Centre for Economic Policy Research

STUDY CENTER GERZENSEE

EUROPEAN SUMMER SYMPOSIUM IN ECONOMIC THEORY

Generously hosted by
Study Center Gerzensee

Monday 1 – Friday 12 July 2013

Rationality and Consistent Beliefs: Theory and Experimental Evidence

Terri Kneeland (University of British Columbia)

The views expressed in this paper are those of the author(s) and not those of the funding organisation(s), nor of CEPR, which takes no institutional policy positions.

Rationality and Consistent Beliefs: Theory and Experimental Evidence

Terri Kneeland

June 21, 2013

Abstract

This is an investigation into the behavioral and experimental support for different epistemic conditions that form the foundations for several solution concepts. It employs strategic *choice* data from a carefully chosen set of ring-network games to obtain individual-level estimates of the following three epistemic conditions: rationality; beliefs about the rationality of others; and consistent beliefs about strategies. We find that 94 percent of subjects are rational, 72 percent of subjects are rational and believe others are rational, and 44 percent of subject are rational and hold at least second-order beliefs about the rationality of others. Not a single subject satisfies all three of the sufficient epistemic conditions for Nash equilibrium. The unique design allows us to weigh the relative plausibility of alternatives to Nash equilibrium used to account for laboratory results. The data tend to support the level-k model.

1 Introduction

Over the past 20 years there has been an accumulation of experimental evidence that challenges the empirical validity of Nash equilibrium predictions when subjects interact in novel environments. This evidence not only prompts the need for alternative solution concepts to explain behavior in one-shot games, but an understanding of *why* Nash equilibrium fails

[*]PhD Candidate, Department of Economics, University of British Columbia. E-mail: terri.kneeland@ubc.ca. I am extremely grateful to my supervisor Yoram Halevy for support and guidance throughout the process of developing and writing this paper. I am also grateful to Mike Peters for numerous conversations, without which this paper would not have been possible. In addition, I would like to thank Colin Camerer, Vince Crawford, David Freeman, Amanda Friedenberg, Ed Green, Li Hao, Andrew Hill, Kevin Leyton-Brown, Wei Li, Ryan Oprea, Ran Spiegler and Charles Sprenger for helpful conversations and comments.

in the laboratory can provide valuable insight into the plausibility of alternative solution concepts.

Solution concepts do not simply provide a description of behavioral predictions but embody assumptions about players' rationality and beliefs that underpin these predictions. Three epistemic conditions form the foundations for several common solution concepts. The first is rationality: whether a player plays a best response to her beliefs. The second is higher-order rationality: whether a player believes others are rational, whether she believes others believe others are rational, and so on. And, the third is consistent beliefs about strategies:

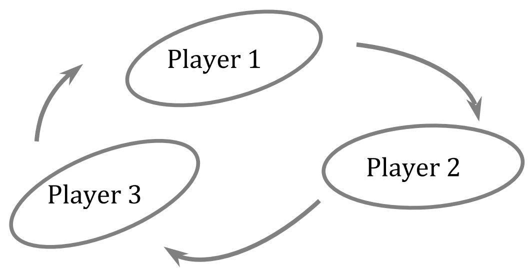

Figure 1: 3-player ring game

The end goal of this paper is to use choice data to make inferences about players' epistemic conditions (i.e. their belief structures). This is an inherently difficult problem as there are generally many different beliefs that lead to the same behavior. This paper focuses on ring games as they allow us to isolate the behavioral implications of different belief structures. A ring game is essentially a series of two-player normal form games. However, the opponent structure is relaxed relative to standard game forms. In the 3-player ring game in Figure 1, Player 1's payoff depends on the action of Player 2. Player 2's payoff depends on the action of Player 3. And, Player 3's payoff depends on the action of Player 1.

The main identification problem comes in identifying a players order of rationality: the behavioral implications of higher orders of rationality are always contained in the behavioral implications of lower orders of rationality. We solve this problem by assuming that a lower order rational player does not respond to changes in higher-order beliefs and by focusing on a carefully chosen set of ring games that exploit this identifying assumption.

The advantage of ring games lies in the relaxed opponent structure. The structure allows one to define games which induce differences only in higher order beliefs (without affecting lower order beliefs). This is not possible in standard game forms (e.g. bimatrix games) because there is a tight link between higher and lower order beliefs. A player's own payoffs affect her 2nd,4th,6th-order beliefs while her opponent's payoffs affect her 1st,3rd,5th-order beliefs and so on. Thus, it is not possible to change payoffs without affecting lower-order beliefs. But, adding a new player into the game creates an additional degree of freedom. In the 3-player ring game, we can change player 3's payoffs without affecting player 1's own payoffs or her 1st-order beliefs. Thus, ring games can be used to isolate the behavioral implications of different orders of rationality because the structure of ring games allows us to isolate beliefs at different orders.

Identification of consistent beliefs is also a problem, as on its own, consistent beliefs has no behavioral implications. To solve this problem, we use a characterization of Nash equilibrium

games) see Jackson (2005) and Kearns (2007)

that requires both rationality and consistent beliefs. Given that a player satisfies the sufficient rationality conditions, identified from behavior in ring games, playing a Nash equilibrium is indication of whether or not a player has consistent beliefs. The characterization used in this paper follows Aumann and Brandenburger (1995) and Perea (2007). We follow the structure of Perea, characterizing behavior in terms of non-interactive belief conditions, which allows us to characterize behavior in terms of one player's beliefs (rather than conditions on multiple players simultaneously as in Aumann and Brandenburger (1995)) which is naturally suited to the application of individual decision data. In addition, we relax the conditions of Perea to make them better suited to identification.

The methodology here differs substantially from previous experimental work that tests between solution concepts. Existing game theoretic work compares different solution concepts by jointly testing relaxations of the three epistemic assumptions in combination with additional structural assumptions that these models impose. Tests of Nash equilibrium, rationalizability and level-k models all allow players to make mistakes according to some error distribution post-hoc. Rationalizability generally requires assumptions about what rationalizable outcome is played when the set is not unique (it is typically imposed that an action is chosen randomly). Level-k models always impose a specification for the behavior of level-0 types. This specification then anchors the beliefs of all other levels by having players play an action that is a finitely iterated best response to the level-0 behavior. Similarly, QRE imposes the structural assumption that players' mistakes are determined according to a logistic error structure. Players then noisily best respond to other players' noisy best responses. These untested structural assumptions are generally difficult to separate from epistemic failures themselves.[5]

The approach in this paper differentiates between alternative solution concepts by testing epistemic conditions directly. Untangling the role that rationality and consistent beliefs play in strategic reasoning will allow direct comparison of alternative solution concepts without relying on structural assumptions that risk introducing bias.[6]

This paper is also closely related to Healy (2011) and Costa-Gomes and Weizsäcker (2008). Both Healy's work and this paper test the epistemic conditions of rationality and consistent beliefs in the lab. The main differences between these papers is the methodology. Healy does an aggregate level analysis and estimates epistemic conditions by eliciting sub-

[5]Wright and Leyton-Brown (2010) provide the most exhaustive test to date between Nash equilibrium and a set level-k and QRE models by performing a meta-analysis of existing normal-form game experimental data using standard structural methodology. Level-k and QRE perform equally well in these tests.

[6]Of note, the estimated order of rationality distribution in this paper is interpretable as an estimate of the level-k level distribution that is independent of the level-0 specification. The only other paper to do provide a level-0 independent distribution is Burchardi and Penczynski (2010). They incentivize communication between group members and analyze the communication in order to determine a subject's depth of reasoning.

jects' beliefs about actions, payoffs, and rationality and eliciting subjects' beliefs about their opponents' beliefs about actions and payoffs. He finds that Nash equilibrium fails because players are often wrong in their beliefs about what other players are doing and that at times subjects are not playing the games we think they are.[7] Costa-Gomes and Weizsäcker elicit beliefs and choices in normal form games. They find that players tend not to best respond to their own stated beliefs, however they also find that players behave differently when asked to state beliefs. Whether or not belief elicitation affects strategic behavior remains an open question. See Healy (2011) for a more complete discussion of these issues. Thus, this paper takes a different approach, obtains individual level estimates, and uses a design specifically chosen to allow us to identify epistemic conditions from choice data. In addition, we are able to get an estimate of a player's order of rationality, which would be difficult under belief elicitation as we would need to elicit beliefs up to the 3rd order.[8]

This paper is also related to the literature on iterated dominance (Beard and Beil 1994; Andrew et al. 1994; Huyck et al. 2002; Ho et al. 1998; and Costa-Gomes et al. 2001), as we estimate a player's order of rationality based on her choices in dominance solvable games. The existing literature tends to focus on violations of iterated dominance, with the exception of Ho et al. (1998) which measures a subject's capability to perform levels of iterated dominance as her level-k level. The current paper measures subjects order of rationality by classifying subjects into levels of iterated dominance using a more general identifying assumption than the level-k structure.

This paper uses strategic choice data to obtain individual-level estimates of the three epistemic conditions. We find that 94 percent of subjects behave consistently with rationality: they do not choose strictly dominated actions. 72 percent of subjects behave consistently with 2nd-order rationality (rational and believe others are rational). 44 percent of subjects behave consistently with 3rd-order rationality (2nd-order rational and hold 2nd-order beliefs that others are rational). And, 20 percent of subjects behave consistently with 4th-order rationality (3rd-order rationality and hold 3rd-order beliefs that others are rational). The epistemic conditions sufficient for Nash equilibrium in two-player games are: rationality,

[7] This means that the assumption of common knowledge of payoffs does not hold. However, Healy finds this is true only for particular games, like the Prisoner's Dilemma. He finds support for the payoff assumptions imposed in this paper in games where there seems to be no obvious role for other-regarding preferences, particularly in dominance solvable games.

[8] In addition, while it is possible to test rationality by testing whether a subject's choices are a best response to her stated 1st-order beliefs about actions, it is not possible to test higher-order rationality in this matter (i.e. failure of a subject's stated 1st-order beliefs about actions to be a best response to her stated 2nd-order beliefs about actions does not indicate a failure of 2nd-order rationality). To avoid this problem, Healy directly estimated 2nd-order rationality by eliciting subjects beliefs about the rationality of their opponent. This has the added difficulty of requiring players to understand the concept of rationality directly.

belief in others' rationality and consistent beliefs about strategies. This means that approximately 28 percent of subjects do not satisfy the sufficient rationality conditions for Nash equilibrium and approximately 72 percent satisfy the sufficient rationality conditions. However, of those 72 percent of subjects that satisfy the sufficient rationality conditions for Nash equilibrium, none of them satisfy consistent beliefs. Not a single subject satisfies all three sufficient epistemic requirements for Nash equilibrium.

The results of this experiment provide support for the key features of the level-k and cognitive hierarchy models; heterogeneous levels of rationality and inconsistent beliefs are important features of strategic reasoning at the individual level.

This paper proceeds as follows. The next subsection discusses related literature. Section 2 introduces the ring game and motivates the experimental design through an example. Section 3 formalizes the epistemic concepts of rationality and consistent beliefs and the identifying assumptions used to separately identify the three epistemic conditions. Section 4 discusses the experimental design. Section 5 gives the experimental results. And, Section 6 provides a discussion of the solution concepts of rationalizability, quantal response equilibrium, and level-k and cognitive hierarchy models in relation to the results. Omitted proofs can be found in Appendix A.

2 Example

The main goal of this paper is to separate to which extent the three assumptions: (1) rationality, (2) beliefs about rationality, and (3) consistent beliefs contribute to strategic reasoning. Are people rational? To what order do people believe others are rational? And, do people have consistent beliefs? In this section, I illustrate the challenge in separately estimating these three features of reasoning and motivate the ring game as a solution to this problem.

Consider the problem of estimating a subject's order of rationality. Bernheim (1984), Pearce (1984), and Tan and Werlang (1988) establish the strategies that can be played by a subject who is kth-order rational:[9] in any complete information game, a subject who is kth-order rational must play a kth-order rationalizable action. In many games (and in all the games considered in this paper) a strategy is kth-order rationalizable if and only if it survives k rounds of iterated deletion of strictly dominated strategies.

