P_C is $\{Y\}$ and the support of P_R is $\{C\}$. It is obvious that Y is the unique best response for the column player. The row player evaluates their strategies as follows:

$$U(A) = 1 * 14 + 0 * (25 - 14) = 14$$

$$U(B) = 1 * 14 + \frac{3}{11} * (25 - 14) = 17$$

$$U(C) = 18$$

C is the unique best response. □

Lemma A.5. *A fully mixed equilibrium can be supported by the equilibrium sub-additive probability P_C with $P_C(X) = P_C(Y) = \frac{4}{11}$ and additive probability P_R with $P_R(C) = \frac{4}{9}$, $P_R(A) = \frac{3}{9}$, $P_R(B) = \frac{2}{9}$.*

Proof of lemma A.5. The support of P_R is $\{A,B,C\}$ and the support of P_C is $\{X,Y\}$. Given that P_R is additive we can calculate the column player's utility in the standard fashion, and verify that $U(X) = U(Y) = \frac{148}{9} \sim 16.4$. The row player evaluates their strategies as follows:

$$U(A) = 1 * 14 + \frac{4}{11} * (25 - 14) = 18$$
$$U(B) = 1 * 14 + \frac{4}{11} * (25 - 14) = 18$$
$$U(C) = 18$$

Therefore, both players are indifferent between all strategies in the supports of the distributions. □

Appendix B

Appendix for "Uncertainty Aversion in Game Theory: Experimental Evidence"

B.1 Experimental Methodology Appendix

In every experimental design there are a myriad of decisions that the experimenter must make. This section discusses some of the more important choices that were made in the design of this experiment. The focus is on issues that are likely to be less familiar to readers, rather than standard experimental practices such as ensuring payments are anonymous.

B.1.1 Piloting and framing effects

This experiment has two distinct sections. First, we measure the subjects preferences and beliefs over their opponent's preferences. Second, we give the subjects a variety of normal form games. Normal form games are normally presented to subjects in a bi-matrix format. In contrast, decision theorists use a range of procedures to measure preferences, none of which are bi-matrices. A natural concern is that the framing of the task might affect the measurement (and it turns out that this concern is well founded). In a series of pilot experiments, we observed that the proportion of subjects that display ambiguity aversion is significantly higher when the task is presented as a worded decision problem (when compared to a bi-matrix game).[1]

The existence of a framing effect precludes presenting the preference measurement tasks and normal form games in differing formats. To resolve this problem, we presented all tasks as games. As we are interested in measuring both a subject's own preferences and their beliefs regarding their opponent's preferences, we formulate the decision problems as a game where the row player reveals their own preferences and the column player attempts to predict the row player's behaviour.

A natural concern is that, by moving the preference measurement tasks into a game format, our preference measures will be distorted relative to previous research on ambiguity aversion. The pilot experiments suggested that adding a series of comprehension questions, as described in the next section, does a reasonable job of reducing the gap between worded decision problems and bi-matrix games.[2]

B.1.2 Assumptions underlying the classification procedure

There are several assumptions underlying the classification procedure used in section Section 2.2.1, which are discussed in this section. Briefly, the assumptions are: source independence of ambiguity preferences; CRRA utility; and, for column players, a strength of beliefs assumption. The results are robust to varying the CRRA utility assumption. The strength of beliefs assumption may be violated in the data, and the additional treatments in Section 2.4 investigate this possibility.

Assumptions on ambiguity preferences

Underlying the classification mapping is an assumption that a subject will (subjectively) view an Ellsberg urn as a source of ambiguity if and only if they view

[1] Similar evidence of framing effects for ambiguity aversion measurement can be found in Chew et al. (2013).

[2] Despite the inclusion of the comprehension questions, the measured level of ambiguity aversion is still slightly lower than that measured in previous 2-colour Ellsberg urn experiments. The results reported here do appear to be consistent with recent evidence from Chew et al. (2013), who argue that we should expect to measure less ambiguity aversion in more complex environments.

their opponent's strategy as a source of ambiguity.[3] Using an Ellsberg urn to produce the classification provides a stable point of comparison to the vast literature on ambiguity aversion in individual decision problems.

Furthermore, it is possible that some subjects who chose S in in the game in figure 2.2 would switch to M if the payoff for S was lowered to, say, \$30.01. Such a subject has ambiguity averse preferences but would be classified as ambiguity neutral given the payoff structure of my game. However, the results in previous work, such as Epstein and Halevy (2014), suggest that this is not a significant source of misclassification even when the payoff for S is as high as \$31.00. Therefore, for these two reasons, we conclude that our classification provides a lower bound on the proportion of subjects who may exhibit ambiguity aversion with respect to their opponent's preferences.

Assumptions on risk preferences

The mapping from choices in the risk measurement game to types will depend on the parameterization of the utility function. We consider two common utility functions: constant relative risk aversion (CRRA, equation B.1) and constant absolute risk aversion (CARA, equation B.2).[4] Figures B.1 and B.2 present the implications

[3] Heath and Tversky (1991) were the first to provide evidence of differing source preferences for ambiguity, comparing bets on political outcomes with bets on football matches. They found evidence that subjects prefer betting on events for which they have a greater familiarity. While it is unlikely that the subjects in this experiment were substantially more familiar with either the playing of normal form games or betting on Ellsberg urns, it is still possible that source dependent preferences could be present in the current environment. In particular, some subjects may find the Ellsberg urn to be less ambiguous than their opponent's behaviour because an Ellsberg urn is symmetric and determined solely by 'nature', while an opponent has their own agency – an argument first put forward by Kelsey and le Roux (2015a). To the extent that we mis-classify subjects, we should expect there to be some subjects who are measured as ambiguity neutral with respect to the urn yet are ambiguity averse with respect to their opponent's action.

[4] The CRRA utility function is given by the equation

$$u(c) = \frac{c^{1-\beta} - 1}{1-\beta} \tag{B.1}$$

and the CARA utility function by

$$u(c) = 1 - e^{-\rho c}. \tag{B.2}$$

Note the implicit assumption that subjects treat the experiment as an independent event, and do not integrate their outside wealth into their utility functions. This is the standard approach in experimental economics, although recently some authors have begun to consider partial wealth integration of the CARA and CRRA parameterizations, respectively.

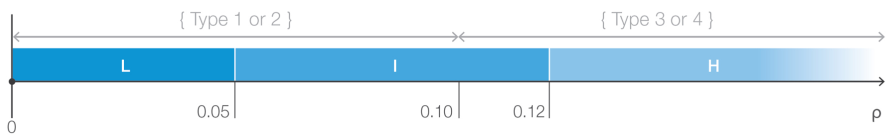

Figure B.1: The role of risk aversion in the risk preference measuring game and the testing game assuming SEU preferences with CARA utility: $u(c) = 1 - e^{-\rho c}$. Subjects who select L in the game in figure 2.5 have type 1 or 2 preferences, while subjects who select H have type 3 or 4 preferences. Subjects who select I are of an indeterminate type.

As can be seen from figure B.1, the CARA parameterization provides a clear classification for subjects that choose either L or H, but is indeterminate for subjects who choose I. Meanwhile, figure B.2 indicates that, under the CRRA parameterization, subjects that choose L and I should be pooled into a low risk category while subjects who choose H should be considered high risk (although there is a small region of misclassification for subjects that choose H). Given that both models essentially agree on the treatment of subjects that choose L and H, we shall implement the CARA parameterization and remove subjects that selected I from the data analysis. Robustness checks indicate that moving to a CRRA parameterization has no effect on the statistical inferences implied by the results, which is unsurprising given that only 10% of subjects choose I in the risk measurement game.

Additional assumptions on beliefs

The mapping for the column players is calculated analogously to the mapping for row players, although it requires an additional assumption regarding the strength of beliefs. The process for eliciting beliefs regarding opponent's preferences elicits the subject's 'best guess' regarding his opponent. In the worst case scenario a subject may be completely unsure about their opponent's preferences, implying

(Holt and Laury (2014)). A theoretically more robust formulation would use $u(c+w)$ where w is the subjects initial wealth level, but we follow the experimental standard here. As suggested by Simon Grant, readers who are troubled by this formulation should focus on the CARA utility parameterization (and not the CRRA parameterization).

Figure B.2: The role of risk aversion in the risk preference measuring game and the testing game assuming SEU preferences with CRRA utility: $u(c) = \frac{c^{1-\beta}-1}{1-\beta}$. Subjects who select L or I in the game in figure 2.5 have type 1 or 2 preferences, while subjects who select H are likely to have type 3 or 4 preferences.

that their choices as the column player in the classification games may contain no information. The hypothesis of a relationship between column player behaviour in the classification games and the testing game can, therefore, be seen as a joint test of the theory outlined in section 2.1 and the assumption that beliefs are strong enough to be informative. Section 2.4 introduces an additional treatment in which subjects are shown a signal of their opponent's preferences, allowing us to separate this joint hypothesis into its component parts.

B.1.3 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 2.2, 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.[5] The process of interpreting the bi-matrix game and converting it into a worded decision problem helped to ameliorate the framing effects induced by the bi-matrix game.

In addition, the drop down menus also served as a comprehension check. This test of subject comprehension was quite stringent. There were a handful of subjects who had extreme difficulty with the drop down menus, and only 25% of subjects earned the maximum possible bonus payment of $7. We used the level of bonus payment earned by the subject as a measure of comprehension, and removed subjects who performed poorly from our sample.

Given the extensive piloting process and experimenter degrees of freedom involved in using the comprehension data to exclude subjects, the possibility of data mining for statistically significant results needs to be addressed. Two approaches were taken to alleviate these concerns. First, pilot sessions were run using only the classification games of section 2.2.1, and not the testing game of section 2.2.2.[6] Second, the rules for excluding subjects from the data, based on comprehension scores, were determined ex-ante rather than ex-post.

Two rules for excluding subjects were determined ex-ante. A strict rule, and a relaxed rule. There was an expectation that the strict rule may be too strict, so the relaxed rule was designed as a back up. The expectations were correct, and after applying the strict rule only 8 ambiguity averse and risk neutral subjects remained (and only 4 risk averse and ambiguity neutral subjects); therefore, the relaxed rule was used. The comprehension data is reported in further detail in appendix B.3.3.

[5] In one set of pilot sessions subjects were required to fill in the drop down menu only for the strategy that they selected. This approach was soon abandoned, as it turned out that subjects were making their decisions without referring to the drop down menus, and then only filling in the drop down menus after they had already come to a decision. The drop down menus thereby failed to break down the framing effect of presenting the choices in a bi-matrix format. In the final experimental design subjects were required to fill in the drop down menus before they could select a strategy.

[6] The only exceptions to this were the first two pilot sessions which used a 3-colour Ellsberg urn instead of the 2-colour Ellsberg urn used in the final design. Only 1 out of 14 subjects was found to be ambiguity averse, meaning that it was not possible to perform any tests on the effects of ambiguity aversion on behaviour, and the 3-colour Ellsberg urn was removed from the design.

B.1.4 Order of realization of randomizations

As discussed in section 2.1, many models of ambiguity aversion induce preferences with a strict preference for randomization. In an experiment, such as this one, where we are asking subjects to respond to multiple games there are two ways in which a preference for randomization could undermine the incentive structure. In both cases, however, the results of Eichberger et al. (2016) (and similar ideas found in earlier papers, dating back to at least Kreps (1988)) can be used to overcome any concerns: if objective randomizations occur before subjective states are realized then dynamic consistency implies no preference for randomization.

The first difficulty, as discussed in Chapter 4, occurs within a game. If subjects have a strict preference for randomization then the availability of mixed strategies can change the equilibrium set (relative to the game where only pure strategies are available). The strict preference for randomization can be negated via two arguments. First, subjects were required to select pure strategies in each of the games (i.e. the computer interface did not allow for mixed strategies to be chosen explicitly). In order to play a mixed strategy subjects would need to provide their own randomization device, whether it be a mental mixing, tossing of a coin or rolling of a die.[7] Second, following the arguments of Eichberger et al. (2016) there should be no preference for randomization if a subjects own randomization is resolved before they view their opponents strategy choice. This requirement was satisfied during the experiment as a consequence of requiring subjects to enter pure strategy choices into the computer, and not revealing their opponent's strategy choices until the end of the experiment.

The second difficulty occurs between games. If subjects have a strict preference for mixing between games, then using a random payment mechanism[8] will provide the subject with an opportunity to hedge their ambiguity across games and they will not treat each game as an independent decision problem. On the other hand, paying subjects for their decisions in all games will obviously give rise to hedging opportunities across games, particularly given that several of the games

[7] No subjects were observed to flip coins or roll dice during the experiment.

[8] A random payment mechanism chooses, at random, a subset of the tasks in the experiment to be the tasks for which subjects are actually paid. Usually, the number of tasks that are selected is one, and this is sometimes referred to as a "pay-one" incentive scheme.

include a draw from an Ellsberg urn that determines the ‘play’ of nature. Baillon et al. (2014) provides some solace, and a joint reading of Azrieli et al. (2014) and Eichberger et al. (2016) reinforces this view. Baillon et al. (2014) demonstrate that if objective randomizations are resolved after subjective randomizations then the elicitation is not incentive compatible; if objective randomizations are realized first then incentive compatibility may be restored. Azrieli et al. (2014) suggest that, absent an assumption that subjects are indifferent between ex-ante and ex-post randomizations, there are no strong objections to be made against using a random payment mechanism in an experiment of this nature. Following this, it is clear from Eichberger et al. (2016) that the random payment mechanism should be resolved prior to the resolution of play in the games. Furthermore, Johnson et al. (2014) demonstrate experimentally that an elicitation process which they call PRINCE, including this order of realizations, can improve the quality of elicited preferences.

Resolving the random payment mechanism prior to the resolution of play in the games requires some care. It is obviously not appropriate to allow the subjects to know which game will be chosen for payment before the subjects play the games. The solution implemented here was as follows: when subjects entered the experimental lab there were 7 flash cards stuck to one of the walls. On the front (visible to subjects) side of the flash cards were the letters A through G, one letter per flash card. On the back (not visible to subjects) were the numbers 1 through 7, one number per flash card. The matching of letters to numbers was randomly determined by the experimenter.[9] During the instructions, a subject was asked to choose a letter from A to G, and the number was recorded. Subjects were informed that, at the end of the experiment, the number on the back of the chosen letter would determine which game would be paid. From the subjects perspective the choice of game was then fixed, but unknown. At the end of the experiment, the letter was flipped over to reveal the game to be paid before any balls were drawn from the urns.

All randomizations and ball draws were conducted using physical devices that were as procedurally transparent as possible.

[9]There were actually two sets of flash cards, one with letters and one with numbers. Letter flashcards were then randomly blu-tacked to number flashcards. This enabled the flash cards to be easily re-randomized between sessions.

B.1.5 Related experimental literature

There is very little experimental evidence regarding ambiguity aversion in game theory. There are, as far as I am aware, only a handful of papers that have collected experimental data that was designed to investigate ambiguity aversion in game theory (Kelsey and le Roux (2015a), Kelsey and le Roux (2015b), Ivanov (2011), Eichberger et al. (2008), Pulford and Colman (2007), and Camerer and Karjalainen (1994)).[10] Additionally, Eichberger and Kelsey (2011) re-analyses previous experimental data and Heinemannn et al. (2009) study strategic uncertainty through the lens of global games. We proceed chronologically.

Camerer and Karjalainen (1994) present evidence from 4 games. The first two games are designed to replicate the standard Ellsberg tasks, but replacing the subjective ball draw with a choice by an opponent. The third game requires subjects to predict the outcome in a co-ordination game that had previously been played by other subjects. In each of these games Camerer and Karjalainen found, on average, a consistent yet small amount of ambiguity aversion along with a large amount of heterogeneity between subjects. These games, however, do not really have strategic interaction in a meaningful sense; in each game the strategic uncertainty only affects one of the players.