Consider the game B1 in Figure 2 (the first matrix represents player 1's payoffs and the

[9] kth-order rationality corresponds to holding finite-orders of beliefs about the rationality of others. For example, 1st-order rationality means you are rational, 2nd-order rationality means you are rational and you believe your opponent is rational. 3rd-order rationality means you are rational and believe that your opponent believes her opponent is rational, and, so on. This is formalized in Section 3.

Player 1 (rows: Player 1's actions; columns: Player 2's actions)

	a	b
a	15	5
b	5	10

Player 2 (rows: Player 2's actions; columns: Player 1's actions)

	a	b
a	10	5
b	5	0

Figure 2: B1: rationalizable bimatrix game

second matrix represents player 2's payoffs. This game is 2nd-order rationalizable for player 1 and 1st-order rationalizable for player 2. If player 2 is rational, she must play action a. If player 1 is rational she can play either a or b. If she satisfies 2nd-order rationality (is rational and believes player 2 is rational) then she must play action a.

Suppose we observe a subject play the action a as player 1. She may have played a because she is 2nd-order rational. However, we cannot rule out the possibility that the subject only satisfies lower orders of rationality. This is because the sets of rationalizable actions necessarily contain one another. The 1st-order rationalizable set contains the 2nd-order rationalizable set which contains the 3rd-order rationalizable set and so on. This leads to an identification problem.

Player 1 (rows: Player 1's actions; columns: Player 2's actions)

	a	b
a	15	5
b	5	10

Player 2 (rows: Player 2's actions; columns: Player 1's actions)

	a	b
a	5	0
b	10	5

Figure 3: B2: rationalizable bimatrix game

This problem can be resolved by observing behavior in related sets of games. Consider the game B2 in Figure 3. This game is 2nd-order rationalizable for player 1 and 1st-order rationalizable for player 2. If player 2 is rational, she must play action b. If player 1 is 2nd-order rational then she must play action b. In addition, B2 is related to B1 in a structured way: player 1 has the same payoffs in both B1 and B2. In other words, player 1 has equivalent 1st-order payoffs in games B1 and B2 (but has different higher-order payoffs).

This structure along with an identifying assumption allows us to solve our identification problem. *We assume that lower-order rational subjects do not respond to changes in higher-order payoffs.* For example, if a rational subject does not respond to changes in 2nd-order

payoffs, then she should play the same action in both games.[10] Under this assumption, a rational subject would play either (a, a) or (b, b) as player 1 in games B1 and B2 respectively, however a subject who is 2nd-order rational (or higher) would play action profile (a, b). Observing the action profile of a subject in both games B1 and B2 would then allow us to separate the behavioral implications of 2nd-order rationality from 1st-order rationality.

Following this logic, the behavioral implications of kth-order rationality can be separated from lower-orders of rationality by looking at games that differ only in kth-order payoffs (and higher) but have different kth-order rationalizable implications. However, no such bimatrix games exist. This is because there is a tight link between higher- and lower-order payoffs in bimatrix games. A player's payoffs determine her 1st-, 3rd-, 5th-order payoffs, and so on. While her opponent's payoffs determine her 2nd-, 4th-, 6th-order payoffs, and so on. Higher-order payoffs cannot change independently of lower order payoffs. Therefore, if any two bimatrix games have the same payoffs up to at least the 2nd-order they must be the same game (and hence have the same rationalizable implications).

This paper solves this problem by making use of a novel class of games: ring games. A ring game is essentially a series of two-player games that have a unique opponent structure. In a standard 2-player game, player 1 and player 2 are each other's mutual opponent. But, in a ring game, player 1's opponent is player 2 but player 2 has an entirely different opponent, player 3. The opponent structure of ring games allows us to solve the identification problem because it allows us to induce changes in higher-order payoffs independently of lower-order payoffs.

Consider the game R1 in Figure 4. Player 1's payoff depends upon her own action and the action of player 2. Player 2's payoff depends upon her own action and the action of player 3. And, player 3's payoff depends upon her action and the action of player 1.

R1 is 1st-order rationalizable for player 3, 2nd-order rationalizable for player 2, and 3rd-order rationalizable for player 1. Thus, player 3 must play a if she is rational, player 2 must play a if she is 2nd-order rational and player 1 must play a if she is 3rd-order rational (rational, believes her opponent is rational and believes her opponent believes her

[10] This identification assumption seems empirically valid. In our experimental data, subjects follow the restrictions of this assumption 83 percent of the time. In addition, the assumption is supported by the limited depth of reasoning literature. If a subject has a finite depth of reasoning k, then she is rational and satisfies (k-1)th-order belief in rationality and will ignore any information that has to be processed at k+1 depths of reasoning. For example, if a subject is rational (and not higher-order rational) because she has a depth of reasoning of 1 then she will not base her decision on any payoffs besides her own. There is empirical support for this type of behavior from experiments that analyze the information search patterns of subjects. Costa-Gomes et al. (2001), Costa-Gomes and Crawford (2006), Wang et al. (2009), Brocas et al. (2009), Camerer et al. (2002), and Johnson et al. (1993) all analyze strategic behavior by investigating the information search pattern of subjects. They find a correlation between the play of kth-order rationalizable strategies and patterns of search that are associated with k depths of reasoning.

Player 1

Player 1's actions \ Player 2's actions	a	b
a	10	5
b	5	15

Player 2

Player 2's actions \ Player 3's actions	a	b
a	10	5
b	5	15

Player 3

Player 3's actions \ Player 1's actions	a	b
a	10	5
b	5	0

Figure 4: R1: rationalizable ring game

opponent is rational). Game R2, in Figure 5, is 1st-order rationalizable for player 3, 2nd-order rationalizable for player 2, and 3rd-order rationalizable for player 1. Player 3 must play b if she is rational, player 2 must play b if she is 2nd-order rational and player 1 must play b is she is 3rd-order rational.

Player 1

Player 1's actions \ Player 2's actions	a	b
a	10	5
b	5	15

Player 2

Player 2's actions \ Player 3's actions	a	b
a	10	5
b	5	15

Player 3

Player 3's actions \ Player 1's actions	a	b
a	5	0
b	10	5

used to estimate whether or not a player has consistent beliefs. On its own, consistent beliefs has no behavioral implications. Thus, we use a characterization of Nash equilibrium, based on Aumann and Brandenburger (1995) and Perea (2007),[11] that combine assumptions about consistent beliefs and rationality to give us testable behavioral implications. In any 2-player complete information game if a player satisfies 2nd-order rationality and consistent beliefs, then she has to play an action consistent with a Nash equilibrium. Thus to estimate whether consistent beliefs hold, we will consider a bimatrix game that has a unique Nash equilibrium but where all actions are rationalizable. For all subjects that satisfy 2nd-order rationality, if they play the Nash equilibrium we will assign them as having consistent beliefs. For any subject that satisfies 2nd-order rationality, failure to play a Nash equilibrium is an indication of a failure of consistent beliefs.

In the following section, we define what we mean by rationality, higher-orders of rationality and consistent beliefs and formalize our identification assumptions.

3 The model

In this section, we apply the tools of epistemic game theory in order to model a player's beliefs about the strategies and payoﬀs of others. This structure then allows us to formally define rationality and player's beliefs about each others' rationality and consistent beliefs. We then formalize the assumptions needed to identify orders of rationality and the characterization of Nash equilibrium which will allow identification of consistent beliefs.

Type spaces are regularly used to model incomplete information games. In such a type space, a player's type represents her beliefs about the payoﬀ types of others. An epistemic type space is analogous, except a player's types represent her beliefs about the *payoﬀs and strategies* of others.

Our experimental design involves both 2-player bimatrix games and n-player ring games. Thus, we formally consider only n-player ring games which contain the case of the 2-player bimatrix game when n=2.

[11] This characterization is based Aumann and Brandenburger (1995) and Perea (2007). We rewrite Aumann and Brandenburger's characterization, which is based on mutual knowledge conditions, using the epistemic type space framework in order to characterize Nash behavior in terms of non-interactive conditions that place restrictions on only one player's belief hierarchy. The non-interactive framework is better suited to using individual decision date. We use a similar framework as Perea, however we provide a tighter characterization that is inline with Aumann and Brandenburger's characterization. We also relax the characterization of a belief with probability 1 to only belief with probability p. This relaxation is necessary for testing epistemic conditions in finite games, as the behavioral implications of a belief with probability 1 are always equivalent to belief with probability p as long as p is sufficiently close to 1. The relaxation of belief with probability 1 is discussed in Appendix B.

Definition. An **n-player ring game** Γ is a tuple $\Gamma = \langle I = \{1, \ldots, n\}; S_1, \ldots, S_n; \pi_1, \ldots, \pi_n; \Theta_1, \ldots, \Theta_n; p\rangle$ where I is a finite set of players, S_i is a finite set of actions for player i, Θ_i is a set of payoff types for each player, $\pi_i : S_i \times S_{p(i)} \times \Theta_i \to \mathbb{R}$ represents the payoff function for each player and $p : I \to I$ is function representing the opponent of player i with the restriction that $p(i) = 1 + i \mathrm{mod} n$.

Given a game, we can define an epistemic type space which describes the beliefs of each player about the strategies and payoffs of the other players.

Definition. Let $\Gamma = \langle I = \{1, \ldots, n\}; S_1, \ldots, S_n; \pi_1, \ldots, \pi_n; \Theta_1, \ldots, \Theta_n; p\rangle$ be an n-player game. A finite Γ-based **epistemic type space** is a set $\langle T_1, \ldots, T_n; b^1, \ldots, b^n; \hat{s}_1, \ldots, \hat{s_n}; \hat{\theta}_1, \ldots, \hat{\theta}_n\rangle$ where T_i is a finite set, $b^i : T_i \to \Delta(T_{-i})$ is an injective function, $\hat{s}_i : T_i \to \Delta(S_i)$, and $\hat{\theta}_i : T_i \to \Theta_i$.

Every player has a set of types T_i. The function $\hat{s}_i$ defines a strategy for each type. The function $\hat{\theta}_i$ defines a payoff type for each type. The function b^i represents each type's beliefs about the types of her opponents, T_{-i}. Therefore $b^i(t_i)$ together with the function $\hat{s}$ defines type t_i's beliefs about the strategies of her opponents and $b^i(t_i)$ together with the function $\hat{\theta}$ defines type t_i's beliefs about the payoff types of her opponents.

3.1 Epistemic conditions

We can define the expected utility of a type t_i by[12]

$$u_i(s_i, t_i) \equiv \sum_{t_{p(i)} \in T_{p(i)}} b^i(t_i)(t_{p(i)})\pi(s_i, \hat{s}(t_{p(i)})).$$

Given this specification for utility, a type is rational if the action $\hat{s}_i(t_i)$ is a best response for player t_i given her beliefs.

Definition. A type t_i is **rational** if $\hat{s}_i(t_i)$ maximizes player i's expected payoff under the measure $b^i(t_i)$. That is, if

$$u_i(\hat{s}_i(t_i), t_i) \geq u_i(s', t_i)$$

for all $s' \in S_i$.

We are interested in whether types are rational, but also if they believe others are rational and so on. First, we say a player believes an event if she places probability 1 on that event happening.

[12] Let $\pi(s_i, \hat{s}(t_{p(i)}))$ be redefined in the standard way whenever $\hat{s}_i$ is a mixed-strategy.

Definition. A type t_i **believes** an event $E \subseteq T_{p(i)}$ if $b^i(t_i)(E) = 1$. Let the set $B^i(E) = \{t_i \in T_i | b^i(t_i)(E) = 1\}$ be the set of types for player i that believe event E.

Higher-order beliefs about the rationality of others can then be defined in recursively in the following way. Define

$$
\begin{aligned}
R_i^1 &= \{t_i \in T_i | t_i \text{ is rational}\} \\
R_i^{m+1} &= R_i^m \cap B^i(R_{p(i)}^m)
\end{aligned}
$$

Definition. If $t_i \in R_i^m$ then we say that t_i satisfies **mth-order rationality.**

We apply the assumption of consistent beliefs only to 2-player games, and hence we define the concept only for the case when n=2. The assumption of consistent beliefs places restrictions on player's beliefs about the strategies of others. Consistent beliefs ensure that player i believes her opponent believes she is playing the action she is actually playing.