The fourth game that Camerer and Karjalainen study is a matching pennies game with an extra risky option for the row player. This is a very clever design that allows the row player to, effectively, choose between a bet on their opponent's strategy choice or a bet on a random draw. Again, the average subject was measured to be ambiguity averse.

Pulford and Colman (2007) took an approach that is somewhat opposed to that of Camerer and Karjalainen (1994). In Pulford and Colman (2007) subjects faced either a complete information game where their opponent was completely indifferent between all outcomes, or a game of incomplete information where their opponent had one of two types and each type of opponent had a differing dominant strategy. Pulford and Colman argue that the complete information game is a risky game and the incomplete information game is an ambiguous game (when subjects

[10]Di Mauro and Castro (2011) study a voluntary contribution game where subjects play against virtual agents. It is debatable whether this should be viewed as a game, or simply a specially structured decision problem.

were not informed about the probability distribution of types).

Pulford and Colman argue, using the principle of insufficient reason, that the only reasonable beliefs are that an opponent who is indifferent everywhere should be expected to choose each strategy precisely half the time. This assumption is very strong. A weaker, more reasonable, assumption would be that we require beliefs in this case to be symmetric (but possibly non-additive). Under this alternative assumption it becomes very difficult to interpret the results of Pulford and Colman (2007).

Eichberger et al. (2008) is a very entertaining paper. Subjects played a variety of normal form games against either a granny, a game theorist, or another subject. Eichberger et al. hypothesized, and it was confirmed in the data, that subjects would view playing against the game theorist as being less ambiguous than playing against the granny.[11] There was no evidence of any differences in behaviour between subjects playing against the granny or against another subject.

Eichberger et al. structured their games in a manner which allowed the ambiguity involved in a game to be quantified. Fixing the identity of the opponent, it was usually the case that subjects played the Nash equilibrium strategy less in games with higher levels of ambiguity.[12]

The results in Eichberger et al. (2008) are clearly *consistent* with ambiguity averse behaviour, although without observing individual level ambiguity measures it is difficult to rule out other potential explanations of the data. Eichberger and Kelsey (2011) revisit the data from Goeree and Holt (2001), and establish that the data is also broadly consistent with ambiguity averse behaviour. The analysis in both papers is, however, at the aggregate level, and given the large amount of between subject heterogeneity that has been observed in decision theory ambiguity aversion experiments (see Halevy (2007) for example) it seems that tracking individual subject level behaviour across games would be a natural next step.

Ivanov (2011) takes this next step. He uses a considerably more complicated design, and makes extensive use of stated beliefs in his identification process.[13]

[11] Personally, I find this result to be surprising. My prior was that undergraduates would feel more comfortable predicting the behaviour of a granny than a game theorist, given that most undergraduates would presumably have spent more time interacting with grannies than game theorists.

[12] The exceptions occurred for subjects that were playing the game theorist.

[13] Eichberger et al. (2008) also asked for stated beliefs, but they did not put them to work in the

First, subjects played a series of normal form games. Next, subjects were asked to state their beliefs regarding their opponent's behaviour in the normal form games.[14] Finally, subjects were asked to choose between lotteries that were constructed from their beliefs in part two and the payoffs from the games in part one.

The central question posed in Ivanov (2011) is quite different from the questions raised in this paper. Ivanov seeks to classify subjects as ambiguity averse/ambiguity neutral/ambiguity loving using only their responses to normal form games; the headline result is that 22/46/32 percent of subjects fall into each category, respectively. Unlike Ivanov (2011), this paper does not consider ambiguity seeking behaviour (ambiguity seeking subjects will be indistinguishable form ambiguity neutral subjects given the experimental design in the current paper). On the other hand, Ivanov (2011) requires an assumption that stated beliefs are equal to true beliefs (troublingly, if we model ambiguity aversion using the MEU model, then the notion of 'true beliefs' may not even be be well defined) in order to identify ambiguity preference, whereas the current paper measures ambiguity preference directly.

Kelsey and le Roux (2015a) is the only other papers that compares behaviour in Ellsberg urn tasks to behaviour in normal form games at the individual level. The results reported in Kelsey and le Roux (2015a) are mostly negative, however. They find significant levels of ambiguity aversion in their normal form game (a battle of the sexes game augmented with a safe option for the column player), but find much less ambiguity aversion (and even some ambiguity seeking) in their Ellsberg urn tasks. This contrasts with the results here, where Ellsberg urn behaviour was correlated with behaviour in the game in figure 2.6.

There are three potential reasons for the differences in results between the two papers. Firstly, Kelsey and le Roux (2015a) do not control for risk aversion. The behaviour in their normal form game might be driven by risk aversion, rather than ambiguity aversion. Secondly, Kelsey and le Roux (2015a) used a three-colour Ellsberg urn, rather than the two-colour Ellsberg urn used here. Chew et al. (2013) report that measured levels of ambiguity aversion are often much lower in two-

fashion of Ivanov.

[14] In one treatment beliefs were elicited at the same time as the games were played.

colour rather than 3-colour Ellsberg tasks.[15] The reasons for this difference are not exactly clear, although Chew et al. argue that it is a confusion effect (that subjects do not recognise the 3-colour urns as being ambiguous).

The third difference between Kelsey and le Roux (2015a) and the current paper is the use of comprehension questions. The results of Chew et al. (2013) suggest that, in some situations, subjects will have trouble recognizing and responding to ambiguity. To ensure that this does not affect results, we used a series of drop down menu comprehension questions to screen for understanding. Kelsey and le Roux (2015a) did not use any tests for comprehension, and this may make it harder to identify a relationship between ambiguity averse behaviour in the games and ambiguity averse behaviour in the Ellsberg tasks.

The identification of the link between individual choice tasks and normal form game behaviour is always going to be difficult because the Nash equilibrium remains an equilibrium for ambiguity averse subjects. This implies that even in ideal circumstances the power of a test to reject the null hypothesis of independence may be low (particularly if the Nash equilibrium is focal). Although, as demonstrated in this paper, the effect can be recovered with a clean experimental design. On balance, the evidence strongly suggests that ambiguity aversion has an influence on behaviour in normal form games.

Kelsey and le Roux (2015b) study games of public good provision. In the 'best shot' formulation of the game ambiguity averse subjects are expected to over-contribute to the public good and in the 'weakest link' formulation ambiguity averse subjects are expected to under-contribute. On aggregate the data supports the hypothesis that ambiguity aversion is prevalent in the subject population, although there is significant heterogeneity and no direct measures of ambiguity preferences are reported. Kelsey and le Roux also hypothesized that British subjects would behave in a more ambiguity averse fashion when paired with Indian subjects (when compared to British subjects paired with other British subjects) although they do not find evidence to support this in the data.

[15] In piloting for this experiment, two sessions were run using a 3-colour Ellsberg urn. Only 1 out of 14 subjects was classified as being ambiguity averse. Given this extremely low observed rate of ambiguity aversion, 3-colour Ellsberg urns were not used in any other sessions.

B.2 Proofs

The proofs in this section use 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 2.1. 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 $\Phi(X)_R > \frac{1}{2}$, is B when $\Phi(X)_R < \frac{1}{2}$ and is either A or B when $\Phi(X)_R = \frac{1}{2}$. 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, 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, as are X and Y. Therefore all strategies are rationalizable. $\square$

Proof of proposition 2.2. We use iterated elimination of never best response strategies.

In the first round, the type 1 row player eliminates C. All other strategies for all other players are best responses.

In the second round, the payoff to the column from playing Y is increasing in the probability that the row player plays C, and the payoff from playing X is decreasing the in the probability that the row player plays C. Therefore, because the largest probability that the column player can assign to the row player playing C is α, Y can only be sustained as a best response whenever $U(Y) - U(X) = u(22)(1-\alpha)+u(12)\alpha-(u(12)(1-\alpha)+u(20)\alpha) > 0$. Therefore, X dominates Y whenever $\alpha > \frac{u(22)-u(12)}{u(22)+u(20)-2u(12)} = \bar{\alpha}$, so that Y is eliminated as never being a best response iff $\alpha > \bar{\alpha}$.

A consequence of this is that no more strategies can be eliminated if $\alpha \leq \bar{\alpha}$. In this case, both column player strategies enter the rationalizable set, as do all three row player strategies, as all three strategies may be played by the non-type 1 row player. The Type 1 row player may play either A or B.

If $\alpha > \bar{\alpha}$, then we must consider the third round. In this case, Y has been eliminated in the second round, and both types of row players now have a unique best response of A. Therefore, $\{A, X\}$ is the only rationalizable strategy when $\alpha > \bar{\alpha}$.

If we normalize the utility function so that $u(22) = 1$ and $u(12) = 0$ then we have that $\bar{\alpha} = \frac{1}{1+u(20)}$. It is clear that $\bar{\alpha}$ is bounded below by $\frac{1}{2}$ and is decreasing in the degree of risk aversion. $\square$

B.3 Results Appendix

B.3.1 Demographic analysis

Table B.1 presents an analysis of the demographic data for row players. Specifications are equivalent to those presented in table 2.6 in the main text, with and without demographic controls: an indicator for female subjects, and an indicator for STEM (plus economics) majors. The sample size is reduced from that in the main text for two reasons: first, there were two sessions where software errors prevented the recording of demographic information and, second, for some subjects the self-reported major of study was not identifiable. As the table makes clear, neither gender nor area of study affects row player decisions in the testing game and the inclusion of demographic controls does not change the estimates of the effects of preferences on behaviour.

Table B.2 presents an analysis of the demographic data for column players. As for the row player data, there is no effect of either gender or area of study on behaviour in the testing game.

B.3.2 Preferences and beliefs

Table B.3 presents the relationship between a subject's own ambiguity aversion and their beliefs regarding their opponent's ambiguity aversion. Again, the sample is

Effect of	Conditional on	$\Delta Pr(A)$	$\Delta Pr(A)$	$\Delta Pr(A)$	$\Delta Pr(A)$
Ambiguity		-0.14	-0.15		
Aversion		(0.181)	(0.165)		
		[0.10]	[0.11]		
Risk		-0.13	-0.14		
Aversion		(0.224)	(0.211)		
		[0.11]	[0.11]		
Ambiguity Aversion	Low Risk Aversion			-0.13	-0.13
				(0.332)	(0.332)
				[0.14]	[0.14]
Ambiguity Aversion	High Risk Aversion			-0.16	-0.17
				(0.358)	(0.313)
				[0.17]	[0.17]
Risk Aversion	Ambiguity Neutrality			-0.12	-0.12
				(0.399)	(0.413)
				[0.14]	[0.14]
Risk Aversion	Ambiguity Neutrality			-0.14	-0.16
				(0.379)	(0.336)
				[0.16]	[0.17]
Risk Aversion & Ambiguity Aversion				-0.28*	-0.29*
				(0.040)	(0.033)
				[0.13]	[0.14]
Female			0.02		0.02
			(0.869)		(0.876)
			[0.10]		[0.10]
STEM			-0.08		-0.08
			(0.454)		(0.448)
			[0.10]		[0.10]

Table B.1: Change in proportion of subjects playing A in the testing game, restricted to subjects that passed the comprehension tests and that have demographic information recorded. N=100. All estimates calculated using a linear probability model, with p-values in brackets and standard errors in square brackets. * indicates value is significantly different from 0 using a non-directional test at the 5% level and ** indicates significance at the 1% level.

Effect of beliefs over	Conditional on	$\Delta Pr(X)$	$\Delta Pr(X)$	$\Delta Pr(X)$	$\Delta Pr(X)$
Ambiguity Aversion		0.03	0.04		
		(0.740)	(0.737)		
		[0.10]	[0.10]		
Risk Aversion		-0.17	-0.17		
		(0.112)	(0.104)		
		[0.10]	[0.10]		
Ambiguity Aversion	Low Risk Aversion			0.12	0.11
				(0.393)	(0.422)
				[0.14]	[0.14]
Ambiguity Aversion	High Risk Aversion			-0.08	-0.07
				(0.623)	(0.672)
				[0.16]	[0.16]
Risk Aversion	Ambiguity Neutrality			-0.08	-0.09
				(0.541)	(0.496)
				[0.14]	[0.14]
Risk Aversion	Ambiguity Aversion			-0.28	-0.27
				(0.083)	(0.091)
				[0.16]	[0.16]
Risk Aversion & Ambiguity Aversion				-0.16	-0.16
				(0.216)	(0.217)
				[0.13]	[0.13]
Female			0.09		0.08
			(0.350)		(0.410)
			[0.10]		[0.10]
STEM			0.00		0.01
			(0.974)		(0.903)
			[0.10]		[0.10]

Table B.2: Change in proportion of subjects playing X in the testing game, as a function of beliefs over opponent's preferences, restricted to subjects that passed the comprehension tests and have demographic data recorded. $N = 87$. All estimates calculated using a linear probability model, with p-values in brackets and standard errors in square brackets. * indicates value is significantly different from 0 using a non-directional test at the 5% level and ** indicates significance at the 1% level.

restricted to subjects who filled in the comprehension drop down menus correctly on the first attempt; 95 of the 206 subjects failed this test, with most of those subjects failing comprehension for the game that measured their beliefs. Table B.4 presents the relationship between a subjects own risk aversion and their beliefs regarding their opponents risk aversion. Only 35 subjects failed the comprehensions tests for the risk games, leaving a sample size of 171 subjects.

The most striking feature of tables B.3 and B.4 is the extremely strong relationship between the subject's own preferences and their beliefs regarding their opponents preferences. On both the risk and ambiguity dimension the Fisher exact test rejects the null hypothesis that preferences and beliefs are independent at all reasonable significance levels.

The results in the main text suggest that subjects were not very confident in their predictions of their opponent's preferences, while the results here demonstrate that subjects predicted their opponent to behave like themselves. A reconciliation of these facts is that when faced with predicting behaviour in an uncertain environment the subject's own behaviour acts as a very strong focal point.

		Predicted ambiguity preference		
		Neutral	Averse	
Own ambiguity preference	Neutral	63	4	67
	Averse	12	32	44
		75	36	111

Table B.3: Subject's own ambiguity preferences crossed with subject's predictions of their counterpart's ambiguity preferences, restricted to subjects who passed the comprehension test for the relevant games. Non-directional Fisher's exact test $p < 1 \times 10^{-10}$.

Result B.1. *Subjects believe that their opponents preferences are the same as their own preferences.*

B.3.3 Comprehension data

This section provides an analysis of the comprehension data, and a discussion of the robustness of the results presented in the main body of the paper. Table B.5 presents the subjects comprehension scores broken down by game number, and table B.6 presents the aggregated data. As the tables demonstrate, only about a quar-

		Predicted risk aversion			
		Low	Med	High	
Own risk aversion	Low	73	6	19	98
	Med	5	9	2	16
	High	9	2	46	57
		87	17	67	171

Table B.4: Subject's own risk aversion crossed with subject's predictions of their counterpart's risk aversion, restricted to subjects who passed the comprehension test for the relevant games. Non-directional Fisher's exact test $p < 1 \times 10^{-10}$.

ter of the subjects answered all of the comprehension questions correctly on their first attempt. Subjects performed better on the risk measurement game than the ambiguity measurement game, and also performed better as the row player (measuring their own preferences) than as the column player (measuring their beliefs regarding their opponent's preferences) in these games.

	Figure number of game					
Comp Score	2.2_R	2.2_C	2.5_R	2.5_C	2.6_R	2.6_C
1	159	131	196	177	168	179
0.75	26	39	5	18	32	21
0.5	7	13	2	5	2	4
0.25	9	4	0	1	1	0
0	5	19	3	5	3	2
Total	206	206	206	206	206	206

Table B.5: Number of subjects attaining each comprehension scores for each of the seven games played. Subjects were awarded a comprehension score of 1 for answering the comprehension questions correctly on the first attempt, and were penalized 0.25 for each incorrect attempt. The column heading 2.2_R (resp. 2.2_C) indicates the responses to the game in figure 2.2 as the row (resp. column) player.