Consistent beliefs are the condition that whenever a type t_i believes an event, then she believes her opponent believes she believes it. In other words, she believes her opponent is correct about her beliefs.

Definition. A type t_i has **consistent beliefs** if for any $t_{-i} \in T_{-i}$ such that $b^i(t_i)(t_{-i}) > 0$, then $b^{-i}(t_{-i})(t_i) = 1$.

3.2 Identification

The concepts of rationality, higher-order of rationality and consistent beliefs are all conditions on only one player's belief hierarchy. This means that the epistemic conditions can be stated separately for each player (rather than across players in terms of mutual or common knowledge assumptions). Characterizing epistemic assumptions in this way is naturally better suited for testing individual decision data. In this section I discuss the identification strategy used in this experiment. I discuss the identifying assumptions used to identify a subject's order of rationality in a set of ring games. I then discuss the assumptions and theoretical result used to identify whether a subject has consistent beliefs.

Formally we make three identifying assumptions. The first two assumptions place restrictions on the type space and the third is the behavioral assumption discussed in Section 2. The first assumption restricts the game to a game of complete information - each player has a single payoff type determined by the payoffs of the game. This assumption imposes that differences in behavior will come from differences in beliefs about rationality and consistency rather than differences in payoff types. Further, it allows us to assume that exogenous

changes in payoffs cause exogenous changes in players beliefs. It is a standard assumption in experimental economics - preferences are induced by the game payoffs. No predictions would be possible without some restrictions on the relationship between the game payoffs and player's preferences. We assume the simplest restriction in this paper.

A1: $\Theta_i = \{\theta_i\}$

The second assumption places a restriction on first-order beliefs. Players form first-order beliefs only about their direct opponents. This assumption only makes sense in ring games (which we focus our attention on), this is because in a ring-game, a player's payoff depends only on the action of her direct opponent. For example, in a 3-player ring game, player 1 cares what player 3 does only to the extent that she cares what player 2 believes player 3 does. In other words, player 1's beliefs about what player 3 does has no influence on her expected payoffs.

A2: $b^i : T_i \to \Delta(T_{p(i)})$ for all $i \in I$

The behavioral assumption discussed in the previous chapter assumes that a player with a finite order of rationality will not respond to changes in higher order beliefs. Under assumption A2, changing the payoffs of a player that is n links away from player i will only affect player i's nth-order beliefs and not her lower order beliefs. Thus A1 and A2, build the epistemic foundation to apply our behavioral assumption A3 (that requires that higher-order beliefs can be changed independent of lower-order beliefs in ring games).

The last assumption is a behavioral assumption that allows us to separate the behavioral implications of different orders of rationality and consistent beliefs. Essentially, this is the assumption that players with finite-orders of rationality do not respond to changes in higher-order beliefs. In order to define this assumption we must define higher-order beliefs about payoffs.

Every game can be characterized by its payoff hierarchy for each player. Under assumption A1 and A2, we can define the payoff hierarchy $\bar{h}^i(\Gamma)$ of player i for any complete information game Γ:

$$\bar{h}^i(\Gamma) = \{\theta_{p^k(i)}\}_{k=0}^{\infty}.$$

The payoffs of player i are determined by $\bar{h}^i_0(\Gamma)$, her 1st-order beliefs about payoffs are determined by $\bar{h}^i_1(\Gamma)$, her 2nd-order beliefs about payoffs are determined by $\bar{h}^i_2(\Gamma)$ and so on. Payoff hierarchies of complete information games are relatively redundant. For a 2-player complete information ring game (i.e. a bimatrix game), $\bar{h}^1(\Gamma) = \{\theta_1, \theta_2, \theta_1, \theta_2, \ldots\}$. However, increasing the number of player's in the ring game permits us to provide a richer set

of payoff hierarchies. For example, in a 3-player ring game, $\bar{h}^1($

set of players for which I can identify consistent beliefs. In addition, Perea has two aspects of consistent beliefs that operate (believing you are correct about your opponent's beliefs and that she is correct about your beliefs), we show that the first part is sufficient, and hence are able to show that if consistent beliefs fail it is the first part that fails.

In Appendix B, I relax the notion of belief from that of belief with probability 1 to belief with probability p. I show that with a sufficiently high p, Proposition 1 still characterizes Nash equilibrium behavior in 2-player games. This allows the identification of consistent beliefs even if we estimate each player to only believe her opponent is rational with probability p, which is the case in our estimates. The robustness of our identification to a generalization of belief with probability 1 to belief with probability p is discussed further in Appendix B.

4 Experimental design

This section discusses the details of the experimental design that allows us to separately identify the three key epistemic conditions: rationality, beliefs about rationality and consistent beliefs. The design is motivated by the results of the previous sections.

4.1 Rationality

Each subject's order of rationality is estimated from strategic choice data in 2 4-round rationalizable 4-player ring games. This separates subjects into five rationality categories: irrational (R0), rational (R1: 1st-order rational), rational and 1st-order belief in rationality (R2: 2nd-order rational), rational and 2nd-order belief in rationality (R3: 3rd-order rational), and rational and 3rd-order belief in rationality (R4: 4th-order rational).

Player 1

Player 2's actions (columns); Player 1's actions (rows)

	a	b	c
a	8	20	12
b	0	8	16
c	18	12	6

Player 2

Player 3's actions (columns); Player 2's actions (rows)

	a	b	c
a	14	18	4
b	20	8	14
c	0	16	18

Player 3

Player 4's actions (columns); Player 3's actions (rows)

	a	b	c
a	20	14	8
b	16	2	18
c	0	16	16

Player 4

Player 1's actions (columns); Player 4's actions (rows)

	a	b	c
a	12	16	14
b	8	12	10
c	6	10	8

Figure 6: G1: 4-player dominance solvable ring game

Consider the games G1 and G2 in Figures 6 and 7. The games are similar to the games R1 and R2 from Section 2 except that G1 and G2 are 4-player, 3 action games. Adding an

Player 1

Player 2's actions

Player 1's actions	a	b	c
a	8	20	12
b	0	8	16
c	18	12	6

Player 2

Player 3's actions

Player 2's actions	a	b	c
a	14	18	4
b	20	8	14
c	0	16	18

Player 3

Player 4's actions

Player 3's actions	a	b	c
a	20	14	8
b	16	2	18
c	0	16	16

Player 4

Player 1's actions

Player 4's actions	a	b	c
a	8	12	10
b	6	10	8
c	12	16	14

has a predicted action (or set of actions) for these 8 games, $a_k \quad S_k \quad \{a, b, c\}^8$. For example, type R4 has a unique prediction. She must play the 4th-order rationalizable strategy in each of the 8 games, $S_4 = \{(a, b, a, a, c, a, b, c)\}$. Type R3 has a unique prediction for 6 of the 8 games, but can play a number of different strategies for the other 2 games, $S_3 = \{(a, b, a, a, a, a, b, c), (c, b, a, a, c, a, b, c), (b, b, a, a, b, a, b, c)\}$. Notice that A1-A3 generate an exclusion restriction and ensure that $S_4 \cap S_3 =$. We can define similar sets for types R1 and R2 such that $S_n \cap S_m =$ for any $n = m \quad \{1, 2, 3, 4\}$.

If a subject's action profile matches one of the predicted action profiles of type Rk, then we would assign them as that type. However, a subject's action profile does not always coincide precisely with a predicted action profile so we would like to allow for the possibility that a player might make an error. We allow subjects to make at most 1 error. If they are within 1 error to an exact type match of type Rk (i.e. the action played in one of the 8 games can be changed so that the action profile matches a predicted action profile of type Rk) then we assign that subject as that type. If a subject makes 2 or more errors, then we assign them as an irrational (R0) type. If a subject's action profile is within 1 error of 2 types, we assign them to the type where the error is due to not following assumption A3 rather than rationality (in practice this means that if an action profile was 1 error away from two different types, they were assigned the lowest type).

Assignment	**# of subjects assigned**	**# of action profiles**

Table 2: Subjects assigned by category

the likelihood that the subject's action profile is generated by the given type. This approach closely follows Costa-Gomes and Crawford (2006). Details are discussed in Appendix C.

4.2 Consistent beliefs

Next, we test whether a player has consistent beliefs about strategies based on strategic choices in game Game G3 in Figure 8. G3 is a 2-player ring game. There is a unique Nash equilibrium (a, a) in game G3.

Player 1

Player 2's actions (columns); Player 1's actions (rows)

	a	b	c
a	4	10	9
b	2	12	2
c	2	6	18

Player 2

Player 1's actions (columns); Player 2's actions (rows)

	a	b	c
a	6	0	8
b	2	10	18
c	2	12	16

Figure 8: G3: 2-player ring game

We apply the characterization of Nash equilibrium to estimate consistent beliefs. If a subject satisfies rationality, 1st-order belief in rationality, and consistent beliefs then she must play the Nash strategy a in game G3 (as both player's 1 and 2). If a subject satisfies rationality and 1st-order belief in rationality as determined from games G1 and G2 (all R2, R3 and R4 subjects), and fails to play to play a in G3, then it must be because the assumption of consistent beliefs fails.

The action a is also a rationalizable action. A player who satisfies any level of rationality and not consistent beliefs could play a. Because of this, our estimate of consistent beliefs will necessarily be an upper bound on the proportion of players who satisfy consistent beliefs. To alleviate overestimation, (a, a) was chosen to be somewhat undesirable. Playing a in G3 only ever gives moderate payoffs and as player 2, there is the possibility of getting a zero payoff. In other words, believing your opponent believes you are playing a should not be immediate unless you have consistent beliefs.

Strategic choices from games G1, G2 and G3 permit us to categorize subjects into groups that are: boundedly rational (R0 and R1 types), sufficiently rational with inconsistent belief (R2, R3, and R4 types who do not satisfy consistent beliefs) and sufficiently rational with consistent belief (R2, R3, and R4 types that satisfy consistent beliefs). This classification categorizes players into 3 epistemic types: boundedly rational types that do not satisfy the sufficient rationality conditions required to play Nash equilibrium in 2-player games, sufficiently rational types that satisfy the sufficient rationality conditions required to play

Nash equilibrium in 2-player games but do not satisfy consistent beliefs, and those that satisfy both the rationality and consistent belief requirements.

4.3 Laboratory implementation

Sessions were conducted in Arts ISIT computer labs with undergraduate students at the University of British Columbia. No subject participated in more than one treatment. Subjects made all decisions through an online interface. In order to ensure independence across subjects, subjects did not interact with one another during the experiment and were not informed of one another's decisions.

We used a within-subjects design in which subjects played all games in each session. Each subject played the games G1, G2 and G3 in each of the player positions, for a total of 10 games.[14] The within subject design allows us to estimate the level of rationality and consistent beliefs for each subject.

Subjects played the games in a random order without feedback. Subjects were required to spend at least 90 seconds on each of the games. Once subjects made choices in all games they were given the opportunity to revise their choices (without any feedback). If subjects chose to revise their choices they could review each of the games (in the same original order) and make changes to their choices.

Subjects were paid a $5 showup fee and received payment from their choices in one of the games. One of the games was randomly selected for payment at the end of the experiment. Subjects were randomly and anonymously assigned to 2 or 4 player groups (depending on whether the game to be paid was a 2- or 4-player game). Subjects were paid based on their choice and the choices of their group members in the selected game and received the dollar value of their payo in the game. Subjects were paid in cash at the end of the experiment. The average session lasted 45 minutes and the average subject earned approximately $17 dollars (maximum payment was $25 and minimum payment was $7, including showup fee). Payments were in Canadian dollars.

Instructions were read aloud by the experimenter at the beginning of the session. The instructions to all subjects were the same. Subjects then completed a short understanding quiz to make sure they understood the instructions. Complete instructions and quiz can be found in Appendix E.

The dataset includes observations from 6 sessions, with between 12 and 16 subjects in each. In total, 80 subjects are represented in the dataset.

[14] 4 additional, related games were played by each subject. The data from those games is not analyzed in this paper

5 Experimental results

Figure 9 gives the proportion of subjects who satisfy each order of rationality. 94 percent of subjects are rational. That is they play a dominant strategy whenever the game has one. Thus, 6 percent of the population is irrational. The 1st-order rational types are those subjects who are rational but not 2nd-order rational. They account for 22 percent of the subject pool. The 2nd-order rational types satisfy rationality and 1st-order belief in rationality but not 2nd-order belief in rationality. 27 percent of subjects are 2nd-order rational. 25 percent of subjects are 3rd-order rational: satisfying rationality and 2nd-order belief in rationality but not 3rd-order belief in rationality. And, lastly, 20 percent of the subject pool are at least 4th-order rationality: satisfying rationality and at least 3rd-order belief rationality.

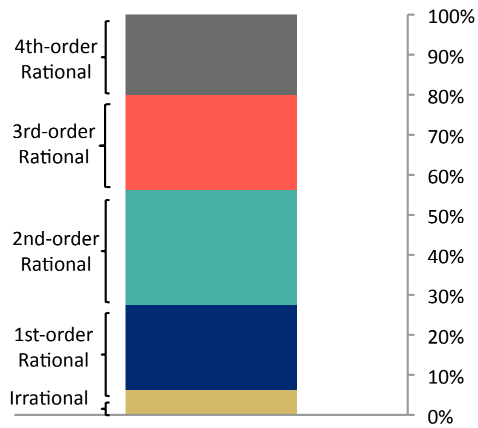

Figure 9: Subjects classified by order of rationality

The estimated levels of rationality suggest that approximately 28 percent of subjects fail to play Nash equilibrium because of a failure of bounded rationality. This includes subjects who are irrational and those who are rational but fail to account for the rationality of others. The sufficient epistemic conditions for Nash equilibrium (in 2-player games) require a subject to be rational and to have at least 1st-order beliefs that others are rational. 72 percent of the population satisfies the sufficient rationality conditions for Nash equilibrium.

Next, we estimate whether each subject satisfies consistent beliefs based on their choices in game G3. Subjects are categorized into 3 epistemic types: (1) boundedly rational types: types that irrational or are rational but fail to account for the rationality of others to the 1st-order, (2) sufficiently rational/inconsistent types: types that satisfy the sufficient rationality conditions of Nash equilibrium but do not satisfy consistent beliefs, and (3) sufficiently rational/consistent types: types that satisfy all three sufficient conditions for Nash equilib-

rium. **Not a single subject in our sample satisfies all 3 sufficient conditions for Nash equilibrium.** It is also important to notice that while failure of Nash equilibrium is due in part to bounded rationality, failure to play Nash equilibrium rarely occurs because subjects are not rational. In addition, the failure to play Nash equilibrium is largely due to a failure of having consistent beliefs. 28 percent of the subject pool is boundedly rational and 72 percent of the subject pool is su□ciently rational with inconsistent beliefs.

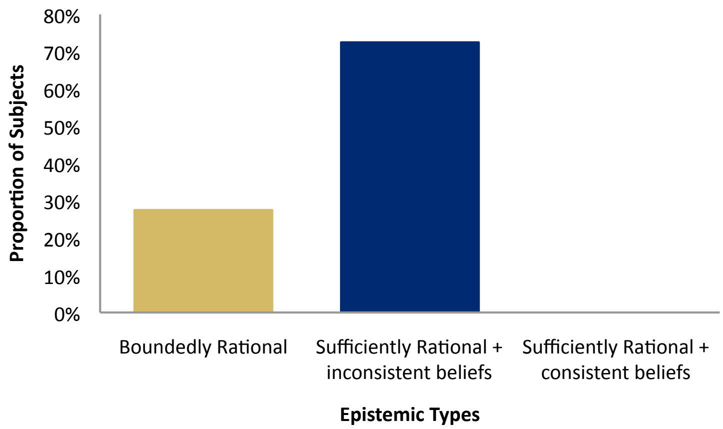

Figure 10: Subjects classified by epistemic type

The proportion of subjects who played each of the three actions $\{a, b, c\}$ for each of the 10 games are given in Figure 11. The bold numbers highlight the Nash action in each game. While a small proportion of subjects did play the Nash equilibrium strategy a as player 1 in game G3, only 1 subject played a as player 2. None of the subjects played a as both player 1 and player 3 in game G3.[15]

5.1 Secondary results

During the experiment, each subject answered a short 5-question quiz after the instructions were read aloud. The quiz was designed to test the subject's understanding of the game structure and to make sure they understood the instructions. The quiz can be found in Appendix E. 16 subjects failed the quiz. Table 3 breaks down the number of subjects who

[15]The Nash equilibrium action is played rarely in G3. Only 1 player played the Nash equilibrium action as player 2. This of course does not mean that no one will play actions consistent with a Nash equilibrium in other games. We can easily come up with games where we would expect over 90 percent of subjects to play the Nash equilibrium (i.e. games that are dominance solvable in 1 round). We however likely do not want to think of players satisfying consistent beliefs in some games and not others. That is because, our estimation of consistent beliefs based on observing Nash equilibrium behavior only determines an upper bound on the proportion of players who have consistent beliefs. Thus, it really only matters what the lowest proportion of Nash equilibrium compliance is across games.

	1				2				3	
	1	2	3	4	1	2	3	4	1	2
	. 25	.075	**. 12**	**. 50**	.562	**.475**	.163	0	**.300**	**.012**
	.025	**. 00**	.088	.037	.025	.475	**.7 7**	.013	.150	.338
	.150	.025	0	.013	**.413**	.050	.050	**. 7**	.550	.650

Figure 11: Proportion of subjects playing each action by game

failed the quiz by order of rationality. 60 percent of the irrational subjects (R0) failed the quiz. 35 percent of the 1st-order rational (R1) subjects failed the quiz. Approximately, 16

		Final Choices					
		R0	R1	R2	R3	R4	#Review
	R0	4	0	0	1	0	2
Inital Choices	R1	1	16	0	0	0	7
	R2	0	2	22	2	4	17
	R3	0	0	1	16	0	6
	R4	0	0	0	0	12	7

Table 4: Order of rationality determined by initial versus final choices

as player 4 in Games G1 and G2 will enable a player to sort out the dominance solvable nature of the games and lead them to have higher orders of rationality. This does not appear to be the case. Table 5 lists the estimated distributions of boundedly rational (R0 and R1 types) and sufficiently rational (R2, R3, and R4) types under different order assumptions. The estimated order distribution is largely unaffected by the order in which subjects play the games. There is no systematic effect of experiencing the player 4 (P4) game as the first few games (columns (2)-(4)) or as the last few games (columns (5)-(6)). Also, we cannot reject the null hypothesis that any of the restricted sample distributions (columns (2)-(6)) are equivalent to the full sample distribution (column (1)) under any standard significance levels using the Pearson chi-squared test statistic.

Type	(1) All orders	(2) 1 P4 game in first 4 games	(3) Both P4 in first 6 games	(4) Both P4 in first 8 games	(5) Both P4 in last 6 games	(6) Both P4 in last 8 games
Boundedly Rational	.28	.24	.33	.34	.35	.29
Sufficiently Rational	.72	.76	.67	.66	.65	.71
N	80	45	21	32	17	28

Table 5: Sample distributions under different game orders

The estimated distribution of orders depends upon the assumption A3. If A3 holds then subjects are not misidentified. However, if A3 does not hold, a subject's order of rationality may be mispecified. However, it is possible to bound our observed rationality distribution under different assumptions on subjects' behavior. For example, suppose a player satisfied kth-order of rationality, did not satisfy A3, but played all actions in the kth-order rationalizable set with equal probability. Then the probability that a kth-order rational player is really a (k-1)th-order rational player is at most 1/9. Additionally, the behavior of irrational types (R0) have not been specified. If an irrational type plays randomly, then there is a 1/81 chance that an irrational type would get assigned as a R1 type, a 1/243 chance an

irrational type would get assigned as an R2 type, and so on. We might expect then that 1 of our 80 subjects was misidentified as a R1 type rather than a R0 type.[17]

6 Discussion

There are a number of alternative solution concepts to Nash equilibrium. Popular alternatives are rationalizability, quantal response equilibrium (QRE), and level-k and cognitive hierarchy models. Each of these solution concepts relaxes at least one of the assumptions underlying Nash equilibrium. Rationalizability relaxes the consistent belief assumption of Nash equilibrium but requires players to satisfy common knowledge of rationality: satisfy rationality and kth-order belief in rationality as k tends to infinity (Bernheim, 1984). QRE relaxes the rationality requirements of Nash equilibrium but maintains the consistent belief component. Level-k and cognitive hierarchy models maintain the assumption that subjects are rational but relax the consistent belief assumption of Nash equilibrium by allowing players to hold heterogenous beliefs about the rationality of others.

In addition to imposing different assumptions on players' rationality and belief structures, these alternative solution concepts impose additional structural assumptions. Rationalizability, in the most general terms, imposes no additional assumptions, but only makes precise predictions in rationalizable games. When testing the model it is often supposed that players play all possible rationalizable actions with equal probability. The QRE model imposes a logistic error structure to describe the probability of error in best responding. Cognitive hierarchy and level-k solution concepts impose a particular behavior for level-0 types and particular assumptions about the beliefs of each level about the levels of other players. The success of these alternative models is typically judged by the likelihood of a given model to explain a given data set evaluated at likelihood-maximizing parameter estimates. Thus, each solution concept is judged based on the complete package of assumptions about rationality, beliefs and additional structural assumptions.

The standard methodology makes it difficult to separately identify the role that each of these assumptions plays in explaining a given data set. It is particularly difficult to separate between the level-k, cognitive hierarchy and QRE models because the structural assumptions these solution concepts impose tend to lead to the same behavior in many

[17] It is also possible that an irrational type does not play randomly and rather responds to incentives in some manner, but is not rational. Wright (2013) finds that defining L0 as taking into account salient features like highest payoff improves the fit of level-k models in normal-form games. For example, it is possible to imagine an irrational type as one that chooses the action that gives him the chance at the highest payoff in each game. In our design, this type of irrationality is not distinguished from rationality. At the extreme, this could shift our type estimation down by one, thus 45 percent of subjects will be identified as 2nd-order rational and higher while 55 percent identified as rational or irrational.

environments. For example, because subjects always make mistakes under the QRE model, players end up best responding to a mixed strategy that puts positive weight on all actions. But, in the level-k model, if level-0 types play all actions with equal probability, as is the typical specification for level-0 behavior, level-1 subjects end up best responding to a mixed strategy with positive weight on all actions. Further, both models allow for players to believe others are making mistakes (and to make mistakes themselves). This is explicitly modeled in the QRE framework as systematic mistakes when best responding. However, because level-0 types can be irrational, as long as higher levels always put some weight on level-0 types (this is true in cognitive hierarchy models), every level-k type believes his opponents make errors with some probability. And, each type in the level-k and cognitive hierarchy model is allowed to make an error when best responding when the models are fitted to the data under maximum likelihood.

The approach in this paper allows us to abstract from the additional structural assumptions and directly assess the assumptions of rationality and beliefs that underlie these solution concepts. Individual level rationality estimates rule out the ability of QRE and rationalizability to explain the experimental data based simply on the assumptions these solution concepts embody. We estimate that 94 percent of subjects are rational.[18] This rules out the assumption that subjects make systematic mistakes. Likewise, rationalizability requires homogeneity in players' beliefs about the rationality of others. This again is ruled out by our rationality estimates. The level-k model is the most promising as it incorporates the assumption that players are rational, relaxes the assumption of consistent beliefs, and allows for heterogeneity in subjects' beliefs about the rationality of their opponents. We take a closer look at these solution concepts below.

Definition. Consider an n-player ring game. The strategy $\sigma \quad (S_1) \times \cdots \times \quad (S_I)$ is **rationalizable** if for all $i \quad I$ and for all $a \quad supp\{\sigma_i\}$, then $a \quad \bigcap_{k=0}^{\infty} \bar{R}_i^k$, where $\bar{R}_i^k$ is the kth-order rationalizable set defined earlier.