There are obviously many different inclusion criteria that could be constructed using the comprehension data. The results in this paper were not data-mined by choosing the most convenient inclusion criteria. Ex-ante, two candidate inclusion criterion were identified. One was the criterion used in the body of the paper. The other criterion was to only use subjects who score a perfect 7 out of 7 on the comprehension questions. The latter criterion was unable to be used because of the small and unbalanced sample sizes that it produced. For example, of the 52

Total Comp Score	Number of Subjects
7	52
6.75	59
6.5	35
6.25	12
6	15
5.75	8
5.5	6
5.25	2
5	3
4.75	5
4.5	1
4.25	1
4	1
3.75	1
3.5	2
3.25	1
1.25	2
Total	206

Table B.6: Total comprehension score by subject. A score of 7 indicates that the subject answered the comprehension questions correctly on the first attempt for each of the 7 games.

subjects who satisfied the criterion only 8 were risk neutral and ambiguity averse, while only 4 were risk averse and ambiguity neutral. Therefore, this criterion was rejected for not providing enough power, and the criterion in the main text was adopted.

We also include, as an additional robustness check, the estimates of the main tables on the full sample of subjects, including those who failed the comprehension tests. Generally, the results do not differ markedly from those presented in the main text. The largest differences occur for column player behaviour in the original treatment in the main testing game, where subjects who failed the comprehension tests show a much larger effect of beliefs over their opponents' preferences on behaviour than was presented in the main text. It is unwise to ascribe intentionality to these subjects, given that they failed to demonstrate a sound knowledge of the payoffs. Nevertheless, we report the estimates here for completeness.

Table B.7 is the full sample analogue of Table 2.6. The only difference in statistical inference between the two samples is a stronger effect of ambiguity aversion conditional on high risk aversion.

Effect of	Conditional on	$\Delta Pr(A)$
Ambiguity Aversion		-0.21**
		(0.005)
Risk Aversion		-0.18*
		(0.017)
Ambiguity Aversion	Low Risk Aversion	-0.12
		(0.229)
Ambiguity Aversion	High Risk Aversion	-0.28*
		(0.027)
Risk Aversion	Ambiguity Neutrality	-0.07
		(0.485)
Risk Aversion	Ambiguity Aversion	-0.23
		(0.051)
Ambiguity Aversion & Risk Aversion		-0.35**
		(0.001)

Table B.7: Change in proportion of subjects playing A in the testing game, as a function of subject preferences. $N = 185$. Pearson's χ^2 p-value shown in brackets. * indicates value is significantly different from 0 using a non-directional test at the 5% level and ** indicates significance at the 1% level.

Table B.8 is the full sample analogue of table 2.7. The effect of ambiguity aversion being associated with a move from A to C amongst subjects with low risk aversion is no longer apparent in the full sample. This change is caused by subjects who failed the comprehension tests playing B with an essentially equal probability across all subject Types.

Table B.9 is the full sample analogue of table 2.9. We see a strong effect difference in behaviour between Type $1'$ and Type $4'$ subjects (bottom row of the table) in the full sample. This effect is being driven almost entirely by subjects who failed the comprehension tests, and therefore should be treated with caution.

B.4 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 meet the formating requirements of this dissertation, but are materially the same as those presented to the subjects.

B.4.1 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.

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.

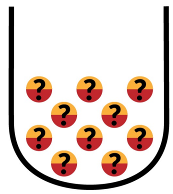

Figure B.3: U bag.

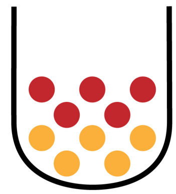

Figure B.4: K bag.

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.

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.

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

	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

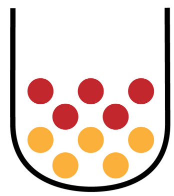

Figure B.5: K bag.

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 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

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).

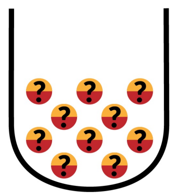

Figure B.6: U urn.

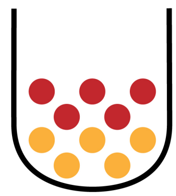

Figure B.7: K bag.

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

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.

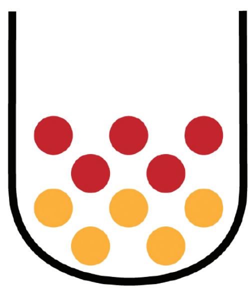

K bag

Figure B.8: Game X.

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.

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

	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 B.9: Game XT.

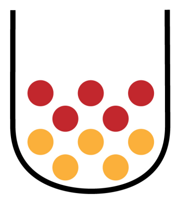

Figure B.10: K bag.

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 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 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.

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.

Effect of	Conditional on	$\Delta\frac{N_B}{N_B+N_A}$	$\Delta\frac{N_C}{N_C+N_A}$
Ambiguity Aversion		0.13	0.19*
		(0.053)	(0.013)
Risk Aversion		-0.02	0.23**
		(1.000)F	(0.003)
Ambiguity Aversion	Low Risk Aversion	0.12	0.06
		(0.186)F	(0.510)
Ambiguity Aversion	High Risk Aversion	0.16	0.27*
		(0.288)F	(0.039)
Risk Aversion	Ambiguity Neutrality	-0.06	0.11
		(0.675)F	(0.269)
Risk Aversion	Ambiguity Aversion	-0.02	0.31*
		(1.000)F	(0.015)
Ambiguity Aversion & Risk Aversion		0.10	0.38**
		(0.373)F	(0.000)

Table B.8: Change in the values of $\frac{N_B}{N_B+N_A}$ and $\frac{N_C}{N_C+N_A}$ in the testing game as a function of subject type, where N_A is the number of subjects selecting A. $N = 185$ Figures in brackets are p-values calculated under a null hypothesis that the coefficient is equal to zero under either Person's χ^2 test or a Fisher exact test (denoted by F). The coefficients are equivalent to the probabilities implied by a saturated multinomial logit regression with a choice of A denoted as the base outcome.

Effect of beliefs over opponent's	Conditional on	$\Delta Pr(X)$
Ambiguity Aversion		-0.13
		(0.069)
Risk Aversion		-0.15*
		(0.033)
Ambiguity Aversion	Low Risk Aversion	-0.06
		(0.491)
Ambiguity Aversion	High Risk Aversion	-0.18
		(0.111)
Risk Aversion	Ambiguity Neutrality	-0.09
		(0.268)
Risk Aversion	Ambiguity Aversion	-0.21
		(0.083)
Ambiguity Aversion & Risk Aversion		-0.28**
		(0.006)

Table B.9: Change in proportion of subjects playing X in the testing game, as a function of reported beliefs over opponents' preferences. $N = 184$. Pearson's χ^2 p-value shown in brackets. * indicates value is significantly different from 0 using a non-directional test at the 5% level and ** indicates significance at the 1% level.

Appendix C

Appendix for “Continuity, Inertia and Strategic Uncertainty: A Test of the Theory of Continuous Time Games”

C.1 Instructions to Subjects

Instructions

You are about to participate in an experiment in the economics of decision-making. If you follow these instructions carefully and make good decisions, you can earn a CONSIDERABLE AMOUNT OF MONEY, which will be PAID TO YOU IN CASH at the end of the experiment.

Your computer screen will display useful information. Remember that the information on your computer screen is PRIVATE. To insure best results for yourself and accurate data for the experimenters, please DO NOT COMMUNICATE with the other participants at any point during the experiment. If you have any questions, or need assistance of any kind, raise your hand and the experimenter will come and help you.

In the experiment you will make decisions over several periods. At the end of the last period you will be paid, in cash, the sum of your earnings over all periods.

The Basic Idea.

In each of several **periods**, you will be randomly matched with another participant for 60 seconds and you will each decide when to **Enter the market**. If you both enter at the same time, you will both earn the same amount, which will depend on the time at which you both entered. If one player enters earlier than the other, she will earn more money while her counterpart will earn less. The longer the second player waits to enter after the first player enters, the greater the difference in their earnings.

Screen Information.

A vertical dashed line marks the passage of time, moving from left to right over the course of the period until both players have chosen to **Enter the Market.** A dot moving left to right labeled 'Me' and another labeled 'Other' shows payment information for each player, although before either player has entered the market the two dots will be precisely on top of each other. The dots show the amount of money each subject will earn if both players have entered now If nobody has entered yet the dots (which are on top each other) show what will happen if both enter now. If one player has entered already the dots (which are now separated) will show the amount each player will earn if the second player enters the market now.

You can choose to Enter the market at any time by pressing the space bar. The time of entry will be shown on the screen as a dashed vertical line.

The screen gives you information on your potential earnings under three possible scenarios:

If you both enter in the same sub-period

If you and your counterpart enter at the same time, you will earn exactly the same amount as your counterpart. The **black line** that looks like a hill shows

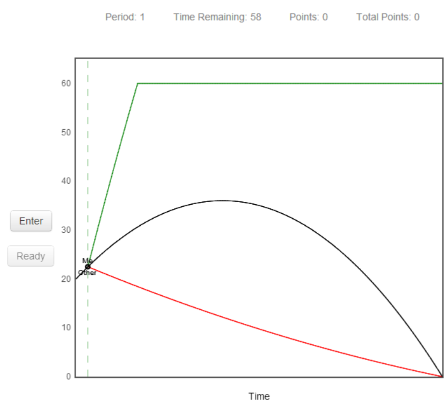

Figure C.1: Nobody has entered.

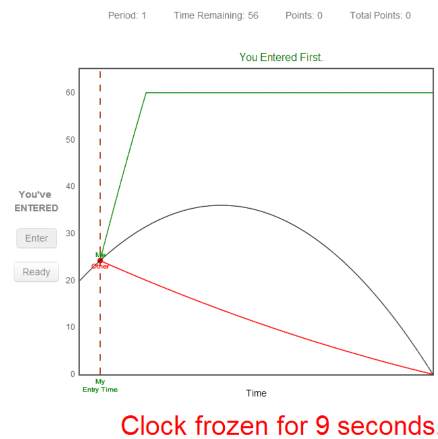

Figure C.2: One player has decided to enter and the clock has frozen.

exactly what you and your counterpart would both earn if you both entered in each moment of the period.

Notice that your joint earnings depend on when you both choose to enter in. For example, if you and your counterpart both entered at time 0, you would both earn 20 points. However, if you both entered at time 24, you would both earn approximately 36 points.

If you enter first

If your counterpart enters later than you, she will earn less and you will earn more than if she had entered at the same time you did. At every moment the screen tells you what would happen if you entered now, and your counterpart entered at a later time than you:

- The green line shows you what you would earn if you entered now and your counterpart entered in each of the remaining moments in the period.
- The red line shows you what your counterpart would earn if you entered now and she entered in each of the remaining moments in the period.

Notice that the longer your counterpart waits to enter after you, the less she earns. For instance in the example in Figure 1, if you entered now and your counterpart entered 5 seconds later, you would earn approximately 38 points (the amount on the green line 5 seconds later) and your counterpart would earn approximately 20 points (the amount on the red line 5 seconds later). If, instead, your counterpart waited 10 seconds to enter, you would earn approximately 50 points and your counterpart approximately 18 points.

Note that these lines will change as you move along: the green line will always be above the current point on the black hill and the red line will always be below, reflecting the fact that you earn more (and your counterpart less) than if your counterpart entered when you did.

If you enter second

If you enter at a later time than your counterpart, you will earn less and your counterpart more than if you had moved at the same time as your counterpart. Importantly, these graphs and payoffs are symmetric: your counterpart sees the same screen (at least prior to anyone entering) and faces the same payoff consequences as you do. Thus it is also true that:

- The green line shows you what your counterpart would earn if (s)he entered now and you entered later at a future time.
- The red line shows you what you would earn if your counterpart entered now and you entered at a later time.

For instance in the example in Figure 1, if your counterpart entered now and you entered in 5 seconds, your counterpart would earn approximately 38 points (the amount on the green line 5 seconds to the right) and you would earn approximately 20 points (the amount on the red line 5 seconds to the right). If, instead, you waited 10 seconds to enter, you would earn approximately 18 points and your counterpart approximately 50 points.

Time Freeze

After a player first enters, the game will **freeze** for 5 seconds. During these 5 seconds the player's counterpart can choose whether to enter too by pressing the spacebar. If (s)he does, the software will treat both entry decisions as occurring **at the same time** and both players will earn the exact same amount (the amount shown on the black line at the moment of entry). If she does not choose to enter during the time freeze, the clock will resume and her earnings will drop, following the red line. The time freeze is demonstrated in Figure 2.

After Entry Occurs

If one player enters before the other, you will see a dotted green line mark the time of first entry and you will see the dots separate, one following the green line and the other red line. A label next to each dot will tell you which corresponds to you (labeled "Me") and which to your counterpart (labeled "Other"). A message at the top of the screen will also remind you whether you were the first to enter (in green) or the second (in red). If you were not the first to enter, the timing of your entry decision will now determine both of your earnings. Figure 3 shows an example.

After both players enter, horizontal lines will appear showing your earnings. At the end of these lines (on the right side of the screen) you will see your and your counterparts' exact earnings (if you both entered at the same time, you will both earn the same amount). Figure C.4 shows an example in which one player has entered at time 10 and the other later at time 14. As a consequence the first to enter (in green) earns 45.01 points and the second (in red) player earns 26.23 points.

Other Information and Earnings.

At the top of the screen you will see the current period number, the time remaining in the current period, the number of points you will earn this period based on current decisions and the total number of points you have accumulated over all periods so far. You will be paid cash for each total point you have earned at the end of the experiment, at a rate given by the experimenter.

At the beginning of each new period, you will be randomly matched with a new participant from the room to play again. All matchings are anonymous – you will never know who you were matched with and neither will they.

Summary

- Your and your counterparts' earnings depend on the time you each decide to Enter the market (by pressing the spacebar).
- When either player enters, the clock will freeze and her counterpart will have the opportunity to enter too, at the same time.
  - If you both enter at the same time, you will earn the same amount, shown on the black hill-like line. As the black hill line shows, your joint earnings depend on the time at which you both enter.
  - If your counterpart enters at a later time than you, you will earn more and she will earn less than if she had entered when you did. These amounts are shown via the green and red lines.
  - Likewise, if you enter at a later time than your counterpart, you will earn less and your counterpart more than if you had entered when she did.
- You will be paid based on the total number of points you earn, over all periods.

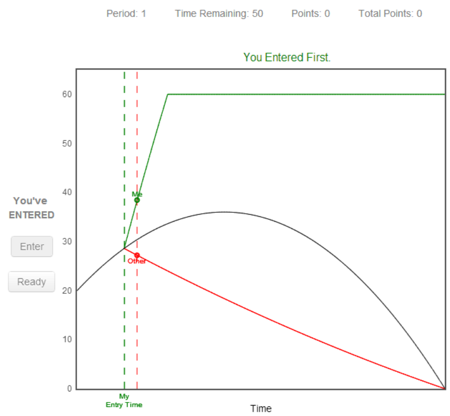

Figure C.3: The green player has entered but the red player has not.

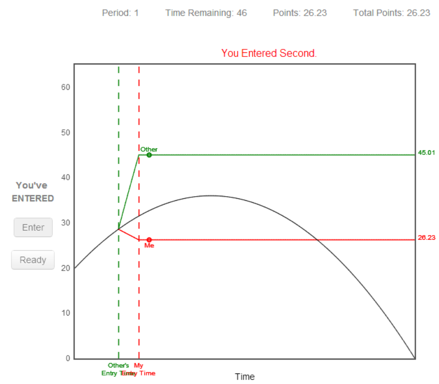

Figure C.4: Both have entered.

C.2 Nash Equilibrium

In section C.2.1 we provide proofs for propositions stated in section 3.2.2 of the main paper. In section C.3 we prove a set of propositions characterizing ε-equilibrium for our games and provide proofs for propositions stated in section 3.2.3 of the main paper.

C.2.1 Nash equilibrium

In subsection C.2.1 we prove proposition 3.1, which relies exclusively on standard theoretical tools.

In subsection C.2.1 we provide more details on the theoretical schema of Simon and Stinchcombe (1989) and provide a self contained heuristic proof of proposition 3.2 following the intuition of Simon and Stinchcombe (1989). We close the subsection by providing a full proof of proposition 3.2 that draws directly on theoretical machinery developed in Simon and Stinchcombe (1989). Because this machinery is substantial, this full proof is not self contained and relies on definitions and lemmas developed in Simon and Stinchcombe (1989).

In subsection C.2.1 we provide some details on the theoretical apparatus of Bergin and MacLeod (1993), an alternative way of modeling continuous time games, and specialize it to our setting. We then re-prove proposition 3.2 using this method. Again, because the required tools are substantial, this proof is not self contained but draws on definitions and results from Bergin and MacLeod (1993). Finally, we use ideas from Bergin and MacLeod (1993) to prove proposition 3.3.

Proof of proposition 3.1

Proof of Proposition 3.1. We prove by contradiction. First, consider the case of pure strategies: Suppose that there is an equilibrium where player i enters at a weakly later grid point than player j, and that this is not the first grid point. Player i has a profitable deviation, which is to enter one grid point before player j or at the first grid point if there are no grid points before player j's entry time. Contradiction.

The modification to allow for mixed strategies is straightforward. Call the last grid point that player i places positive weight on k_i. Suppose that there is an equilibrium with $k_i \geq k_j$ and $k_i \neq 0$ (where we are labeling $t = 0$ as the zeroth grid point). Player i can improve their payoff by moving weight from entry at k_i to entry at the grid point identified by $\max\{k_j - 1, 0\}$. □

Simon and Stinchcombe (1989)

In this section, following Simon and Stinchcombe (1989), we model our game by considering a set of strategies that are unambiguously defined in the limit of a discrete time grid as the grid becomes infinitely fine. Together, the two player game generated by the utility function in equation 3.1, the histories H_t and strategies s_i given below define what we call the Perfectly Continuous time game.[1] We begin by defining a history, and use the history to define a set of strategies.

In our timing game, a history consists of two pieces of information: the current time, and a record of players that have already entered the market. Therefore, a history at time t, H_t, is an object in $[[0,t] \cup \{1\}]^2$, where the first element denotes the entry time for firm i, and the second element the entry time for firm j, with 1 indicating that a firm has not yet entered. For example, a history at time $t = 0.5$ can be written $H_{0.5} = \{0.4, 1\}_{0.5}$, indicating that the first firm entered at time $t = 0.4$, the second firm is yet to enter, and the subscript denotes the current time.

[1]The strategies also implicitly define the action sets and the allowable order of moves, and the game is of complete and perfect information.

Given this definition of histories, a strategy is a mapping $s_i : H_t \mapsto \{0,1\}$, where 1 indicates an action of "enter" at that history. We fix $s_i(\{t', \cdot\}_t) = 1$ for all $t' < t$ (after a firm has entered they cannot leave the market again), require that s_i is a piecewise continuous function, and also require that if $s_i(\{1,t'\}_t) = 1$ then $s_i(\{1,\hat{t}\}_t) = 1$ for all $\hat{t} \leq t'$ (if, at time t, I wish to enter in response to an opponent who entered at time t' then I must also enter in response to any opponent that entered prior to t'). Notice that these strategies satisfy assumptions F1 - F3 of Simon and Stinchcombe (1989).[2]

Most of the technical complications that arise from modeling games in continuous time are related to the "well-ordering paradox" that, in continuous time, the notions of "the next moment" or "the previous moment" have no meaning. As Simon and Stinchcombe (1989) point out this means that sensible sounding strategies such as "I play *left* at time zero; at every subsequent time, play the action I played last period" are not well defined. The approach that Simon and Stinchcombe take is to restrict attention to strategies that are uniquely defined as the limit point for *any* discrete time grid as the grid becomes arbitrarily fine; their conditions F1-F3 are designed to generate such strategies.

We present two proofs of proposition 3.2 in this section that follow this modeling approach. The first proof, which we call the "heuristic proof" does not directly address many of the complexities associated with continuous time, other than to restrict the strategy set to the strategies defined above. The question of whether this restriction is appropriate is not undertaken here. For this reason we also provide an alternative proof, resting almost entirely on results in Simon and Stinchcombe (1989), which demonstrates that our equilibria are sound. We also provide a heuristic proof of remark 3.1.

Heuristic proof of proposition 3.2. Consider, the strategies:

$$s_i^{t'}(H_t) = \begin{cases} 0 & \text{if } H_t = \{1,1\}_t \text{ with } 0 \leq t < t' \leq t^* \\ 1 & \text{otherwise} \end{cases} \tag{C.1}$$

[2]The conditions require a bounded number of changes in action for each player, piecewise continuity with respect to time and strong right continuity with respect to histories. The first two are obviously satisfied here, and the requirement that $s_i(\{1,t'\}_t) = 1$ then $s_i(\{1,\hat{t}\}_t) = 1$ for all $\hat{t} \leq t$ ensures strong right continuity with respect to histories.

Each of these strategies generates simultaneous entry at t'. To aid intuition, we focus on these strategies precisely because they are strategies such that there is no delay in responding to an opponent's entry in any history. Such a delay is (weakly) payoff decreasing, and certainly cannot be observed in an equilibrium. More formally, this implies ruling out all strategies $s_i(H_t) = 0$ where $H_t = \{1, t'\}$ for $t' < t$.

We claim that the strategies given in equation C.1 form a set of (symmetric) equilibria. That is, the pair of strategies $(s_i^{t'}(H_t), s_j^{t'}(H_t))$ are an equilibrium of our Perfectly Continuous time game. To see that this is true, consider two cases.

In the first case, consider a strategy that deviates from $s_i^{t'}(H_t)$ by replacing $s_i^{t'} = 1$ with $s_i'^{t'} = 0$ for some set of histories. Given our restriction that a firm cannot re-enter after it has exited, this can only occur for histories in which the firm has not yet entered. This implies that these must be histories in which either our opponent has entered, or no one has entered. Histories in which no one has entered at times after t' are off the equilibrium path, so the change in strategy profile has no effect on payoffs. Histories in which the opponent has entered will either be off the equilibrium path, or on a path for which $U(s_i^{t'}, s_j^{t'}) > U(s_i'^{t'}, s_j^{t'})$ (this is because the payoff for the second entrant is a decreasing function of their entry time).

In the second case, consider a strategy that deviates from $s_i^{t'}(H_t)$ by replacing $s_i^{t'} = 0$ with $s_i'^{t'} = 1$ for some set of histories. It must be the case that at least one of these histories is on the equilibrium path: s' will cause the firm to enter before t'. Given the equilibrium strategies this deviation will generate an instantaneous reply from the firm's opponent. We shall now observe joint entry at some time prior to t'. Given that the payoff to joint entry is increasing on the interval $[0, t^*]$, such a deviation reduces the firms payoff.

$\square$

Heuristic proof of remark 3.1. We proceed somewhat unusually and consider the second round of iterated elimination of weakly dominated strategies first.

Consider the strategies used to build the equilibrium in the heuristic proof of proposition 3.2. Given that we restrict attention to only these strategies, each of the strategies of the type given in equation C.1 is weakly dominated by the strategy

$$s_i(H_t) = \begin{cases} 0 & \text{if } H_t = \{1,1\} \text{ with } 0 \le t < t^* \\ 1 & \text{otherwise} \end{cases} \tag{C.2}$$

which generates joint equilibrium entry at t^*.

We now show that any strategy not in equation C.1 is weakly dominated by a strategy that is in equation C.1.

Consider the class of strategies that has $s_i(\{1,1\}_t) = 0$ for all $t < t'$ and $s_i(\{1,1\}_{t'}) = 1$. Each strategy in this class is weakly dominated by $s_i^{t'}$ as defined in equation C.1.[3] If their opponent is yet to enter by $t = t'$ then all strategies in the class provide the same utility. The only way that strategies in the same class can differ is in how they respond to entry from their opponent: because $s_i^{t'}$ responds instantly in all cases, each other strategy in the class must have a delay in response for at least some opponent entry time. Given that payoffs are decreasing in the entry time of the second firm, the dominance relation is established.

All strategies must fall into one of the above classes.

The strategy in equation C.2 is therefore the unique strategy that survives iterated elimination of weakly dominated strategies.

□

We conclude this section with a more formal (but less self contained) proof of proposition 3.2, drawing more directly on definitions and results from Simon and Stinchcombe (1989).

Formal proof of Proposition 3.2. The proof consists entirely of demonstrating that Theorem 3 of Simon and Stinchcombe (1989) is applicable to our environment. We begin by stating this theorem. The definitions of all relevant notation can be found in Simon and Stinchcombe (1989).

Theorem C.1 (Theorem 3 from Simon and Stinchcombe (1989).)**.** *Consider a continuous-time game with a d^H-continuous valuation function. Let f be a continuous-time strategy profile satisfying F1-F3. Suppose that there exists a sequence of δ^n-fine grids, (R^n), where $\delta^n \to_n 0$ and a sequence (g^n, ε^n) such that $\varepsilon^n \to_n 0$ and for each n, g^n is an ε^n-SGP equilibrium for the game played on R^n. Further suppose*

[3] Or, if $t' > t^*$ in violation of equation C.1, then the strategy is weakly dominated by $s_i^{t^*}$.

that g^n is defined by further restricting $f_{|R^n}$ to the R-admissible decision nodes. Then f is an SGP equilibrium for the continuous time game.

We proceed in 5 steps:

1. The game has a d^H-continuous valuation function because all games with integrable flow payoff functions, such as ours, have a d^H-continuous valuation function.

2. Label the strategy profile that induces entry at time t in equation C.1 as $f(t)$. $f(t)$ satisfies Simon and Stinchcombe (1989)'s conditions F1-F3.

3. Take the sequence of grids to be grids with $G = n$ uniformly spaced points.

4. Define $g(t)^n$ to be the strategy profile where each agent plays "enter" at all grid points that occur (weakly) after t. At all grid points (strictly) before t, each agent plays "wait" if both agents have played "wait" at all previous grid points and otherwise they play "enter". This strategy is the restriction $f_{|R^n}$ to the R-admissible decision nodes.

5. It follows immediately from the proof of proposition C.3 that $g(t)^n$ forms an $\varepsilon(t)^n$-equilibrium where $\varepsilon(t)^n$ is defined by replacing the inequality in equation C.4 with equality. Furthermore, $\varepsilon(t)^n \to 0$ as $G = n \to \infty$.

We therefore conclude that $f(t)$ is an equilibrium of the continuous time game. We finish by noting that the strategy $f(t)$ induces an outcome where both agents enter at time t, and that the definition of $f(t)$ allows $0 < t \leq t^*$. Joint entry at $t = 0$ is obviously also supported by an equilibrium (if your opponent enters at $t = 0$ then your best response is also to enter at $t = 0$). Therefore, we can support entry at any time $t \in [0, t^*]$ as an equilibrium in the continuous time game.

□

Bergin and MacLeod (1993)

Bergin and MacLeod (1993) take an alternative approach to modeling continuous time games by introducing a general notion of inertia and modeling continuous time as the limit as inertia disappears.[4] In this subsection, we first formally define a narrower notion of reaction-lag-based inertia appropriate for the setting of our experiment. We then prove proposition 3.3 which claims that short of the condition limit (i.e. when inertia is greater than zero) firms must enter immediately in equilibrium. Finally, we provide an alternative proof of proposition 3.2 using proposition 3.3 and Bergin and MacLeod (1993)'s Theorem 3. By doing so, we show that in our game the two approaches to modeling Perfectly Continuous time – e.g. as the limit of Perfectly Discrete time á la Simon and Stinchcombe (1989) or as the limit of Inertial Continuous time á la Bergin and MacLeod (1993) – make identical predictions.[5]

We introduce a simplified version of Bergin and MacLeod (1993)'s definition of inertia that is appropriate for our game, capturing the key idea of inertia as a reaction lag.[6]

Definition C.1 (Inertia)**.** *Fix a time $\hat{t} \in [0,t^*]$. Suppose that $s_i(\{1,1\}_t) = 0$ for all $t < \hat{t}$ and that $s_i(\{1,1\}_{\hat{t}}) = 1$. A strategy satisfies* inertia *if there exists a $\delta > 0$ such that $s_i(\{1,t'\}_t) = 0$ for all t,t' such that $t' \leq t < \min\{\hat{t}, t' + \delta\}$.*

Our inertia condition prevents firms from responding immediately to the entry decisions of their opponents - responses are delayed by at least $\delta > 0$ of the game.[7] We can think of $\hat{t}$ as a firm's planned entry time – if their opponent has not entered yet, then the firm plans to enter at $\hat{t}$. The inertial requirement then states that the firm cannot enter within δ of their opponent's entry time, unless they are entering at their planned entry time.

We define Inertial Continuous time simply as Perfectly Continuous time restricted to strategies that satisfy the inertia condition of definition C.1 and can use

[4]A key technical innovation of Bergin and MacLeod (1993) is that it allows one to model continuous time in infinitely repeated games. Infinitely repeated games violate the Simon and Stinchcombe (1989) assumption that an agent will change their behaviour a finite number of times.

[5]We conjecture that this is also true more generally, whenever the relevant conditions of both Bergin and MacLeod (1993) and Simon and Stinchcombe (1989) are satisfied.

[6]Bergin and MacLeod (1993) assume that the action space is constant across time, violating a condition in our experiment: we do not allow firms to exit the market once they have entered (i.e. once a firm has exited their action space shrinks from two actions to only one action). This technical problem can be neatly resolved by imposing an arbitrarily large amount of inertia to all histories in which the firm has already entered.

[7]Recall that we define $\delta \equiv \frac{\delta_0}{T}$, where δ_0 is a subject's natural reaction lag in real time and T is the game length in real time.

this to provide a proof of proposition 3.3.

Proof of proposition 3.3. The pure strategy case provides the intuition for the full proof, so we first prove that there are no pure strategy equilibria with delayed entry.

We proceed by contradiction. Suppose that there is a pure strategy equilibrium where the earliest entry is not at time 0. Without loss of generality assume that firm i enters weakly before firm j at time $t > 0$. It is easy to verify that firm j has a profitable deviation: entering at a point that is δ before t (or at 0 if $t < \delta$). This contradicts the assumption that the strategies form an equilibrium.

To deal with mixed strategies, the proof by contradiction needs only minimal changes. Write $t_j > 0$ and t_i for the latest entry time used by firms j and i in a mixed strategy equilibrium for histories where their opponent has yet to enter. We proceed by demonstrating that there is a strategy in the support of player j's mix that is not optimal.

Suppose that $t_j \geq t_i$. Consider the history where we arrive at time $\max\{t_i - \delta, 0\}$ and neither firm has entered yet. Clearly firm j's strategy is not optimal: given firm i's strategy, and the fact that we reach the current subgame with positive probability, firm j should enter immediately with certainty. This contradicts the assumption that the strategies form an equilibrium. □

The logic behind the proofs of propositions 3.1 and 3.3 does not apply to proposition 3.2. The reason is that in continuous time a firm can respond instantaneously to its opponent's entry. The logic of the proof by contradiction thus does not hold: firm j's best response is no longer to preempt firm i, but to enter simultaneously with firm i. We conclude the section with an alternative proof of proposition 3.2.

Alternative proof of Proposition 3.2 using Bergin and MacLeod (1993). The proof consists entirely of demonstrating that Theorem 3 of Bergin and MacLeod (1993) is applicable to our environment. There are some technical hurdles to be surmounted as Bergin and MacLeod (1993) define strategies as being mappings from outcomes to actions, rather than the more standard mappings from histories to actions. Fortunately, in our game, there is a straightforward mapping from histories to outcomes that simplifies matters greatly.