QRE is typically modeled by assuming players make mistakes according to a logistic error function.

Definition. Consider an n-player ring game. The strategy $\sigma \quad (S_1) \times \cdots \times \quad (S_I)$ is a **quantal response equilibrium (QRE)** of Γ if for all $i \quad I$ and for all $a \quad supp\{\sigma_i\}$, then

[18] In addition, 51 subjects are assigned to R1-R4 types because they match a predicted type profile exactly (no mistakes in best-responding in 8 games). The remaining matched subjects make 1 error, though this error is predominately (21 of 24 subjects) in not matching the A3 assumption (thus not necessarily a mistake in best-responding).

$$\sigma_i(a) = \frac{exp(\lambda \cdot u_i(a, \sigma_{p(i)}))}{\sum\limits_{s_i \in S_I} exp(\lambda \cdot u_i(s_i, \sigma_{p(i)}))}$$

The level-k model is defined below. This paper does not differentiate between level-k and cognitive hierarchy models (hence we consider only the former for simplicity).

Definition. Consider an n-player ring game $\Gamma =< I = \{1, \ldots, I\}; S_1, \ldots, S_I; \pi_1, \ldots, \pi_i; p >$. The strategy $\sigma = \sigma_1 \times \cdots \times \sigma_I$, where $\sigma_i : \{0, \ldots, k\} \rightarrow \quad (S_i)$ is a **level-k equilibrium** in Γ if for all $i \quad I$, for all $n \quad \{1, 2, \ldots, k\}$ and for all $a \quad supp\{\sigma_i(l)\}$

$$u_i(a, \sigma_{p(i)}(n-1)) \geq u_i(a', \sigma_{p(i)}(n-1)) \qquad a' \quad S_i$$

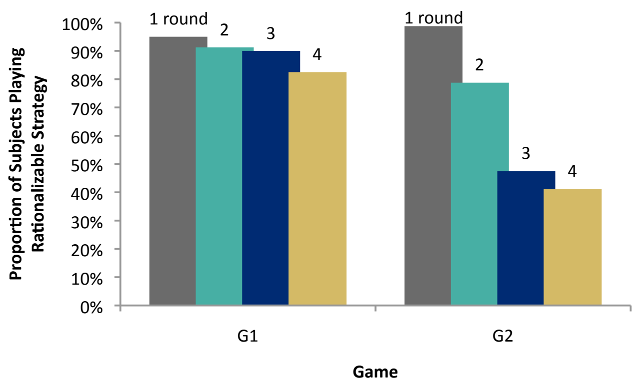

Figure 12: Proportion of subjects who play rationalizable action in G1 and G2

The majority of our subject pool satisfies rationality, which does not fit QRE assumptions unless the precision parameter, λ, is high (people do not make mistakes in best response often). But, even if a high λ could explain the high proportion of rational subjects, QRE is not able to predict the pattern of behavior across the 8 games described by G1 and G2 for any value of the precision parameter. Figure 12 gives the proportion of subjects who play the rationalizable actions in each of the player positions in games G1 and G2. The empirical frequency of playing the rationalizable strategy increases across the player positions (position 1 to 4) in both games G1 and G2. The QRE predictions are not consistent with aggregate data for any λ. For any value of λ, the QRE equilibrium does not predict that the rationalizable action will be played most often by player 4, then player 3, player 2 and player 1 in descending order. The QRE predictions are given in Figure 13.[19] The behavior is clearly inconsistent with the rationalizable solution concept as well.

[19]In addition, a high precision would not account for the failure of consistent beliefs.

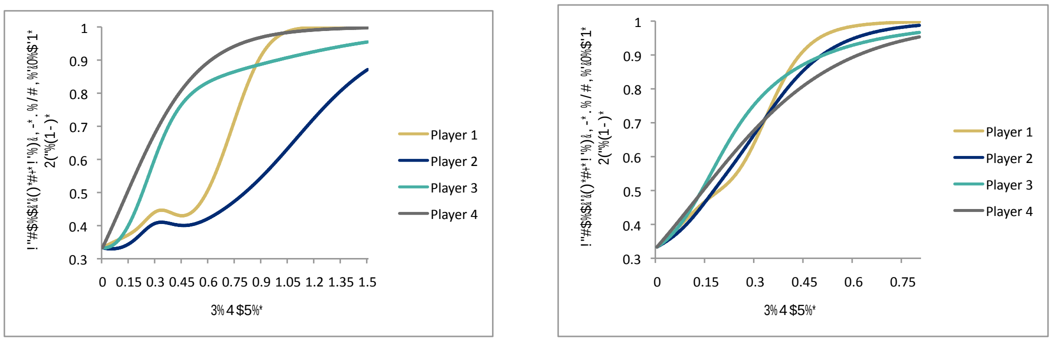

Figure 13: QRE predictions for games G1 and G2

Level-k models predict the aggregate behavior in Figure 12. Consider any distribution over levels L1, L2, L3 and L4. L4 types play the rationalizable strategy as player 1. L3, and L4 types play the rationalizable strategy as player 2. L2, L3 and L4 types play the rationalizable strategy as player 3. L1, L2, L3 and L4 types play the rationalizable strategy as player 4. No matter what the distribution over levels, or the specification for L0, the proportion of subjects playing the rationalizable strategy will weakly increase over the player positions 1-4. Further, depending on the specification of L0, the level-k model can capture the differences in rationalizable play across games G1 and G2. If we specify L0 behavior as uniformly random, all levels will play the rationalizable action in G1. However, in G2 only L4 subjects play the rationalizable action as player 1, only L4 and L3 subjects play the rationalizable action as player 2, only L2, L3, and L4 subjects play the rationalizable action as player 3, and all levels play the rationalizable action as player 4.

In addition, kth-order rationality and assumption A3 captures the main assumptions defining the behavior of an Lk type. L1 types are rational and do not respond to changes in 1st-order beliefs as they best respond to a fixed level-0 behavior. L2 types are 2nd-order rational and and do not respond to changes in 2nd-order beliefs. L3 types are 3rd-order rational and do not respond to changes in 3rd-order beliefs, and so on. Thus, the distribution of orders of rationality estimated in Figure 9 gives us a 'L0' independent estimate of the level-k distribution. The distribution of levels in Figure 9 puts more weight on higher levels than is generally estimated in level-k experiments. This could suggest that the usual assumption of uniform L0 biases the estimation of the level distribution.

The experimental design in this paper also allows for a direct test of the limited depth of reasoning assumption that underlies the level-k and cognitive hierarchy models. These models assume that players do not base optimal behavior on higher-order beliefs. This assumption has not previously been tested directly.[20] Our unique experimental design allows

[20] Analysis of subjects' search patterns (over payoffs) suggests that reasoning conforms to differential and limited depths of reasonings (Costa-Gomes et al. (2001); Costa-Gomes and Crawford 2006; Brocas et al.

us to asses this assumption because it isolates changes in higher-order beliefs while keeping lower-order beliefs constant. Approximately 80 percent of our subject pool does not respond to changes in 3rd-order beliefs.

References

Andrew, S., Keith, W., and Charles, W. (1994). A Laboratory Investigation of Multiperson Rationality and Presentation Effects. *Games and Economic Behavior*, 6:445–468.

Aumann, R. and Brandenburger, A. (1995). Epistemic Conditions for Nash Equilibrium. *Econometrica*, 63:1161–1180.

Beard, T. R. and Beil, R. O. (1994). Do People Rely on the Self-Interested Maximization of Others? An Experimental Test Management Science. *Management Science*, 40(2):252–262.

Bernheim, D. (1984). Raitonalizable strategic behavior. *Econometrica*, 52:1007–1028.

Brocas, I., Carrillo, J. D., Wang, S. W., and Camerer, C. F. (2009). Measuring Attention and Strategic Behavior in Games with Private Information. *cEPR Discussion Paper No. DP7529*.

Burchardi, K. B. and Penczynski, S. P. (2010). Out of your mind: estimating the level-k model. *working paper*.

Camerer, C., Johnson, E. J., Rymon, T., and Senc, S. (2002). Detecting Failures of Backward Induction: Monitoring Information Search in Sequential Bargaining. *Journal of Economic Theory*, 104:16–47.

Camerer, C. F., Ho, T.-H., and Chong, J.-K. (2004). A Cognitive Hierarchy Model of Games. *Quarterly Journal of Economics*, 119(3):861–898.

Costa-Gomes, M. and Crawford, V. P. (2006). Cognition and Behavior in Two-Person Guessing Games: An Experimental Study. *American Economic Review*, 96(5):1737–1768.

2009). But in standard 2-player games, a subject with any depth of reasoning greater than two must look up both her own payoffs and her partners payoffs (tight link between higher- and lower-order beliefs about payoffs in standard games). Thus, structural assumptions over the order in which players search the same payoff information are necessary to differentiate between different depths of reasoning. The ring game simplifies this inquiry because a player's lower- and higher-order beliefs are no longer linked and players with different depths of reasoning base behavior on different payoffs (assumptions about the order of search is no required) and we get insight into reasoning from choice data alone.

Costa-Gomes, M., Crawford, V. P., and Broseta, B. (2001). Cognition and Behavior in Normal-Form Games: An Experimental Study. *Econometrica*, 69(5):1193–1235.

Costa-Gomes, M. A. and Weizsäcker, G. (2008). Stated Beliefs and Play in Normal-Form Games. *The Review of Economic Studies*, 75:729–762.

Healy, P. J. (2011). Epistemic Foundations for the Failure of Nash Equilibrium. *unpublished*.

Ho, T.-H., Camerer, C., and Weigelt, K. (1998). Iterated Dominance and Iterated Best Response in Experimental 'p-Beauty Contests'. *American Economic Review*, 88(4):947–969.

Huyck, J. B. V., Wildenthal, J. M., and Battalio, R. C. (2002). Tacit Cooperation, Strategic Uncertainty, and Coordination Failure: Evidence from Repeated Dominance Solvable Games. *Games and Economic Behavior*, 38(1):156–175.

Jackson, M. (2005). *Advances in Economics and Econometrics, Theory and Applications: Ninth World Congress of the Econometric Society*, chapter The economics of social networks. Cambridge University Press.

Johnson, E. J., Camerer, C., Rymon, T., and Sen, S. (1993). *Frontiers of Game Theory*, chapter Cognition and Framing in Sequential Bargaining for Gains and Losses, pages 27–47. MIT Press.

Kearns, M. (2007). *Algorithmic Game Theory*, chapter Graphical Games. Cambridge University Press.

McKelvey, R. D. and Palfrey, T. R. (1995). Quantal Response Equilibria for Normal Form Games. *Games and Economic Behavior*, 10:6–38.

Nagel, R. (1995). Unraveling in Guessing Games: An Experimental Study. *American Economic Review*, 85(5):1313–1326.

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

Perea, A. (2007). A One-Person Doxastic Characterization of Nash Strategies. *Synthese*, 158:251–271.

Stahl, D. O. and Wilson, P. W. (1995). On Player's Models of Other Players: Theory and Experimental Evidence. *Games and Economic Behavior*, 10(1):218–254.

Tan, T. C.-C. and Werlang, S. R. C. (1988). The Bayesian foundations of solution concepts of games. *Journal of Economic Theory*, 45:370–391.

Wang, J. T.-y., Spezio, M., and Camerer, C. F. (2009). Pinocchio's Pupil: Using Eyetracking and Pupil Dilation to Understand Truth Telling and Deception in Sender-Receiver Games. *American Economic Review*, 100(3):984–1007.

Wright, J. R. (2013). Modeling Salience Improves Predictions of Human Strategic Behavior. *wo*.

Wright, J. R. and Leyton-Brown, K. (2010). Beyond Equilibrium: Predicting Human Behavior in Normal-Form Games. *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence*, page 901Ð907.

Yildiz, M. and Weinstein, J. (2007). A Structure Theorem for Rationalizability with Application to Robust Predictions of Refinements. *Econometrica*, 75:365–400.