We first state Bergin and MacLeod's Theorem 3:

Theorem C.2 (Theorem 3 from Bergin and MacLeod (1993)). *A strategy $x \in S^*$ is a subgame perfect equilibrium if and only for any Cauchy sequence, $\{x^n\}$ converging to x, there is a sequence $\varepsilon^n \to 0$, such that x^n is an ε^n-subgame perfect equilibrium.*

To start, we note that although the theorem implies that all convergent Cauchy sequences need to form a convergent sequence of ε-equilibrium it is clear from Bergin and MacLeod's proof that it is sufficient to find a single satisfactory sequence. Once a single sequence is found, part (b) of their proof implies that x is an equilibrium and then part (a) of their proof implies that all other convergent sequences must also have associated convergent ε paths.

We note again that Bergin and MacLeod (1993) use a special domain for their strategies. The domain of x is the object (A,T) where $A = \{t_i, t_j\}$ is a set of outcomes (one for each player) and $T = [0,1]$ is the set of feasible entry times. Clearly, we can translate each of our strategies into the Bergin and MacLeod (1993) formulation. For example, the strategy that induces player i to enter at the earlier of time τ and immediately after their opponent enters could be written as:

$$x_1((\cdot,t_j),t) = \begin{cases} 0 & \textit{if } t < \tau \textit{ and } t < t_j \\ 1 & \textit{otherwise} \end{cases} \tag{C.3}$$

The set of strategies that satisfy our inertia condition (definition C.1) maps into a subset of Bergin and MacLeod's set of strategies S. If we add the set of strategies that are formed in the limit as $\delta \to 0$, then we have a set of strategies that are a subset of Bergin and MacLeod's set of strategies S^*.[8]

Now, for each time $\tau \in [0,t^*]$, let $x(\tau)^n$ be a sequence of strategies for both players that satisfy definition C.1 with $\tau = \hat{t}$ and $\delta = \delta^n$ such that $\delta^n \to 0$. Then $x(\tau)^n$ Cauchy converges to the strategy given in equation C.3. Furthermore, from the proof of proposition C.4, for each $x(\tau)^n$ we can identify an ε^n, which is found by replacing the inequality in equation C.7 with an equality, such that the strategies form an ε^n-equilibrium. Clearly, the sequence $\varepsilon^n \to 0$ as $\delta^n \to 0$.

Therefore, we conclude that the strategy $x(\tau)$, which generates joint entry at $t = \tau$, forms an equilibrium of the continuous time game. Given that this is true for

[8]Note that the strategy presented in equation C.3 is in S^* but not in S because it does not satisfy the inertia condition.

all $\tau \in [0,t^*]$, we can sustain all such entry times in equilibrium.

□

C.3 ε-equilibrium

In subsection C.3.1 we prove three propositions that completely characterize the ε-equilibrium sets for each of our timing protocols. In appendix C.3.2 we use these characterizations to prove comparative static propositions stated in section 3.2.3 that form the basis of alternative hypotheses for the experiment.

Following Radner (1980) we assume that players are willing to tolerate a payoff deviation from best response of size ε and treat their counterparts as having the same tolerance. We can then form a set of propositions that provide alternative predictions to Nash equilibrium as a function of ε for each of our three main protocols.

C.3.1 Characterization of ε-equilibrium sets

Again, we begin by considering Perfectly Discrete Time. Recall that for a game with grid size G, dates begin at $t = \left\{0, \frac{1}{G}, \frac{2}{G}, ..., 1\right\}$ (a total of $G+1$ dates). For notational simplicity, assume that the grid is such that $t^* = 1 - \frac{\Pi_D}{4c}$ lies exactly on some grid point (this assumption simplifies the following expressions and is imposed in the parameterization of our experiment), and label this gridpoint $\frac{k}{G}$.

Proposition C.3. *Suppose that all agents enter at or before t^*.*[9] *In a Perfectly Discrete time game with $G+1$ periods, $t^* = \frac{k}{G}$ and tolerance ε the set of entry times that can be sustained in a pure strategy sub-game perfect ε-equilibrium is given by $\left\{0, \frac{1}{G}, ..., \frac{\kappa}{G}\right\}$ where κ is the largest integer that satisfies $0 \le \kappa \le k$ and*

$$\frac{1}{G}\left[\frac{3\Pi_F}{2} - c(2+\frac{1}{G})\right] + \frac{\kappa}{G^2}\left[2c - \frac{\Pi_F}{2}\right] \le \varepsilon. \tag{C.4}$$

If no non-negative integer satisfies equation C.4 then $\kappa = 0$ (i.e. the unique equilibrium is immediate entry).

[9] If the ε-equilibrium set includes t^*, it is possible for entry times after t^* to also be supported as ε-equilibria. Such entry times never arise in the data and violate the conceptual spirit of ε-equilibrium outlined by Radner and we assume them away purely to simplify the discussion and notation.

Proof of proposition C.3. We begin with two assumptions. Firstly, we assume that all firms enter at or before t^*. Secondly, when one firm enters the other firm shall enter as soon as possible afterwards (in the next sub period). We shall demonstrate at the end of the proof that the second assumption can be disposed of, but imposing it simplifies the proof.

We proceed by backwards induction. Suppose that firms arrive at period $k-1$ and neither firm has entered yet. We are interested in determining whether cooperation can be sustained for one more period. If it can, then the payoff to each firm will be $U(\frac{k}{G},\frac{k}{G})$. If, however, a firm defects and enters immediately then they will earn $U(\frac{k-1}{G},\frac{k}{G})$.

After some algebra, we see that $U(\frac{k-1}{G},\frac{k}{G})-U(\frac{k}{G},\frac{k}{G})=\frac{1}{G}\left[\frac{3\Pi_F}{2}-c(2+\frac{1}{G})\right]+\frac{k}{G^2}\left[2c-\frac{\Pi_F}{2}\right]$.

Therefore, we conclude that if neither firm has entered when firms arrive at period $k-1$, then joint entry at period k can be sustained as an ε-equilibrium if

$$\varepsilon \geq U(\frac{k-1}{G},\frac{k}{G})-U(\frac{k}{G},\frac{k}{G}). \tag{C.5}$$

Now, rollback to period $k-2$. Suppose that firms arrive at period $k-2$ and neither firm has entered yet. Can cooperation be sustained for one more period? There are two cases.

In the first case equation C.5 holds, so that if both firms wait at period $k-2$ then there is an equilibrium continuation where each firm earns $U(\frac{k}{G},\frac{k}{G})$. If, however, a firm enters at period $k-2$ they will earn $U(\frac{k-2}{G},\frac{k-1}{G})$. Waiting can therefore be sustained as an ε-equilibrium if $\varepsilon \geq U(\frac{k-2}{G},\frac{k-1}{G})-U(\frac{k}{G},\frac{k}{G})$. This inequality must hold whenever equation C.5 holds because the utility function is increasing in (joint delay of) entry times. We therefore conclude that if both firms may wait in an ε-equilibrium at period $k-1$ then they may also wait in an ε-equilibrium at period $k-2$ (and the same logic implies this must be true for any earlier period as well).

In the second case equation C.5 does not hold. Therefore, if both firms wait at period $k-2$ they must both enter at period $k-1$ and the continuation payoff is $U(\frac{k-1}{G},\frac{k-1}{G})$ for both firms. Defecting to immediate entry will earn $U(\frac{k-2}{G},\frac{k-1}{G})$.

After some algebra, we see that $U(\frac{k-2}{G},\frac{k-1}{G})-U(\frac{k-1}{G},\frac{k-1}{G})=\frac{1}{G}\left[\frac{3\Pi_F}{2}-c(2+\frac{1}{G})\right]+$

$\frac{k-1}{G^2}\left[2c-\frac{\Pi_F}{2}\right] < U(\frac{k-1}{G},\frac{k}{G}) - U(\frac{k}{G},\frac{k}{G})$.

Therefore, we conclude that if neither firm has entered when firms arrive at period $k-2$, then joint entry at period $k-1$ can be sustained as an ε-equilibrium if $\varepsilon \geq U(\frac{k-2}{G},\frac{k-1}{G}) - U(\frac{k-1}{G},\frac{k-1}{G})$. Notice that the ε cutoff at period $k-2$ is less than the cutoff at $k-1$. Therefore, there are some ε values where co-operation can be sustained at period $k-2$, but not at period $k-1$.

The proof continues by induction, rolling back one period at a time and noting that $U(\frac{k-n-1}{G},\frac{k-n}{G}) - U(\frac{k-n}{G},\frac{k-n}{G})$ is decreasing in n, and substituting $n=\kappa$ delivers the equation in the proposition.

We still need to demonstrate that relaxing our initial assumption that entry by one firm would be immediately followed by entry from the other will not change the equilibrium set found above. There are two concerns: firstly, that allowing lengthier delays might destroy some equilibria described above and, secondly, that allowing lengthier delay might create new equilibria.[10]

No equilibria can be destroyed. Take an equilibrium found above, and allow an agent to unilaterally deviate to a strategy that involves lengthy delays after their opponent's entry (in at least some subgames). This deviation must weakly decrease the agent's payoff, so it is not a profitable deviation and the original equilibrium survives.

New equilibria can be created, but none of these equilibria will involve an initial entry time that is later than the set of entry times described in the proposition. We demonstrate this with a proof by contradiction. For a given ε, define κ as given by equation C.4. Now, suppose that there is an equilibrium with entry at the pair of nodes $(\frac{\kappa+j}{G},\frac{\kappa+j+i}{G})$ with $j \geq 1$ and $i \geq 2$. Then, from the definition of ε-equilibrium it must be the case that $U(\frac{\kappa+j-1}{G},\frac{\kappa+j}{G}) - U(\frac{\kappa+j+i}{G},\frac{\kappa+j}{G}) \leq \varepsilon$ (otherwise the second entrant would have a profitable deviation).

Now, it is also true that $U(\frac{\kappa+j-1}{G},\frac{\kappa+j}{G}) - U(\frac{\kappa+j+i}{G},\frac{\kappa+j}{G}) > U(\frac{\kappa+j-1}{G},\frac{\kappa+j}{G}) - U(\frac{\kappa+j+1}{G},\frac{\kappa+j}{G}) \geq U(\frac{\kappa}{G},\frac{\kappa+1}{G}) - U(\frac{\kappa+2}{G},\frac{\kappa+1}{G}) > U(\frac{\kappa}{G},\frac{\kappa+1}{G}) - U(\frac{\kappa+1}{G},\frac{\kappa+1}{G}) > \varepsilon$.

The second inequality is demonstrated by writing out $U(\frac{k-1}{G},\frac{k}{G}) - U(\frac{k+1}{G},\frac{k}{G})$, collecting terms and noting that the coefficient on k is $\frac{8c-(\Pi_F+\Pi_S)}{2G^2} > 0$ (intuitively,

[10]Intuitively, one way to view this is to think of longer delays as being analogous to a game where the grid becomes coarser after one player has entered. This increases returns to defection and therefore shrinks the equilibrium set.

this is demonstrating that the payoff difference between preempting rather than being preempted is increasing in time). The final equality follows from the definition of κ. This establishes the contradiction.

□

We now characterize the ε-equilibrium set for Inertial Continuous time.

Proposition C.4. *Assume that both firms enter at some time $t \leq t^*$.*[11] *In an inertial continuous time game with inertial delay δ, and tolerance ε, all entry times $t \in (0, \bar{\tau}]$ can be sustained in a subgame perfect ε-equilibrium where $\bar{\tau}$ is the solution to:*

$$\arg\max_{\tau \in [\delta, t^*]} \tau \tag{C.6}$$

$$s.t.\ \delta \left[\left(\frac{1-\tau}{2} + 1 \right) \Pi_F - c(2 - 2\tau + \delta) \right] \leq \varepsilon. \tag{C.7}$$

If no such τ exists then the unique equilibrium is immediate entry.[12]

Proof of proposition C.4. The proof technique is the same as for proposition C.3. Again, we begin by assuming immediate response (or as soon as possible given the reaction lag) to an entry, and establish that it can be discarded ex-post.

Consider a strategy such that each firm enters at $t = \tau$ if their opponent has yet to enter at τ, and they enter as soon as possible if their opponent enters before τ. A best response to such a strategy is to to enter at $\tau - \delta$, and enter as soon as possible if their opponent enters before $\tau - \delta$. Notice that entering in range $(\tau - \delta, \tau)$ is not a best response because of the assumption implicit in definition C.1 which tells us that the opponent could still enter at τ in this case.

The payoff from waiting both firms waiting until τ is given by $U(\tau, \tau)$, and the best response strategy of entering at $\tau - \delta$ earns a payoff of $U(\tau - \delta, \tau)$.

[11] See Footnote 9.

[12] In our experimental implementation of inertial time immediate entry typically involved subjects entering at δ, as it took subjects a reaction lag to respond to the start of the period.

Observe that $U(\tau-\delta,\tau)-U(\tau,\tau)=\delta\left[\left(\frac{1-\tau}{2}+1\right)\Pi_F-c(2-2\tau+\delta)\right]$, and that this expression is increasing in τ.

We therefore conclude that the strategy considered above can be sustained as an ε-equilibrium if $\varepsilon\geq\delta\left[\left(\frac{1-\tau}{2}+1\right)\Pi_F-c(2-2\tau+\delta)\right]$.

This establishes the proposition.

As dealt with in the proof of proposition C.3, it is possible to discard the assumption of immediate response. Suppose that there is an ε-equilibrium where there is a delay $\hat{\delta}>\delta$ between entry times and that the first entry time occurs at $t>\bar{\tau}$ where $\bar{\tau}$ is the largest solution to equation C.7. The payoff for the second mover is then given by $U(t+\hat{\delta},t)$. The assumption that this constitutes a subgame perfect ε-equilibrium implies that $U(t-\delta,t)-U(t+\hat{\delta},t)\leq\varepsilon$. But, we also have $U(t-\delta,t)-U(t+\hat{\delta},t)>U(t-\delta,t)-U(t+\delta,t)>U(\bar{\tau}-\delta,\bar{\tau})-U(\bar{\tau}+\delta,\bar{\tau})>U(\bar{\tau}-\delta,\bar{\tau})-U(\bar{\tau}+\delta,\bar{\tau})>U(\bar{\tau}-\delta,\bar{\tau})-U(\bar{\tau},\bar{\tau})=\varepsilon$. Note that the second inequality follows from the derivative $\frac{d(U(t-\delta,t)-U(t+\delta,t))}{dt}=\frac{\delta}{2}\left[8c-(\Pi_F+\Pi_S)\right]>0$, and that the final equality follows from the definition of $\bar{\tau}$. We have reached a contradiction, and reject the existence of such equilibria.

□

Finally, we note that the central motivation for ε-equilibrium – that agents are willing to tolerate small deviations from best response in order to achieve high, cooperative payouts – loses its bite in Perfectly Continuous time where agents can achieve these same payouts without deviating from best response at all. For completeness we state this as a final proposition.

Proposition C.5. *Assume that all agents enter at or before t^*.*[13] *In Perfectly Continuous Time, the set of first entry times that can be supported in Nash equilibrium and the set of first entry times that can be supported in ε-equilibrium are identical for any ε: any entry time $t\in[0,t^*]$ can be supported in either case.*

Proof of proposition C.5. The Nash equilibrium set is given by proposition 3.2. All Nash equilibria are ε-equilibria, so all entry times $t\in[0,t^*]$ can also be supported

[13] See Footnote 9.

in an ε-equilibrium. The proposition rules out entry times after t^* by assumption, so the two sets are identical over the allowable range of entry times.[14] $\square$

C.3.2 ε-equilibrium: Proofs of propositions stated in section 3.2.3

In this subsection we use the preceding characterizations to prove propositions from section 3.2.3.

Proof of proposition 3.4. Proposition 3.4 is an application of proposition C.4.

Substituting t^* into equation C.7, we see that t^* is an equilibrium for any $\varepsilon > 0$ when δ satisfies:

$$\delta\left[\left(\frac{1-t^*}{2}+1\right)\Pi_F - c(2-2t^*+\delta)\right] \leq \varepsilon. \tag{C.8}$$

Equation C.8 bounds a quadratic equation in δ with negative leading coefficient, so that any δ less than the smaller root will satisfy our requirements. Denote this root by $\tilde{\delta}$ and write $b = (\frac{1-t^*}{2}+1)\Pi_F - c(2-2t^*)$ for convenience. The ever-handy quadratic formula delivers:

$$\tilde{\delta} = \frac{b}{2c} - \sqrt{\frac{b^2}{4c^2} - \frac{\varepsilon}{c}}$$

Noting that $\tilde{\delta}$ is always positive whenever $\varepsilon > 0$ completes the proof of the first part of the proposition.

For the second part of the proposition, note that because inertia cannot exceed the total game length $\delta \leq 1$. Substituting $\tau = \delta$ into equation C.7 we find that there will exist no ε-equilibrium with delay if:

$$\delta\left[\frac{3}{2}\Pi_F - 2c + \delta(c - \frac{\Pi_F}{2})\right] > \varepsilon.$$

The left hand side of this equation is increasing in δ over the interval $[0,1]$ and clearly less than ε at $\delta = 0$. We can therefore find a δ that satisfies the inequality

[14] Again, in Perfectly Continuous time, it is possible to support entry times after t^* in an ε-equilibrium. All such equilibria would, however, be ruled out by applying Simon and Stinchcombe's 1989 iterated elimination of weakly dominated strategies argument.

only if the inequality holds when $\delta = 1$. When $\delta = 1$ the inequality simplifies to $\Pi_F - c > \varepsilon$.

$\square$

The upper bound on ε in the second part of the proposition is a natural consequence of the "thick" indifference curves that are associated with large values of ε. For any game with bounded payoffs there will exist a ε large enough that an agent is indifferent between all outcomes, and all outcomes may therefore be sustained in a ε-equilibrium. The restriction on ε can therefore be viewed as a non-triviality requirement.

C.4 Decision Making Under Uncertainty

This appendix contains further information on the decision making under uncertainty rules that are discussed in the main text. In section C.4.1 we give more details on the three decision rules considered in section 3.5.1. Section C.4.2 provides a proposition showing that the asymptotic effect inertia has on cooperation in our game under the MRA heuristic holds for a broader class of dilemma-like games.

C.4.1 Three decision rules

In this section we give more details on the three non-parametric heuristics discussed in Milnor (1954) and examined in 3.5.1 of the body of the paper. Suppose agents play trigger strategies[15] where a trigger time t is the strategy that enters at the earliest of time t and the soonest available entry time after their opponent enters. We shall write $\bar{U}(t,t')$ to identify the payoff associated with agents using trigger strategies with entry times t and t' so that, for example, $\bar{U}(t,t') = U(t,t+\delta)$ in Inertial Continuous Time when $t+\delta < t'$. Notice that the payoffs associated with a particular pair of trigger times will therefore vary with the continuity and inertia of the game.

Our key assumption on beliefs is that agents assume that their opponent will use a trigger strategy with $t \in [0,t^*]$ but have complete uncertainty regarding which trigger strategy within the set will be played (this uncertainty is imposed in Milnor (1954) and (Stoye, 2011a) via a symmetry axiom argued for in Arrow and Hurwicz (1972)).[16] In order to maintain parsimony (particularly important here because we are conducting this analysis *ex post*), we focus on simple applications of three decisions rules without adding additional potentially *ad hoc* structure (for example

[15]For our game, a restriction to admissible strategies implies trigger strategies.

[16]We interpret the term "complete uncertainty" under a "no priors" interpretation (Stoye, 2011a). In the context of our game this implies that the agent believes that any pure strategy in the set $[0,t^*]$ may be used, but the agent has no additional information. Alternatively, it would be possible to use an "exogenous priors" interpretation (Stoye, 2011a) which allows for any probability distribution over the set of states – a multiple priors framework. For our game, this interpretation requires the inclusion of mixed strategy beliefs. It can be demonstrated, however, that the results are robust to introducing beliefs over mixed strategies (e.g. see proposition 3.19 in Halpern and Pass (2012) for the MRA case). As a consequence our predictions do not change under an exogenous priors interpretation of beliefs.

equilibrium structure).[17]

In a first model, which we call Laplacian Expected Utility, agents respond to strategic uncertainty as if they are expected utility maximizers with uniform beliefs over counterpart strategies. Milnor (1954) provides an axiomatization of this decision model, showing that this representation is the result of combining expected utility axioms with the symmetry axiom that underlies the type of extreme strategic uncertainty we are considering here. In our case, the utility of an LEU agent can be written as

$$U_{LEU}(t) = \int_0^{t^*} \bar{U}(t,s)ds,$$

and finding the LEU maximizing entry time is then straightforward.

A second model, Maximin Ambiguity Aversion (or Maxmin Expected Utility) has been axiomatized for endogenous priors by Gilboa and Schmeidler (1989) and for exogenous priors by Stoye (2011b). The model provides a straightforward conservative heuristic for decision making under uncertainty. As with our other decision rules, we assume complete uncertainty over the interval $[0,t^*]$. Adapted to our setting, the maximin expected utility of an agent entering at time t can be written as:

$$U_{MAA}(t) = \min_{t' \in [0,t^*]} \bar{U}(t,t').$$

It is straightforward to see that the argument of the minimization can be taken to be $t' = 0$ for all possible entry times t: having your opponent enter immediately is (weakly) the worst possible thing that can happen for any strategy in every treatment. The best response to $t' = 0$ is $t = 0$ so that the maximum of the minimum

[17] Though we do not impose equilibrium structure on our decision rules, we note that there exist equilibrium concepts that produce predictions that are equivalent to our MRA and MAA decision rules. For ambiguity averse agents, Lo (2009) introduces an equilibrium concept that generates identical predictions to our MAA decision rule. For agents with minimax regret preferences, the Halpern and Pass (2012) equilibrium notion of "iterated regret minimization with prior beliefs" generates the same predictions as our MRA decision rule when prior beliefs are restricted to trigger strategies. While we elect to maintain parsimonious decision rules with exogenous beliefs, it is possible to build equilibrium frameworks (with varying degrees of endogenous belief formation) that produce the same behavioral predictions.

payoffs is given by $U_{MAA}(0) = U(0,0)$.[18]

Finally, the third (and main) model discussed in the paper, Minimax Regret Avoidance, can be traced back to Wald (1950) and Savage (1951) and was axiomatized by Milnor (1954) for the case of discrete states and by Stoye (2011a) for the case of continuos states.[19] Applied to a strategic setting, it is a decision rule whereby agents choose a strategy that minimizes the worst case regret with respect to the range of possible counterpart strategies. More formally, the regret between two strategies, $R(t,t')$, is the difference between the best response payoff and the realized payoff, so that

$$R(t,t') = (\max_{\hat{t} \in [0,t^*]} \bar{U}(\hat{t},t')) - \bar{U}(t,t').$$

The maximal regret associated with a strategy is then defined as

$$R(t) = \max_{t' \in [0,t^*]} R(t,t').$$

Finally, the minimax regret strategy is the strategy that minimizes $R(t)$. Writing $\underline{t}$ for the minimax regret entry time[20] we have

$$\underline{t} = \arg\min_{t \in [0,t^*]} R(t) = \arg\min_{t \in [0,t^*]} \max_{t' \in [0,t^*]} R(t,t'). \tag{C.9}$$

[18] Again, while we focus on a very simple application of maximin to the strategy space, predictions reported in the text are unchanged in equilibrium variations. For example, applying the epistemically founded equilibrium model of Lo (2009) we find that the unique equilibrium entry time is $t = 0$ for all of our treatments (with the exception of PC where it also mirrors the Nash equilibrium), just as in our simpler, non-equilibrium approach.

[19] Interest in minimax regret has increased in recent years with general theories of minimax regret equilibrium provided by Renou and Schlag (2010) and Halpern and Pass (2012), and specific applications to bargaining in Linhart and Radner (1989) and monopoly pricing are found in Bergemann and Schlag (2008) and Bergemann and Schlag (2011).

[20] Notice that there are two places in equation C.9 that we have implicitly restricted attention to pure strategies. The first is in the max operator, where we consider only beliefs over pure strategies. This restriction is without loss of generality because the regret maximizing strategy will always be a pure strategy in our game (see Halpern and Pass (2012), particularly proposition 3.19, for a detailed and general discussion). The second is in the arg min operator, where we require agents to select a pure strategy. This deliberate modeling choice reflects the experimental design, where subjects must implement a pure strategy. While the introduction of mixed strategies can affect the set of regret minimizing strategies, it is not appropriate to allow for mixing in that fashion when modeling experimental behavior with non-expected utility preferences.

We calculate specific MRA predictions in the body of the paper numerically.[21] In section C.4.2, below, we show that for a broad class of games in trigger strategies (including ours) the main effect of inertia observed in these calculations (and in our experiments) must hold under the MRA decision rule.

C.4.2 Minimax regret in a broader class of games

In our experiment, behavior grows more efficient as inertia falls to zero, a result that we show is consistent with subjects using a minimax regret avoidance rule to choose among strategies. In this section we show that there is a general class of dilemma-like inertial time games for which the minimax regret decision rule generates very similar predictions: when inertia is large, minimax regret predicts behavior that is bounded away from the socially optimal outcome and as inertia shrinks towards 0 behavior approaches the socially optimal outcome. Our experimental design can be rewritten as a special case of this result, as can the prisoner's dilemma, Bertrand competition, Cournot competition and public goods games.

We begin by defining the structure of the game and the assumptions on the payoff functions that are necessary for the result to hold, and then define our class of strategies. We restrict attention to trigger strategies, following a suggestion in Halpern and Pass (2012), for both normative and positive reasons. On the positive side, trigger strategies are typically used by subjects in dilemma games, a point emphasized in Friedman and Oprea (2012). On the normative side, allowing for fully general strategies leads to agents using minimax regret decision rules to

[21] Specifically, we define two types of regret: type 1 regret R_1 from entering too early, and the type 2 regret R_2 from entering too late:

$$R_1(t,t') = \begin{cases} R(t,t') & \text{if } t < t' - \delta \\ 0 & \text{if } t' - \delta \leq t \end{cases} \tag{C.10}$$

and

$$R_2(t,t') = \begin{cases} R(t,t') & \text{if } t' - \delta < t \\ 0 & \text{if } t \leq t' - \delta. \end{cases} \tag{C.11}$$

We note that $R_1(t) = R(t,t^*) = U(t^* - \delta, t^*) - U(t,t^*)$ and $R_2(t) = R(t, t-\delta) = U(t - 2\delta, t-\delta) - U(t, t-\delta)$, giving us each type of regret as a function of δ. Then, by noting that $R_1(t)$ is decreasing and $R_2(t)$ is increasing, and noting that $R(t) = \max\{R_1(t), R_2(t)\}$, we can find $R(t)$ by solving the equation $R_1(t) = R_2(t)$. The resultant solution is a function of δ and it is then straightforward to numerically compute minimax regret predictions.

"believe" that their opponent is using strategies that are both non-admissible and involve making large sacrifices with no hope of reciprocation.[22]

Suppose that we have a game being played in inertial continuous time between two players on the interval $[0,1]$. Payoffs are defined via a symmetric flow utility function. The (instantaneous) action space for both players is constant with representative elements $a_i, a_j \in A$, and the flow utility is denoted by $u_i^t(a_i, a_j)$. We identify some key action profiles, which shall be assumed to exist:

- $\arg\max_{a \in A} u^t(a,a) = \hat{a}^t$, with $\hat{a}^t$ unique for all t. We shall call $\hat{a}^t$ the cooperative action at time t. Furthermore, assume that $\hat{a}^t$ is constant so that $\hat{a}^t = \hat{a}$ for all t.[23]

- The instantaneous game at t has a unique Nash Equilibrium that is denoted by a^{*t}. We shall call a^{*t} the Nash equilibrium action at time t. Furthermore, assume that a^{*t} is constant so that $a^{*t} = a^*$ for all t.[24]

- If $u^t(\tilde{a}, \hat{a}) > u^t(\hat{a}, \hat{a})$, then we shall call $\tilde{a}^t$ an exploitative action at time t.[25] We shall assume that there exists at least one exploitative strategy at every t. It will typically be the case that a^* is an exploitative strategy.

We now define a class of generalized trigger strategies. Our trigger strategies collapse to the equivalent of grim trigger in games with only two strategies. We shall require that each trigger strategy uses a fixed, pre-determined and constant exploitation strategy $\tilde{a}$.[26] A trigger strategy with trigger time t_1 and a constant $\tilde{a}$ satisfies the following conditions:

[22] As an example, consider a repeated prisoners dilemma in continuous time. Consider an agent who believes that their opponent will respond to any defection with permanent defection unless the initial defection occurs precisely $\sqrt{e^\pi}$ seconds into the game, in which case the opponent will continue to cooperate. Such an agent will experience extremely large regret unless they defect at precisely $\sqrt{e^\pi}$. Restricting beliefs to trigger strategies removes such strategies from the consideration set of agents when assessing regret.

[23] This final assumption is without loss of generality. For example, if we identify actions with choosing a real number, we can simply relabel the cooperative action to be 1 at every instant.

[24] When $a^{*t} \neq \hat{a}^t$ for all t, then this is also without loss of generality.

[25] Where it will not result in confusion, we will sometimes write $\tilde{a}$ instead of $\tilde{a}^t$.

[26] While this assumption seems quite restrictive it turns out not to affect outcomes very much. The results presented below hold for any pair of exploitation strategies, and the difference in regret minimizing trigger times for differing exploitation strategies is arbitrarily small for small δ.

1. At time 0, play $\hat{a}$.

2. At each time $t < t_1$ if the history is such that both players have always played $\hat{a}$, or if the agent has always played $\hat{a}$ and the opponent has played $\hat{a}$ at all t' such that $t' \leq t - \delta$, then play $\hat{a}$.

3. At each time $t < t_1$ if the agent has played anything other than $\hat{a}$ then play a^* at all remaining moments.

4. At each time $t < t_1$, if the opponent has played anything other than $\hat{a}$ at any time t' such that $t' < t - \delta$ then play a^* at all remaining moments.

5. At time t_1, if the history is such that both players have always played $\hat{a}$ at all $t < t_1$ then play $\tilde{a}$ over the period $t \in [t_1, t_1 + \delta)$ followed by a^* at all $t \geq t_1 + \delta$.

6. At time t_1, if the history is such that the agent has always played $\hat{a}$ at all $t < t_1$ and the opponent played $\hat{a}$ at all $t < t'$ with $t_1 - \delta \leq t'$ and anything other than $\hat{a}$ at t' then play $\tilde{a}$ over the period $t \in [t_1, t' + \delta)$ and play a^* at all $t \geq t' + \delta$.

We can therefore identify each trigger strategy by its trigger time and the exploitative strategy that it implements: a trigger strategy is a pair $(t_1, \tilde{a})$, although we note again that because $\tilde{a}$ is fixed it is not a choice variable, and for this reason we shall often refer to a trigger strategy solely by its trigger time. We will use the expression $U((t_1, \tilde{a}), (t_1', \tilde{a}'))$ or $U(t_1, t_1')$ to denote the payoff induced by the use of trigger strategies. We also interpret the reaction lag, δ, as being a fixed parameter of the game rather than a strategic choice of the agents.

We shall impose the following assumptions on the instantaneous flow payoff functions. While the assumptions may appear to be onerous, we shall demonstrate that several standard applications satisfy them.

- $u^t(\hat{a}, \tilde{a}) \leq u^t(a^*, a^*) < u^t(\hat{a}, \hat{a}) < u^t(\tilde{a}, \hat{a})$ and $u^t(\hat{a}, \tilde{a}) \leq u^t(\tilde{a}, \tilde{a}) < u^t(\hat{a}, \hat{a})$ for all t and all $\tilde{a}$, and that each of these flow payoff functions are continuously differentiable in t. This can be interpreted as the game having a 'dilemma' structure at every instant.