A Omitted proofs

Proof of Proposition 1

Let t_i be a type that satisfies conditions (i) and (ii).
Let $\bar{T} = \{t_{-i} \quad T_{-i} | b^i(t_i)(t_{-i}) > 0\}$.
Since t_i is rational it must be true that for any $s \quad supp\{\sigma_i\}$

$$u_i(s, \sigma_{-i}, \theta_i) \geq u_i(a', \sigma_{-i}, \theta_i)$$

$a' \quad S_i$. In other words, σ_i is a best response to σ_{-i}
Further, since $t_i \quad R_i^2 \quad t_{-i} \quad R_{-i}^1$ for all $t_{-i} \quad \bar{T}$
And, since t_i has consistent beliefs it must be that $b^{-i}(t_{-i})(t_i) = 1$ for all $t_{-i} \quad \bar{T}$.
Therefore, for any $t_{-i} \quad \bar{T}$ and any action $s \quad supp\{\hat{s}_{-i}(t_{-i})\}$

$$u_{-i}(s, \sigma_i, \theta_{-i}) \geq u_{-i}(s', \sigma_i, \theta_{-i})$$

$s' \quad S_{-i}$. In other words, $\hat{s}_{-i}(t_{-i})$ is a best response to σ_i for all $t_{-i} \quad \bar{T}$
But, since $supp\{\sigma_{-i}\} = supp\{\bigcup_{t_{-i} \in \bar{T}} \hat{s}_{-i}(t_{-i})\}$. It must be that for any $s \quad supp\{\sigma_{-i}\}$

$$u_{-i}(s, \sigma_i, \theta_{-i}) \geq u_{-i}(s', \sigma_i, \theta_{-i})$$

$s' \quad S_{-i}$. And, hence σ_{-i} is a best response to σ_i
Therefore, σ is a Nash Equilibrium.□

B Relaxing belief to p-belief

In the experiment ran in this paper, we technically do not estimate whether a subject's beliefs hold with probability 1. This is because for finite games, the behavioral implications of a belief with probability 1 will be the same for a belief with probability p, for high enough p. In this appendix, we show that allowing for a relaxation of belief to that of p-belief (belief with probability 1 to belief with probability p) is consistent with our experimental results.

For example, throughout this paper we have described 1st-order belief in other's rationality as believing you opponent is rational with probability 1. However, because the behavioral implications of belief that your opponent is rational with probability 1 are the same as the epistemic condition of believing that your opponent is rational with at least probability p for some high enough p, we want to relax the notion of belief to that of p-belief. Rationality and consistent beliefs are likewise relaxed to allow for the notion of p-belief.

We can say that a type p-believes an event if he believes that event with probability at least p.

Definition. A type t_i ***p*-believes** an event $E \in T_{p(i)}$ if $b^i(t_i)(E) \geq p$. Let the set $B^i_p(E) = \{t^i \in T^i | b^i(t^i)(E) \geq p\}$ be the set of types for player i that p-believe event E.

We define the following sets in order to define higher-order p-beliefs about the rationality of others. A type satisfies p-belief in rationality if he puts weight at least p on types of his opponent that are rational.

$$\begin{aligned} R^1_{i,p} &= \{t_i \in T^i | t_i \text{ is rational}\} \\ R^{m+1}_{i,p} &= R^m_{i,p} \cap B^i_p(R^m_{p(i),p}) \end{aligned}$$

Definition. If $t_i \in R^m_{i,p}$ then we say that t_i satisfies **mth-order p-rationality**

We can define a relaxed version of consistent beliefs.

Definition. A type t_i has **consistent p-beliefs** if for all $t_{-i} \in T_{-i}$ such that $b^i(t_i)(t_{-i}) \geq 0$ then $b^{-i}(t_{-i})(t_i) \geq p$.

It can be shown that $R^m_{i,p} = R^m_i$ for each of the players in Games G1 and G2 as long as $p \geq \frac{7}{8}$ ($p \geq \frac{3}{4}$ for player 1). In other words, the behavioral predictions under rationality and kth-order $p-$belief in rationality are the same as the behavioral predictions under rationality and kth-order 1-belief in rationality (when p is sufficiently high). Thus, we do not identify 1-belief in rationality but rather a lower-bound $\bar{p}$ such that a player satisfies a $\bar{p}$-belief in

rationality. Given this we observe that a R4 subject satisfies rationality and 3rd-order $\frac{3}{4}$-belief in rationality, a R3 subject satisfies rationality and 2nd-order $\frac{7}{8}$-belief in rationality, a R2 subject satisfies rationality and 1st-order $\frac{7}{8}$-belief in rationality, and a R1 subject satisfies rationality.

Thus, we do not actually estimate whether a subject believes others are rational with probability 1, but rather that they believe others are rational with at least probability $\frac{7}{8}$. Our identification of consistent beliefs, relies along the assumption that subject's satisfied the su□cient rationality conditions for Nash Equilibria in G3. Thus, it is necessary to relax the su□cient epistemic characterization of Nash equilibrium to allow for a p-belief in rationality. This is established in the two results below. It is show that that for any $p \geq \frac{6}{7}$, as long as a subject is rational and satisfies 1st-order p-belief in rationality and consistent p-beliefs then she must play the Nash strategy a in game G3 (as player 1 and 2). Thus, taking subjects who are classified as R2 or higher for games G1 and G2, ensure that they satisfy rationality and 1st-order $\frac{7}{8}$-belief in rationality. Thus, if they fail to play the Nash equilibrium in G3 it must be because they fail to satisfy consistent $\frac{6}{7}$-belief (in other words, consistent p-beliefs requires: if player 1 plays strategy s then she must believe player 2 believes she is playing s with at least probability p).

Proposition. *Consider an 2-player ring game $\Gamma = \{1,2\}; S_1, S_2; \pi_1, \pi_2$ and a Γ-based epistemic type space $T_1, T_2; b^1, b^2; \hat{s}_1, \hat{s}_2$. There exists a $\bar{p}$ such that if t_i satisfies (i) 2nd-order p-rationality and (ii) consistent p-beliefs, for some $p \geq \bar{p}$, then there exists some Nash equilibrium σ such that $\hat{s}_i(t_i) = \sigma_i$.*

Proof.

Type t_i has consistent p-beliefs

Let $\bar{T} = \{t_{-i} \quad T_{-i} | b^i(t_i)(t_{-i}) > 0\}$.

Define $\sigma_{-i}(s) = \sum_{t_{-i} \in T_{-i}} \hat{s}(t_{-i})(s) \cdot b^i(t_i)(t_{-i})$ for all $s \quad S_{-i}$

Since t_i is rational, it must be true that for all $s \quad supp\{\sigma_i\}$

$$u_i(s, \sigma_{-i}) \geq u_i(a', \sigma_{-i})$$

$a' \quad S_i$. Thus, σ_i is a best response to σ_{-i}.

Further, since $t_i \quad R^2_{i,p} \qquad t_{-i} \quad R^1_{-i,p}$ for all $t_{-i} \quad \bar{T}$

Define $\mu_i(t_{-i}) \qquad (S_i)$ by $\mu_i(t_{-i})(s) = \sum_{t \in T_i} \hat{s}(t)(s) \cdot b^{-i}(t_{-i})(t)$ for all $s \quad S_i$

Therefore, for any $t_{-i} \quad \bar{T}$ and for any action $s \quad supp\{\hat{s}_{-i}(t_{-i})\}$

$$u_{-i}(s, \mu_i(t_{-i})) \geq u_{-i}(s', \mu_i(t_{-i}))$$

$s' \quad S_j$.

It must be that $\mu_i(t_{-i})(\sigma_i) \geq p$

Thus, by continuity, there exists a $\bar{p}(t_{-i})$ such that for any $s \quad supp\{\hat{s}_{-i}(t_{-i})\}$

$$u_{-i}(s, \sigma_i) \geq u_{-i}(s', \sigma_i)$$

$s' \quad S_j$ and $\quad p \geq \bar{p}(t_{-i})$. Thus $\hat{s}_i(t_{-i})$ is a best response to σ_i.

Let $\bar{p} = max\{\bar{p}(t_{-i}) | t_{-i} \quad \bar{T}\}$.

Then, since $supp\{\sigma_{-i}\} = supp\{\bigcup_{t_{-i} \in \bar{T}} \hat{s}_{-i}(t_{-i})\}$, σ_{-i} is a best response to σ_i

Therefore, σ is a Nash Equilibrium. □

Proof for G3:

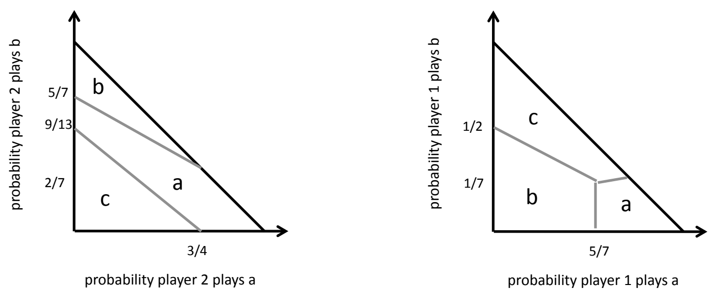

Figure 14: Best response correspondences for game G3

This proof establishes that p-belief with $p > \frac{6}{7}$ is su□cient to guarantee that (a, a) is the unique action profile in G3.

We will regularly refer to the notion of q-dominance.

Definition. In a 2-player game, an action profile $a \quad S_1 \times S_2$ is ***q*-dominant** for player i if for any $\mu \quad (S_{-i})$ with $\mu(a_{-i}) \geq q$, then

$$a_i \quad \underset{a \in S_i}{argmax}\{u_i(a, \mu, \theta_i)\}$$

First, from Figure B notice that for player 1:

(i) (a, a) is $\frac{3}{4}$-dominant

(ii) (b, b) is $\frac{5}{7}$-dominant

(iii) (c, c) is $\frac{4}{13}$-dominant

And, for player 2:

(i) (a, a) is $\frac{5}{7}$-dominant

(ii) (b, c) is $\frac{1}{2}$-dominant

(iii) (c, b) is $\frac{1}{2}$-dominant

Player 1: Now, consider the behavior of a type t_1 that satisfies a (i),(ii) and (iii) with p-belief, $p > \frac{6}{7}$.

There exists a rational type t_2 such that $b^1(t_1)(t_2) \geq p$ and $b^2(t_2)(t_1) \geq p$.

Now, consider 2 cases:

(i) Suppose t_1 plays b i.e. ($\hat{s}_1(t_1)(b) = \mu_1(b) > 0$)

Since t_1 is rational and believes opponent playing according to $\mu_2 = \hat{s}_2(t_2)$:

$\mu_2(b) \geq \frac{5}{7} - \frac{3}{5}\mu_2(a)$

Therefore:

$\mu_2(b) \geq \frac{2}{7}$

$\mu_2(a) \leq \frac{5}{7}$

$\mu_2(c) \leq \frac{2}{7}$

This also means that $\mu_1(a) + \mu_1(b) \geq p \quad \mu_1(c) \leq 1 - p$.

Now, consider the behavior of the rational type t_2. We know that with $\mu_1(c) \leq 1 - p$, as long as $p > \frac{6}{7}$, then any best response will be to mix between only a and c. $\quad \mu_2(a)+\mu_2(c) \geq p \quad \mu_2(b) \leq 1 - p$. Contradiction.

(i) Suppose t_1 plays c i.e. ($\hat{s}_1(t_1)(c) = \mu_1(c) > 0$)

Since t_1 is rational and believes opponent playing according to $\mu_2 = \hat{s}_2(t_2)$:

$u_2(b) \leq \frac{9}{13} - \frac{11}{13}p_1$

Therefore:

$\mu_2(b) \leq \frac{9}{13}$

$\mu_2(a) \leq \frac{3}{4}$

$\mu_2(c) \geq \frac{1}{4}$

This also means that $\mu_1(a) + \mu_1(c) \geq p \quad \mu_1(b) \leq 1 - p$.

Now, consider the behavior of the rational type t_2. We know that with $\mu_1(b) \leq 1 - p$, as long as $p > \frac{6}{7}$, then any best response will be to mix between only a and b. $\quad \mu_2(a)+\mu_2(b) \geq p \quad \mu_2(c) \leq 1 - p$. Contradiction.