- We also have two conditions on the rate of change of the flow payoffs. $u^{t+t'}(a^*,a^*) < u^t(\hat{a},\hat{a})$ and $u^{t+t'}(\tilde{a},\tilde{a}) \geq u^t(\hat{a},\tilde{a})$ for all t and all $t' \leq \delta$ and all $\tilde{a}$. These rather weak conditions ensure that the 'dilemma' nature of the game persists across every stretch of time that is less than or equal to one reaction lag. A game with constant flow payoffs will trivially satisfy this condition if it satisfies the previous dilemma conditions.

- $\frac{du^t(\tilde{a},\hat{a})}{dt} \geq \frac{du^t(\tilde{a},\tilde{a})}{dt} \geq 0$, $\frac{du^t(a^*,a^*)}{dt} \geq 0$, $\frac{du^t(\hat{a},\tilde{a})}{dt} \leq 0$ and $\frac{du^t(\hat{a},\hat{a})}{dt} \leq 0$ for all t and all $\tilde{a}$. Jointly, these conditions imply that playing the socially optimal strategy becomes (weakly) less attractive as the game progresses.

We now proceed to our main result via a series of lemmas. Proofs are collected in appendix C.4.3. Our first lemma and corollary clarify behaviour in the limiting case as inertia approaches zero.

Lemma C.6. *When $\delta = 0$, trigger strategies are identified only by their trigger time, and the class of trigger strategies with $t = 1$ is weakly dominant among the class of all trigger strategies. Furthermore, a trigger strategy with a trigger time $t < 1$ is not weakly dominant.*

Corollary C.7. *When $\delta = 0$ any trigger strategy with a trigger time of $t = 1$ generates a regret of 0. Any trigger strategy with a trigger time of $t < 1$ generates positive regret.*

Next, we establish some useful properties of the best response function.

Lemma C.8. *The best response trigger strategy t' to a trigger strategy $\underline{t}$ satisfies $\underline{t} - \delta \leq t' \leq \underline{t}$. Furthermore, the best response payoff, $U(t',\underline{t})$ is increasing in $\underline{t}$.*

We shall now reintroduce our two types of regret. Type 1 regret is regret from using a trigger time that is earlier than optimal, and type 2 regret is regret from using a trigger time that is later than optimal. In our current setup, we define the type 1 and type 2 regret between two strategies as:

$$R_1(t,\underline{t}) = \begin{cases} U(t',\underline{t}) - U(t,\underline{t}) & \text{if } t < t' \\ 0 & \text{if } t' \leq t \end{cases} \tag{C.12}$$

and

$$R_2(t,\underline{t}) = \begin{cases} U(t',\underline{t}) - U(t,\underline{t}) & \text{if } t' < t \\ 0 & \text{if } t \leq t'. \end{cases} \tag{C.13}$$

We can also define the type 1 and type 2 regret of a trigger strategy as $R_1(t) = \max_{\underline{t}} R_1(t,\underline{t})$ and $R_2(t) = \max_{\underline{t}} R_2(t,\underline{t})$, noting that typically the arguments of the maxima will be different for $R_1(t)$ and $R_2(t)$. Finally, we can write $R(t) = \max\{R_1(t), R_2(t)\}$. Our next two propositions establish some facts about type 1 and type 2 regret.

Lemma C.9. *Type 1 regret is strictly decreasing in trigger time over the interval $t \in [0, 1-\delta)$ and is equal to 0 in the interval $t \in [1-\delta, 1]$. Furthermore, $R_1(t)$ is continuous.*

Lemma C.10. *Type 2 regret is weakly increasing in trigger time. Furthermore, $R_2(t)$ is continuous and $\lim_{\delta\to 0} R_2(t) = 0$ for all t.*

Finally, we arrive at the proposition that establishes that MRA behaviour supports greater cooperation as δ decreases, and that MRA behaviour converges to fully cooperative play in the limit. The proof is a straightforward application of the intermediate value theorem.

Proposition C.11. *When $\delta > 0$ the earliest entry time that minimizes regret, t_δ, is strictly less than 1. Furthermore, t_δ is decreasing in δ and $\lim_{\delta\to 0} t_\delta = 1$.*

Examples

The following are examples that satisfy the requirements of our setup above.

1. Prisoner's dilemma in inertial continuous time, such as Friedman and Oprea (2012).

2. Oligopoly game with a constant flow rate of customers arriving every instant. To illustrate, in the Bertrand case, suppose that the two firms have zero marginal costs and each instant a mass 1 of customers arrives and demand the quantity $Q = 1 - P$ from the lowest priced firm every instant. In this case $a^* = 0$ and $\hat{a} = \frac{1}{2}$ and take the exploitation strategy to be $\tilde{a} =$

0.49. The associated flow payoffs are $u(\hat{a},\tilde{a}_1) = 0, u(a^*,a^*) = 0, u(\tilde{a},\tilde{a}) = 0.12495, u(\hat{a},\hat{a}) = \frac{1}{8}, u(\tilde{a},\hat{a}) = 0.2499$.

3. Cournot oligopoly game. To illustrate suppose that the two firms have zero marginal costs and each instant a mass 1 of customers arrives and demand the quantity $Q = 1 - P$ from the lowest priced firm every instant. In this case $a^* = \frac{1}{3}, \hat{a} = \frac{1}{4}$ and an example exploitation strategy is $\tilde{a} = \frac{3}{8}$. The associated payoffs are $u(\hat{a},\tilde{a}) = \frac{3}{32}, u(a^*,a^*) = \frac{1}{9}, u(\tilde{a},\tilde{a}) = \frac{3}{32}, u(\hat{a},\hat{a}) = \frac{1}{8}, u(\tilde{a},\hat{a}) = \frac{9}{64}$.

4. Public goods game in continuous time. At each instant each player may contribute $a \in [0,1]$ to the public good. The public good provides instantaneous utility to each player that is equal to three-quarters of the total instantaneous contribution. In this case $a^* = 0$, $\hat{a} = 1$ and an example exploitation strategy is $\tilde{a} = 0$. The associated payoffs are $u(\hat{a},\tilde{a}) = -0.25, u(a^*,a^*) = 0, u(\tilde{a},\tilde{a}) = 0, u(\hat{a},\hat{a}) = 0.5, u(\tilde{a},\hat{a}) = 0.75$.

C.4.3 Proofs

Proof of lemma C.6. When $\delta = 0$, the class of trigger strategies becomes much simpler. The trigger strategy with trigger time t' implements the following strategy. At all $t \geq t'$ play a^*. At all $t < t'$ play $\hat{a}$ if your opponent has played $\hat{a}$ at all $s \leq t$, otherwise play a^*. Because there is no exploitation period trigger strategies can be identified by their trigger time alone.

Compare two strategies with trigger times t_1 and t_2 with $t_1 < t_2$. Denote the opponent's trigger time by t_3. If $t_3 \leq t_1$ then both t_1 and t_2 earn the same payoff for all t. If $t_3 > t_1$ then both t_1 and t_2 earn the same payoff on $[0,t_1]$ and $[\min\{t_3,t_2\},1]$ but on the interval $[t_1,\min\{t_3,t_2\}]$ then t_1 earns flow payoffs $u^t(a^*,a^*)$ while t_2 earns flow payoffs $u^t(\hat{a},\hat{a})$. Therefore t_2 earns a strictly higher payoff against some trigger strategies but the converse does not hold, so that t_2 weakly dominates t_1. □

Proof of corollary C.7. A weakly dominant strategy generates 0 regret, and non-weakly dominant strategies generate positive regret. □

Proof of C.8. In the range $t' > \underline{t} + \delta$ the payoff $U(t',\underline{t})$ is independent of t'.

In the range $\underline{t} < t' \leq \underline{t} + \delta$ we have

$$U(t', \underline{t}) = \int_0^{\underline{t}} u^s(\hat{a}, \hat{a})ds + \int_{\underline{t}}^{t'} u^s(\hat{a}, \underline{\tilde{a}})ds + \int_{t'}^{\underline{t}+\delta} u^s(\tilde{a}, \underline{\tilde{a}})ds + \int_{\underline{t}+\delta}^{1} u^s(a^*, a^*)ds \tag{C.14}$$

which is decreasing in t' because $u^t(\hat{a}, \underline{\tilde{a}}) < u^t(\tilde{a}, \underline{\tilde{a}})$.

When $t' < \underline{t} - \delta$ we have

$$U(t', \underline{t}) \quad = \quad \int_0^{t'} u^s(\hat{a}, \hat{a})ds \; + \; \int_{t'}^{t'+\delta} u^s(\tilde{a}, \hat{a})ds \; + \; \int_{t'+\delta}^{1} u^s(a^*, a^*)ds \tag{C.15}$$

Taking the derivative with respect to t' yields

$$u^{t'}(\hat{a}, \hat{a}) + u^{t'+\delta}(\tilde{a}', \hat{a}) - u^{t'}(\tilde{a}', \hat{a}) - u^{t'}(a^*, a^*)$$

which is strictly positive.

This establishes that, for any $\tilde{a}$ and any $(\underline{t}, \underline{\tilde{a}})$, we must have $\underline{t} - \delta \leq t' \leq \underline{t}$.

We can now, therefore, write down the functional form of the best response payoff:

$$\max_{t'} U(t', \underline{t}) = \max_{t'} \int_0^{t'} u^s(\hat{a}, \hat{a})ds + \int_{t'}^{\underline{t}} u^s(\tilde{a}, \hat{a})ds + \int_{\underline{t}}^{t'+\delta} u^s(\tilde{a}, \underline{\tilde{a}})ds + \int_{t'+\delta}^{1} u^s(a^*, a^*)ds \tag{C.16}$$

Now, applying the envelope theorem, we have

$$\frac{d \max_{t'} U(t', \underline{t})}{d\underline{t}} \quad = \quad u^{\underline{t}}(\tilde{a}, \hat{a}) \quad - \quad u^{\underline{t}}(\tilde{a}, \underline{\tilde{a}}) \quad > \quad 0 \tag{C.17}$$

if the constraints $\underline{t} - \delta \leq t' \leq \underline{t}$ are not binding.

If the constraints are binding we have either

$$\frac{d\max_{t'}U(t',\underline{t})}{d\underline{t}} \;=\; u^{\underline{t}}(\hat{a},\hat{a}) \;-\; u^{\underline{t}}(\tilde{a},\underline{\tilde{a}}) \;+\; u^{\underline{t}+\delta}(\tilde{a},\underline{\tilde{a}}) \;-\; u^{\underline{t}}(a^*,a^*) \;>\; 0 \quad \text{(C.18)}$$

or

$$\frac{d\max_{t'}U(t',\underline{t})}{d\underline{t}} \;=\; u^{\underline{t}-\delta}(\hat{a},\hat{a}) - u^{\underline{t}-\delta}(\tilde{a},\hat{a}) + u^{\underline{t}}(\tilde{a},\hat{a}) - u^{\underline{t}}(a^*,a^*) \;>\; 0 \quad \text{(C.19)}$$

We have therefore established that the best response payoff is increasing in the opponent's trigger time.

□

Proof of lemma C.9. We shall use t' to denote the best response trigger strategy to the trigger $\underline{t}$. The type 1 regret between two strategies can then be written as

$$R_1(t,\underline{t}) = \begin{cases} U(t',\underline{t}) - U(t,\underline{t}) & \text{if } t < t' \\ 0 & \text{if } t' \leq t \end{cases} \qquad \text{(C.20)}$$

The type 1 regret of an individual trigger strategy is defined by the maximum of the regret between two strategies:

$$R_1(t) = \max_{(\underline{t})} R_1(t,\underline{t})$$

Now, we wish to establish that $R_1(t,\underline{t})$ is maximized when $\underline{t} = 1$. When $\underline{t} > t + \delta$ it is clear that $R_1(t,\underline{t})$ is increasing in $\underline{t}$ as the positive part of equation C.20 is increasing in $\underline{t}$ (as established in the proof of the previous lemma), and the negative part of equation C.20 is independent of $\underline{t}$. When $t < t' \leq \underline{t} \leq t + \delta$ we have

$$U(t',\underline{t}) = \int_0^{t'} u^s(\hat{a},\hat{a})ds + \int_{t'}^{\underline{t}} u^s(\tilde{a},\hat{a})ds + \int_{\underline{t}}^{t'+\delta} u^s(\tilde{a},\underline{\tilde{a}})ds + \int_{t'+\delta}^{1} u^s(a^*,a^*)ds$$

and

$$U(t,\underline{t}) = \int_0^{t} u^s(\hat{a},\hat{a})ds + \int_{t}^{\underline{t}} u^s(\tilde{a},\hat{a})ds + \int_{\underline{t}}^{t+\delta} u^s(\tilde{a},\underline{\tilde{a}})ds + \int_{t+\delta}^{1} u^s(a^*,a^*)ds$$

so that

$$U(t',\underline{t})-U(t,\underline{t}) = \int_t^{t'} u^s(\hat{a},\hat{a})ds + \int_t^{t'} u^s(\tilde{a},\hat{a})ds + \int_{t+\delta}^{t'+\delta} u^s(\tilde{a},\underline{\tilde{a}})ds + \int_{t+\delta}^{t'+\delta} u^s(a^*,a^*)ds$$

is independent of $\underline{t}$. Therefore $R_1(t,\underline{t})$ is weakly increasing in $\underline{t}$ so that type 1 regret is maximized when the opponent has a trigger time of 1. We can therefore write the type 1 regret associated with a trigger strategy t as

$$R_1(t) = \left(\max_{t'} U(t',1)\right) - U(t,1)$$

whenever $t < t'$. Because the maximization problem is independent of t, the derivative with respect to t is simply

$$\frac{\partial R_1(t)}{\partial t} = \frac{\partial - U(t,1)}{\partial t} = -u^t(\hat{a},\hat{a}) + u^t(\tilde{a},\hat{a}) - u^{t+\delta}(\tilde{a},\hat{a}) + u^{t+\delta}(a^*,a^*) < 0.$$

We have therefore established that type 1 regret is decreasing in trigger time when $t < t'$.

We know from Lemma C.8 that $U(t',\underline{t})$ must be maximized when $1-\delta \leq t' \leq 1$. It is easily seen that $t' = 1-\delta$ because $u(\hat{a},\hat{a}) < u(\tilde{a},\hat{a})$. Therefore, $t \geq 1-\delta$ implies that $t \geq t'$, so that $R_1(t) = 0$ whenever $t \geq 1-\delta$.

We now demonstrate the continuity of $R_1(t)$. $U(t,\underline{t})$ is continuous because the underlying flow payoff functions are assumed to be continuous, implying that the derivatives of $U(t,\underline{t})$ exist via the Leibniz integral rule which also implies continuity. An application of Berge's maximum theorem then establishes that $\max_{t'} U(t',\underline{t})$ is continuous in $\underline{t}$, and a second application establishes that $\max_{\underline{t}}(\max_{t'}(U(t',\underline{t})) - U(t,\underline{t}))$ is continuous in t.

□

Proof of lemma C.10. Type 2 regret between two strategies can be written as

$$R_2(t,\underline{t}) = \begin{cases} U(t',\underline{t}) - U(t,\underline{t}) & \text{if } t' < t \\ 0 & \text{if } t \leq t' \end{cases} \tag{C.21}$$

We shall address three cases.

Case 1 ($t' = \underline{t} < t$):

We start by establishing that $R_2(t,\underline{t})$ is maximized at $\underline{t} = t - \delta$.

There are two sub cases. First, $\underline{t} < t \leq \underline{t} + \delta$.

In this case, $R_2(t,\underline{t}) = \int_{\underline{t}}^{t} u^s(\tilde{a},\underline{\tilde{a}})ds - \int_{\underline{t}}^{t} u^s(\hat{a},\underline{\tilde{a}})ds$. This is decreasing in $\underline{t}$ because $u^t(\tilde{a},\underline{\tilde{a}}) > u^t(\hat{a},\underline{\tilde{a}})$.