Player 2: Consider a type t_2 that satisfies a (i),(ii) and (iii) with p-belief.

So, we know that there exists a rational type t_2 such that $b^2(t_2)(t_1) \geq p$ and $b^1(t_1)(t_2) \geq p$.

Now, consider 2 cases:

(i) Suppose t_2 plays b i.e. $\hat{s}_2(t_2)(b) = \mu_2(b) > 0$

Since t_2 is rational and believes opponent playing according to $\mu_1 = \hat{s}_1(t_1)$:

Type t_2 can mix between all three, so no conditions on μ_2.

Now, consider the behavior of the rational type t_1.

(i) Rational type t_1 is mixing between a and b

Therefore:

$\mu_2(b) \geq \frac{2}{7}$

$\mu_2(a) \leq \frac{5}{7}$

$\mu_2(c) \leq \frac{2}{7}$

And, $\mu_1(a) + \mu_1(b) \geq p$ $\quad \mu_1(c) \leq p$. This means that rational type t_2 can only mix between a and c. Therefore $\mu_2(a) + \mu_2(c) \geq p$ $\quad \mu_2(b) \leq p$. This is a contradiction.

(ii) Rational type t_1 is mixing between a and c

Therefore

$\mu_2(b) \leq \frac{9}{13}$

$\mu_2(a) \leq \frac{3}{4}$

$\mu_2(c) \geq \frac{1}{4}$

And, $\mu_1(a) + \mu_1(c) \geq p$ $\quad \mu_1(b) \leq p$. This means that rational type t_2 can only mix between a and b. Therefore $\mu_2(a) + \mu_2(b) \geq p$ $\quad \mu_2(c) \leq p$. This is a contradiction.

(i) Suppose t_2 plays c i.e. $\hat{s}_2(t_2)(c) = \mu_2(c) > 0$

Since t_2 is rational and believes opponent playing according to $\mu_1 = \hat{s}_1(t_1)$:

Type t_2 can mix between all three, so no conditions on μ_2.

Now, consider the behavior of the rational type t_1.

(i) Rational type t_1 is mixing between a and b

Therefore:

$\mu_2(b) \geq \frac{2}{7}$

$\mu_2(a) \leq \frac{5}{7}$

$\mu_2(c) \leq \frac{2}{7}$

And, $\mu_1(a) + \mu_1(b) \geq p$ $\quad \mu_1(c) \leq p$. This means that rational type t_2 can only mix between a and c. Therefore $\mu_2(a) + \mu_2(c) \geq p$ $\quad \mu_2(b) \leq p$. This is a contradiction.

(ii) Rational type t_1 is mixing between a and c

Therefore

$\mu_2(b) \leq \frac{9}{13}$

$\mu_2(a) \leq \frac{3}{4}$

$\mu_2(c) \geq \frac{1}{4}$

And, $\mu_1(a) + \mu_1(c) \geq p$ $\quad \mu_1(b) \leq p$. This means that rational type t_2 can only mix between a and b. Therefore $\mu_2(a) + \mu_2(b) \geq p$ $\quad \mu_2(c) \leq p$. This is a contradiction.

$\square$

C Assignment Mechanism

Consider 5 types: R0, R1, R2, R3 and R4. The action of each type is determined by assumptions A1-A3 and her order of rationality over the set of 8 games as given in Table 1. Let $a_i \in \{a,b,c\}^8$ be an 8-tuple representing the action of subject i in each of the 8 games. Each of the types, R1-R4, has a predicted action (or set of actions) for these 8 games, $a_k \in S_k \subset \{a,b,c\}^8$. For example, type R4 has a unique prediction. She must play the 4th-order rationalizable strategy in each of the 8 games, $S_4 = \{(a,b,a,a,c,a,b,c)\}$. Type R3 has a unique prediction for 6 of the 8 games, but can play a number of different strategies for the other 2 games, $S_3 = \{(a,b,a,a,a,a,b,c),(c,b,a,a,c,a,b,c),(b,b,a,a,b,a,b,c)\}$. Notice that A1-A3 generate an exclusion restriction and ensure that $S_4 \cap S_3 = \emptyset$. We can define similar sets for types R1 and R2. The following property holds:

Property1: For any $n \neq m \in \{1,2,3,4\}$, $S_n \cap S_m = \emptyset$.

Based on this exclusion restriction, a finite mixture model can be used to estimate a player's order of rationality. This approach assumes that all subjects have some probability, π_k, of being each type k. Given the large number of subjects that follow their type's prediction exactly, we assume a subject follows her type's prediction and make errors according to a spike error structure. This specification closely follows Costa-Gomes and Crawford (2006). If a subject is of type k, she follows one of type k's predicted action profiles with probability $1-\epsilon_k$ and makes an error in any of the 8 games with probability ϵ_k, where $\epsilon_k \in [0,\frac{1}{2})$. If she makes error in one of the 8 actions, she plays either of the other strategies with equal probability (so each other strategy with probability $\frac{1}{2}$). We can define the error density of type $k \in \{1,2,3,4\}$ for a given action profile a_i by

$$d_k(a_i,\epsilon_k) = (1-\epsilon_k)^{n_{ik}}(\frac{\epsilon_k}{2})^{8-n_{ik}}$$

where $n_{ik} = \max\limits_{a_k \in A_k}\{I^T(a_i,a_k)\cdot I(a_i,a_k)\}$ and $I:\{a,b,c\}^8 \times \{a,b,c\}^8 \to \{0,1\}^8$ is an indicator function that equals 1 at index j whenever $a_{ij} = a_{kj}$.

Based on this error density, we can define subject i's log-likelihood by

$$lnL^i(\pi,\epsilon,a^i) = ln\left[\sum_{k=1}^{4}\pi_k d_k(a_i,\epsilon_k)\right]. \quad (1)$$

Whenever $a_i = a_k$ for some $a_k \in A_k$ and some $k \in \{1,\ldots,4\}$, then $\hat{\epsilon}_k = 0$ and $\hat{\pi}_k = 1$ maximizes (1). Thus, if a subject plays an action profile that coincides exactly with a predictions of type k, she would be assigned type k. In addition, maximum likelihood can be used to assign other action profiles that do not match a type's predictions exactly, as

there exist some $a \in \{a,b,c\}^8$ such that $a \notin A_1 \cup A_2 \cup A_3 \cup A_4$. However, we will not be able to fully separate the implications of all action profiles. Depending on the action profile, the maximum likelihood parameters may not be unique. In this circumstance, subjects are assigned to the lowest type. This imposes that the rationality requirements take precedence over the A3 assumption for each type. We assign a subject as R0 whenever her action profile deviates from all of the predicted action profiles of types R1-R4 by more than 1 error.[21]

The following is true of the assignment mechanism.

Assignment Mechanism: *Consider a subject i who plays action profile $a_i \in \{a,b,c\}^8$. The following is true about the parameter $\hat{\pi}$:*

(i) $\hat{\pi}_k = 1$ *if* $n_{ik} = 8$

(ii) $\hat{\pi}_k = 1$ *if* $n_{ik} = max\{n_{i1}, n_{i2}, n_{i3}, n_{i4}\}$, $n_{ik} \geq 7$ *and* $k < j$ *for any* $n_{ij} = n_{ik} = max\{n_{i1}, n_{i2}, n_{i3}, n_{i4}\}$

(iii) $\hat{\pi}_0 = 1$ *if* $max\{n_{i1}, n_{i2}, n_{i3}, n_{i4}\} < 7$

D Beliefs about payoffs

Throughout this paper, we have restricted our attention to complete information games which implicitly assume that common knowledge of payoffs holds (assumption A1). In this section, I relax assumption A1 and discuss in which ways the results of the previous sections are affected by relaxing this assumption.

As discussed earlier, each complete information game can be represented by a payoff hierarchy defined by

$$\bar{h}^i(\Gamma) = \{\theta_{p^k(i)}\}_{k=0}^{\infty}.$$

The sequence $\bar{h}^i$ represents player i's hierarchy of beliefs about payoffs for the specified game payoffs. However, under relaxations of A1 this belief hierarchy may not be consistent with a type t_i's beliefs that are defined from a given type space.

We can formally define each types beliefs about payoff types. For any event $F \subseteq \Theta_{p^m(i)}$, $m \geq 1$, we can define player i's mth-order beliefs about the payoff types of others from the higher-order belief function g_i^m and the payoff function $\hat{\theta}$. Define $h_i^m : T_i \to \Delta(\Theta_{p^m(i)})$

$$h_i^m(t_i)(E) = \sum_{t_{p^m(i)} \in T_{p^m(i)} : \hat{\theta}_{p^m(i)}(t_{p^m(i)}) \in E} g_i^m(t_i)(t_{p^m(i)})$$

[21] Under our model, a player has a higher likelihood of being type R1-R4 than R0 whenever she makes 2 errors or less. However, we assign a subject to R0 based on more than one error to bias the mechanism in the direction of identifying R0 types.

Thus $h^i(t_i) = \{h_i^k\}_{k=0}^{\infty}$ represents type t_i's hierarchy of beliefs about the payoff types of others. Type t_i's payoff type is $h_i^0(t_i)$, she believes that her opponent $p(i)$ has payoffs $h_i^1(t_i) \quad (\Theta_{p(i)})$, she believes that her opponent $p(i)$ believes that her opponent $p(p(i))$ has payoffs determined by $h_i^2(t_i) \quad (\Theta_{p^2(i)})$, and so on.

In theory, a type could hold any hierarchy of beliefs about payoff types Θ as specified by the type space. However, in this paper an important restriction is to types whose beliefs coincide with the actual payoff hierarchy $\bar{h}^i$.

Definition. A type t_i satisfies **correct beliefs about payoffs to the mth-order** if $\{h_i^k\}_{k=0}^{m-1} = \{\bar{h}_i^k\}_{k=0}^{m-1}$.

It is possible to show that for any n-player ring game, we can relax the assumption about common knowledge of payoffs to correct beliefs about payoffs to the nth-order. This result is given in Proposition D.[22]

Proposition. *Consider an n-player ring game* $\Gamma =< I = \{1, \ldots, n\}; S_1, \ldots, S_n; \pi_1, \ldots, \pi_n; \theta_1, \ldots, \theta_n; p >$ *and a* Γ*-based epistemic type space* $< T_1, \ldots, T_n; \Theta_1, \ldots, \Theta_n; b^1, \ldots, b^n; \hat{\theta}_1, \ldots, \hat{\theta}_n; \hat{s}_1, \ldots, \hat{s_n} >$. *Suppose type* t_i *satisfies the following conditions: (i) rationality, (ii)* $(n-1)$*th-order belief in rationality, (iii) consistent beliefs, and (iv) correct beliefs about payoffs to the nth-order. Then it must be that* σ *defined by* $\sigma_i = \hat{s}_i(t_i)$ *and* $\sigma_{p^k(i)}(s) = \sum_{t_{p^m(i)} \in T_{p^m(i)}} \hat{s}_{p^m(i)}(t_{p^m(i)}) \cdot g_i^m(t_i)(t_{p^m(i)})$ *for all* $s \quad T_{p^m(i)}$, $k \quad \{1, \ldots, I-1\}$ *constitutes a Nash equilibrium.*

Thus, if the assumption of common knowledge of payoffs is relaxed to correct beliefs about payoffs to the 4th-order, none of the results in this paper will change. This means that we can allow for any higher-order uncertainty about payoffs and this will not affect the interpretation of the results. As Yildiz and Weinstein (2007) show, allowing for higher-order uncertainty while maintaining correct beliefs about payoffs to some finite-order ensures that the set of Bayesian Nash equilibria is equivalent to the rationalizable set. Thus, higher-order uncertainty (plus 4th-order rationality) will require all subjects to play the rationalizable action in G1 and G2. This means, higher-order uncertainty about payoffs alone cannot rationalize the pattern of data we observe in Figure 12. We would still need to relax rationality assumptions on top of allowing for higher-order uncertainty in payoffs in order to explain the experimental data in games G1 and G2.