In the second sub case $\underline{t} + \delta \leq t$.

In this case, $R_2(t,\underline{t}) = \int_{\underline{t}}^{\underline{t}+\delta} u^s(\tilde{a},\underline{\tilde{a}}) - u^s(\hat{a},\underline{\tilde{a}})ds$ which is weakly increasing in $\underline{t}$ because $u^t(\tilde{a},\underline{\tilde{a}})$ is weakly increasing in t and $u^t(\hat{a},\underline{\tilde{a}})$ is weakly decreasing in t.

Therefore, the type 2 regret of trigger t is maximized when $\underline{t} = t - \delta$, so that

$$R_2(t) = U(t-\delta, t-\delta) - U(t, t-\delta) \tag{C.22}$$

$$= \int_{t-\delta}^{t} u^s(\tilde{a},\underline{\tilde{a}}) - u^s(\hat{a},\underline{\tilde{a}})ds \tag{C.23}$$

which is weakly increasing in t.

Case 2 ($t' = \underline{t} - \delta$):

We begin with two sub cases. In the first, $\underline{t} - \delta < t \leq \underline{t}$ and $R_2(t,\underline{t})$ is decreasing in $\underline{t}$ because, given the constraints, an increase in $\underline{t}$ pushes the actual entry time (t) closer to the optimal entry time (t') which reduces regret.[27]

In the second sub case, $\underline{t} + \delta \leq t$. In this case, $R_2(t,\underline{t}) = U(\underline{t}-\delta,\underline{t}) - U(t,\underline{t}) = \int_{\underline{t}-\delta}^{\underline{t}} u^s(\tilde{a},\hat{a}) - u^s(\hat{a},\hat{a})ds + \int_{\underline{t}}^{\underline{t}+\delta} u^s(a^*,a^*) - u^s(\hat{a},\underline{\tilde{a}})ds$ which is weakly increasing in $\underline{t}$ because the positive integrands are weakly increasing in time and the negative integrands are weakly decreasing in time.

We have thus established that the $\underline{t}$ that maximizes $R_2(t,\underline{t})$ must satisfy $\underline{t} \leq t \leq \underline{t} + \delta$. In this range we have $\frac{dU(t,\underline{t})}{dt} \leq 0$. It is also the case that $\frac{dU(\underline{t}-\delta,\underline{t})}{dt} \geq 0$. It must, therefore, be the case that

$$R_2(t) = U(\underline{t}-\delta,\underline{t}) - U(t,\underline{t}) \tag{C.24}$$

is weakly increasing in t.

[27] Alternatively, we could write out the payoff functions and note that the conditions on the payoff function that cause $t' = \underline{t} - \delta$ also imply that type 2 regret is decreasing in $\underline{t}$.

Case 3 ($\underline{t} - \delta < t' < \underline{t}$)

In this case we can apply the unconstrained envelope theorem, which allows us to write:

$$\frac{dR_2(t)}{dt} = \frac{\partial - U(t,\underline{t})}{\partial t} \tag{C.25}$$

When $t > \underline{t}$ it is clear that $U(t,\underline{t})$ is decreasing in t: the agent is being pre-empted and benefits from reducing the period in which they are being exploited.

When $t \leq \underline{t}$ the case for $U(t,\underline{t})$ being decreasing is more subtle. Recall that t' is the best response to $\underline{t}$ and that $t' < t$. Furthermore, t' is strictly interior to the interval $\underline{t} - \delta < t' < \underline{t}$. Taken together, these imply that t lies in a region of the U function that is decreasing.[28]

It is therefore true that in the third case we also have $\frac{dR_2(t)}{dt} \geq 0$.

The continuity of R_2 can be established in the same fashion as the continuity of R_1.

From the above arguments it is clear that $\lim_{\delta \to 0} \underline{t} = t$. From Lemma C.8 we see that $\lim_{\delta \to 0} t' = \underline{t}$ so that $\lim_{\delta \to 0} t' = \lim_{\delta \to 0} \underline{t} = t$. It is therefore true that $\lim_{\delta \to 0} U(t',\underline{t}) - U(t,\underline{t}) = U(t,t) - U(t,t) = 0$.

□

Proof of proposition C.11. The proof is simply a matter of stitching together the previous results.

Write $T(t) = R_1(t) - R_2(t)$. $T(0) > 0$ and $T(1-\delta) \leq 0$ and $\frac{dT}{dt} < 0$. By the intermediate value function, there is a unique time, $t_\delta \in [0, 1-\delta]$, such that $T(t_\delta) = 0$. We argue that this is the earliest regret minimizing time. When $t < t_\delta$ we have $R_1(t) > R_2(t)$ so that $R_(t) = R_1(t)$. From Lemma C.9 we know that $R_1(t)$ is decreasing, so that $R_(t)$ is decreasing as well. When $t > t_\delta$ we have $R_2(t) > R_1(t)$ so that $R_(t) = R_2(t)$. From Lemma C.10 we know that $R_2(t)$ is non-increasing, so that $R_(t)$ is non-increasing as well. Therefore t_δ must be a regret minimizing trigger time, and must also be the smallest regret minimizing trigger time.

[28] The first order condition for t' is that $u^{t'}(\tilde{a},\hat{a}) + u^{t'+\delta}(a^*,a^*) = u^{t'}(\hat{a},\hat{a}) + u^{t'+\delta}(\tilde{a},\underline{\tilde{a}})$ and the assumptions on the rate of change of the flow payoff function assure that the second order condition is satisfied globally.

From Lemma C.10, we have that $\lim_{\delta \to 0} R_2(t) = 0$ for all t. Therefore $\lim_{\delta \to 0} T(t) = R_1(t)$ for all t. Furthermore, we have $\lim_{\delta \to 0} t_\delta = \lim_{\delta \to 0} 1 - \delta = 1$ as $R_1(t) > 0$ for all $t < 1 - \delta$ and $R_1(1 - \delta) = 0$.

$\square$

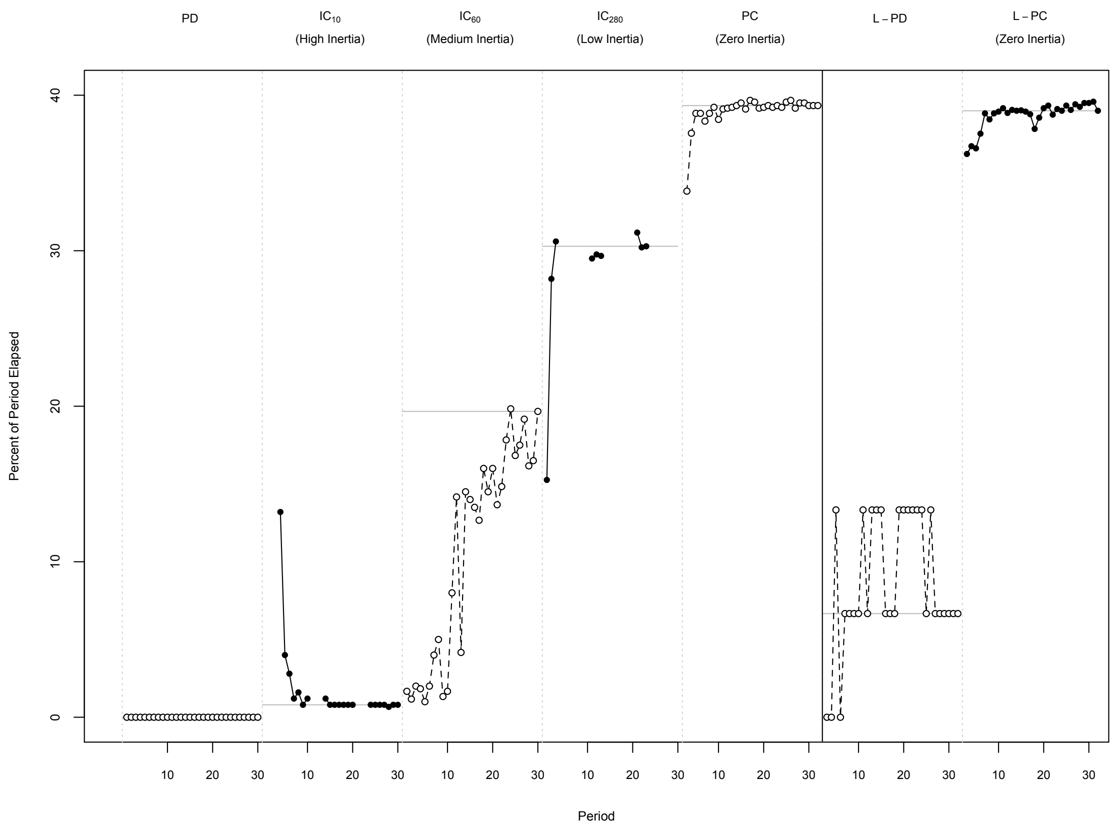

Figure C.5: Median entry time (normalized to the percent of the period elapsed), by period, across all treatments. IC_{60} and PC, PD, $L-PD$ and $L-PC$ treatments ran for 30 periods, while the IC_{10} and IC_{280} treatments were run within-session with three blocks of 3 periods of IC_{280} followed by 7 periods of IC_{10}. Grey horizontal lines for each treatment mark the median final period entry time.

C.5 Time Series of Median Entry Times

Figure C.5 plots period-to-period time series of median raw, observed entry times from all periods in each treatment. Note that unlike the analysis in the main text, these period-to-period results have not been corrected for censoring bias using product limit estimates (since such estimates, at the subject level, require multiple observations per subject and therefore cannot be conducted for each period of the dataset).

Appendix D

Appendix for “Mental Equilibrium and Mixed Strategies for Ambiguity Averse Agents”

D.1 Proofs

D.1.1 Proof of theorem 4.1

Lemma D.1. *Suppose that ν is a belief function and $\Phi = core\ (\nu)$. Then $A \in supp\ \Phi$ if and only if $\nu(A^c) < 1$.*

Proof of lemma D.1. Suppose that $A \in \text{supp}\ \Phi$. Therefore there exists a $\phi \in \Phi$ such that $\phi(A) > 0$ and $\phi(A^c) < 1$. Therefore there exists a $p \in \text{core}\ (\nu)$ such that $p(A^c) < 1$. Therefore $\nu(A^c) < 1$.

Suppose that $A \notin \text{supp}\ \Phi$. Therefore $\phi(A) = 0$ for all $\phi \in \Phi$. Therefore $p(A) = 0$ for all $p \in \text{core}\ (\nu)$. Therefore $p(A^c) = 1$ for all $p \in \text{core}\ (\nu)$. Note that there always exists a $p \in \text{core}\ (\nu)$ such that $p(A^c) = \nu(A^c)$. Therefore $\nu(A^c) = 1$. $\square$

Lemma D.2 (Lo-Nash equilibrium)**.** *Consider a pair $<\sigma,\Phi>$ such that $\Phi_i =$ core ν_i for all i, where ν_i is a belief function. If σ and ν_i satisfy*

$$\sigma(\cdot|a_i) \in core\ \nu_i \ \ \forall a_i \in supp\, \sigma^{A_i}, \forall i \in N \tag{D.1}$$

$$\nu_i(a_{-i}^c) = 1 \Leftrightarrow \prod_{j\neq i} \sigma^{A_j}(a_j) = 0 \ \ \forall a_{-i} \in A_{-i}, \forall i \in N \tag{D.2}$$

and

$$a_i \in \arg\max_{\hat{a}_i \in A_i} \int \hat{a}_i d\nu_i \ \ \forall a_i \in supp\, \sigma^{A_i}, \forall i \in N \tag{D.3}$$

then the pair $<\sigma,\Phi>$ form a Lo-Nash equilibrium. Conversely, if $<\sigma,\Phi>$ are a Lo-Nash equilibrium then σ and ν_i satisfy equations D.1, D.2 and D.3.

Proof of Lemma D.2. We need to demonstrate that $<\sigma,\Phi>$ is a Lo-Nash equilibrium.

Equation D.1 is equivalent to equation 4.5.

The equivalence of equation 4.7 and equation D.3 follows directly from Corollary 4.4 of Gilboa and Schmeidler (1994).

$\prod_{j\neq i} \sigma^{A_j}(a_j) = 0 \Leftrightarrow \exists j \neq i$ s.t. $\sigma^{A_j}(a_j) = 0 \Leftrightarrow a_{-i} \notin \times_{j\neq i}$ supp σ^{A_j} and Lemma D.1 establish that equation 4.6 and equation D.2 are contrapositives.

□

Proof of theorem 4.1. Given the above lemmas, it is sufficient to show an equivalence between equations D.1 and 4.8, equations D.2 and 4.9, and equations D.3 and 4.10. We begin by noting that result 4.1 ensures that each α_i can be associated with a belief function ν_i (and vice-versa) and that Φ_i can be defined as the core of ν_i.

The equivalence of equations D.3 and 4.10 follows from results 4.1 and 4.2.

The equivalence of equations D.1 and 4.8 is a consequence of the definition of the core of a capacity and the fact that $\nu_i(B) = \sum_{T\subseteq B} \alpha_i(T)$ (which follows directly from the definition of α_i).

The equivalence of equations D.2 and 4.9 follows immediately from $\nu_i(a_{-i}^c) = \sum_{T:a_{-i}\notin T} \alpha_i(T)$ which again is an immediate consequence of the definition of α_i.

□

D.1.2 Proof of theorem 4.2

Proof of Theorem 4.2. Suppose that $< \nu_i, \nu_j >$ form an equilibrium. The α_i defined in the theorem is the α_i derived from result 4.1. Given this, Result 4.2 and Definition 4.4 implies that

$$s_i \in \arg\max_{s_i'} \sum_{T \in \Sigma_i} u_i'(s_i', T)\alpha_i(T)$$

for all $s_i \in$ supp σ^{A_i}. Linearity of u' then implies equation 4.13.

To establish equation 4.12, notice that ν_i is additive for all pairs of events that do not form a partition of the state space. Therefore, $\sigma^{A_j}(T) \geq \nu_i(T)$ for all $T \in \Sigma/\Omega$, and $\sigma^{A_j}(\Omega) = \nu_i(\Omega) = 1$. Recalling that $\nu_i(T) = \sum_{\tau \subseteq T} \alpha_i(\tau)$ establishes the result.

For the converse, the linearity of preferences in the mental state space and Equation 4.13 implies that every strategy in the support of σ_i is optimal for agent i with respect to α_i. Then Result 4.2 implies that every strategy in the support of σ_i is also optimal with respect to ν_i.

□

D.1.3 Proof of lemma 4.3

Proof of Lemma 4.3. By construction each Ψ_{x_m} is a partition.

Suppose that Ψ_i is not a filtration. Then there exists a $\Psi_{x_m}^k \in \Psi_{x_m}$ with at least two elements $s, s' \in \Psi_{x_m}^k$ such that $s \in \Psi_{x_{m'}}^{\kappa}$, $s' \in \Psi_{x_{m'}}^{\kappa'}$, $s \notin \Psi_{x_{m'}}^{\kappa'}$ and $s' \notin \Psi_{x_{m'}}^{\kappa}$.

Furthermore we can, without loss of generality, choose m, m' such that $m = m' + 1$. To see this, note that by construction Ψ_{x_0} and Ψ_{x_1} are admissible in a filtration. Suppose that Ψ_{x_m} is the first partition that is not admissible. Then it must be that there is an event in Ψ_{x_m} that is not a subset of any events in a previous partition. Suppose that this previous partition is $\Psi_{x_{m-n}}$ with $n > 1$. But, all events in $\Psi_{x_{m-1}}$ are subsets of events in $\Psi_{x_{m-n}}$, implying that there must also be an event in Ψ_{x_m} that is not a subset of any event in $\Psi_{x_{m-1}}$.

Therefore there exists an information set that precedes x_m through which exactly one of s or s' passes (if s and s' always passed through the same collection of information sets then they would share an event in the$\Psi_{x_{m'}}$ partition) but both s and

s' passes through x_m. This violates perfect recall.

$\square$