Alternatively, instead of allowing only for higher-order uncertainty we could allow for payoff uncertainty at lower-orders. However, the set of payoff types would have to be non-standard in order to rationalize the data in this experiment. To see this, suppose the set of

[22] The proof is contained in the proof of Proposition 1.

payoff types $\Theta_i = \Theta$ for each player includes all payoff types that maintains that preferences over outcomes are monotone. Under this restriction, Θ could include all expected utility types with different risk attitudes. What this may not include is types with other-regarding preferences. Still, allowing for payoff heterogeneity as specified in this form is not capable of explaining this experimental data unless we additionally relax rationality assumptions as well.

To see this, consider the rationalizable sets for each payoff type defined as follows.

$$\bar{R}_i^1(\theta_i) = \left\{ s_i \in S_i | \exists \mu \in \Delta(S_{p(i)}) \text{ such that } s_i \in \underset{s \in S_i}{argmax}\{u_i(s, \mu, \theta_i)\} \right\}$$

The set $\bar{R}_i^1(\theta_i)$ is the 1st-order rationalizable set for payoff types θ_i and is set of all actions that are rational for player i with payoff type θ_i. Define the the kth-order rationalizable set as follows:

$$\bar{R}_i^k(\theta_i) = \left\{ s_i \in S_i | \exists \mu \in \Delta\left(\bigcup_{\theta \in \Theta_{p(i)}} \bar{R}_{p(i)}^{k-1}(\theta) \right) \text{ such that } s_i \in \underset{s \in S_i}{argmax}\{u_i(s, \mu, \theta_i)\} \right\}$$

Therefore the set $\bigcap_{k=1}^{\infty} \bar{R}_i^k(\theta_i)$ represents all rationalizable actions for payoff type θ_i and hence represents all the possible action that could be played in some Bayesian Nash equilibrium determined by the payoff types Θ.

Since player 4 has a dominant strategy in G1 and G2, this means that the 1st-order rationalizable sets for player 4 are singletons, $\bar{R}_4^1(\theta)(\text{G1}) = \bar{R}_4(\theta)(\text{G1}) = \{a\}$ and $\bar{R}_4^1(\theta)(\text{G2}) = \bar{R}_4(\theta)(\text{G2}) = \{c\}$ $\forall \theta \in \Theta_4$. But, then this implies that the 2nd-order rationalizable sets for player 3 are singletons, $\bar{R}_3^2(\theta)(\text{G1}) = \{a\}$ and $\bar{R}_3^2(\theta)(\text{G2}) = \{b\}$ $\forall \theta \in \Theta_3$, and so on with $\bar{R}_2^3(\theta)(\text{G1}) = \{b\}$ and $\bar{R}_2^3(\theta)(\text{G2}) = \{a\}$ $\forall \theta \in \Theta_3$, and $\bar{R}_1^4(\theta)(\text{G1}) = \{a\}$ and $\bar{R}_1^4(\theta)(\text{G2}) = \{c\}$ $\forall \theta \in \Theta_1$. This means that the set of rationalizable outcomes when we allow richness in the payoff types is still the unique set of rationalizable outcomes generated under the assumption of common knowledge of payoffs. Rationality assumptions still need to be relaxed on top of allowing for richness in payoff types in order to explain the experimental data in games G1 and G2.

Allowing for payoff heterogeneity, however, may change our interpretation of the data from G3. If subjects do not play the Nash equilibrium in G3, then we assumed that they had inconsistent beliefs. This, however, is based on the assumption about correct beliefs about payoffs to the 2nd-order. If players have different payoff types than it might be that they have consistent beliefs but different beliefs about payoffs. For example, if player 1 has the

payoff type determined by the monetary payoffs in the experiment but believes that player 2 has some other payoff type. Then, if player 1 plays c, believes player 2 is playing c and believes player 2 believes she is playing c, it is possible for player 1 to be rational, have consistent beliefs and believe player 2 is rational as long as player 2's payoff type is such that c is a best response to c for player 2.

E Instructions and quiz

Do not use the BACK or REFRESH Buttons

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.

To insure best results for yourself 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 one of the experimenters will approach you.

The Basic Idea

You will play 14 games. In each of these games, you will be randomly matched with other participants currently in this room. For each game you will choose one of three actions. Each other participant in your game will also choose one of three actions.

		Your Earnings		
		Player 2's actions		
		d	e	f
Your actions	a	10	4	16
	b	20	8	0
	c	4	18	12

		Player 2's Earnings		
		Player 3's actions		
		g	h	i
Player 2's actions	d	12	16	4
	e	0	12	8
	f	4	4	20

		Player 3's Earnings		
		Your actions		
		a	b	c
Player 3's actions	g	20	12	8
	h	6	8	18
	i	0	16	4

Your earnings will depend on the combination of your action and player 2's action. These earnings possibilities will be represented in a table like the one above. Your action will determine the row of the table and player 2's action will determine the column of the table. **You may choose action a, b, or c** and player 2 will choose action d, e, or f. The cell corresponding to this combination of actions will determine your earnings.

For example, in the above 3-player game, if you chose a and player 2 chooses d, you would earn 10 dollars. If instead player 2 chose e, you would earn 4 dollars.

Player 2 and Player 3's earnings are listed in the other two tables. Player 2 may choose action d, e or f and Player 3 may choose action g, h, or i. Player 2's earning depends upon the action he chooses and the action player 3 chooses. Player 3's earnings depend upon the action he chooses and the action you choose.

For example, if you choose c, player 2 chooses e, and player 3 chooses h then you would earn 18 dollars, player 2 would earn 12 dollars and player 3 would earn 18 dollars.
When you start each new game, you will be randomly matched with different participants. We do our best to ensure that you and your counterparts remain anonymous.

The earnings tables in a new game are not always the same as in the previous game, so you should always look at the earnings carefully at the beginning of each game. You will be required to spend at least 90 seconds on each game. You may spend more time on each game if you wish.

Earnings

You will earn a show-up payment of $5 for arriving to the experiment on time and participating.

In addition to the show-up payment, one game will be randomly selected for payment at the end of the experiment. Every participant in this room, will be paid based on their actions and the actions of their randomly chosen group members in the selected game. Any of the games could be the one selected. So you should treat each game like it will be the one determining your payment.

You will be informed of your payment, the game chosen for payment, what action you chose in that game and the action of your randomly matched counterpart only at the end of the experiment. You will not learn any other information about the actions of other player's in the experiment. The identity of your randomly chosen counterparts will never be revealed.

Frequently Asked Questions

Q1. Is this some kind of psychology experiment with an agenda you haven't told us?
Answer. No. It is an economics experiment. If we do anything deceptive or don't pay you cash as described then you can complain to the campus Human Subjects Committee and we will be in serious trouble. These instructions are meant to clarify how you earn money, and our interest is in seeing how people make decisions.

Do not use the BACK or REFRESH Buttons

Quiz

		Your Earnings		
		Player 2's actions		
		d	e	f
Your actions	a	10	4	16
	b	20	8	0
	c	4	18	12

		Player 2's Earnings		
		Player 3's actions		
		g	h	i
Player 2's actions	d	12	16	4
	e	0	12	8
	f	4	4	20

		Player 3's Earnings		
		Your actions		
		a	b	c
Player 3's actions	g	20	12	8
	h	6	8	18
	i	0	16	4

Consider the above game. You are Player 1. Your earnings are given by the blue numbers. You may choose a or b or c.

1. Your earnings depend on your action and the action of which other player?
   (a) Player 1
   (b) Player 2
   (c) Player 3

2. Suppose you choose a, Player 2 chooses f, and Player 3 chooses i. What will your earnings be?
   (a) 10
   (b) 0
   (c) 16
   (d) 6

3. Suppose Player 2 chooses d and Player 3 chooses h. Which action will give you the highest earning?
   (a) a
   (b) b
   (c) c

4. Suppose you choose c. What is your highest possible earning?
   (a) 20
   (b) 18
   (c) 4

5. Suppose you choose b. What is your lowest possible earning?
   (a) 0
   (b) 4
   (c) 8

Continue

F Raw data

Subject	Rationality	ExactMatch	G1P1	G1P2	G1P3	G1P4	G2P1	G2P2	G2P3	G2P4	G3P1	G3P2
1	0	0	1	3	2	3	1	1	2	3	1	2
2	0	0	2	3	2	2	2	3	3	3	3	3
3	0	0	2	1	1	1	1	2	2	2	3	2
4	0	0	3	1	2	2	1	2	2	3	3	3
5	0	0	3	2	2	1	3	1	1	3	1	3
6	1	0	1	2	2	1	3	2	2	3	1	3
7	1	0	1	2	1	1	1	2	3	3	1	3
8	1	0	1	2	1	1	1	2	3	3	1	3
9	1	0	1	2	1	1	1	2	3	3	3	3
10	1	0	1	1	1	1	1	2	1	3	3	2
11	1	0	1	2	1	1	1	1	1	3	3	3
12	1	0	3	2	1	1	1	2	1	3	2	3
13	1	0	3	2	2	2	3	2	2	3	3	2
14	1	1	1	2	1	1	1	2	1	3	1	3
15	1	1	1	2	1	1	1	2	1	3	2	3
16	1	1	1	2	1	1	1	2	1	3	3	3
17	1	1	1	2	1	1	1	2	1	3	1	2
18	1	1	1	2	1	1	1	2	1	3	1	2
19	1	1	1	2	1	1	1	2	1	3	1	2
20	1	1	1	2	1	1	1	2	1	3	1	3
21	1	1	3	1	1	1	3	1	1	3	3	2
22	1	1	3	2	2	1	3	2	2	3	1	1
23	2	0	1	2	1	1	3	2	2	3	3	3
24	2	0	1	2	1	1	3	2	2	3	3	3
25	2	0	1	2	1	1	3	2	2	3	3	3
26	2	0	1	2	1	1	3	2	2	3	3	2
27	2	0	1	1	1	1	1	2	2	3	3	3
28	2	0	1	2	1	1	3	2	2	3	3	2
29	2	0	1	2	1	1	3	2	2	3	3	3
30	2	0	1	2	1	1	3	2	2	3	3	3
31	2	0	1	2	1	1	3	2	2	3	1	2
32	2	0	3	2	1	1	3	3	2	3	2	3
33	2	1	1	2	1	1	1	2	2	3	3	2
34	2	1	1	2	1	1	1	2	2	3	3	3
35	2	1	1	2	1	1	1	2	2	3	3	2
36	2	1	1	2	1	1	1	2	2	3	2	2
37	2	1	1	2	1	1	1	2	2	3	3	3
38	2	1	1	2	1	1	1	2	2	3	3	2
39	2	1	1	2	1	1	1	2	2	3	3	3
40	2	1	1	2	1	1	1	2	2	3	3	2
41	2	1	1	2	1	1	1	2	2	3	1	3
42	2	1	1	2	1	1	1	2	2	3	2	3
43	2	1	1	2	1	1	1	2	2	3	3	3
44	2	1	1	2	1	1	1	2	2	3	1	2
45	2	1	3	1	1	1	3	1	2	3	1	2
46	3	0	1	2	1	1	2	1	2	3	2	3
47	3	0	3	2	1	1	1	1	2	3	3	3
48	3	0	3	2	1	1	1	1	2	3	3	3
49	3	1	1	2	1	1	1	1	2	3	3	2
50	3	1	1	2	1	1	1	1	2	3	3	3
51	3	1	1	2	1	1	1	1	2	3	2	3
52	3	1	1	2	1	1	1	1	2	3	1	3
53	3	1	1	2	1	1	1	1	2	3	3	3
54	3	1	1	2	1	1	1	1	2	3	1	2
55	3	1	1	2	1	1	1	1	2	3	1	3

Subject	Rationality	ExactMatch	G1P1	G1P2	G1P3	G1P4	G2P1	G2P2	G2P3	G2P4	G3P1	G3P2
56	3	1	1	2	1	1	1	1	2	3	3	2
57	3	1	1	2	1	1	1	1	2	3	2	3
58	3	1	1	2	1	1	1	1	2	3	3	3
59	3	1	1	2	1	1	1	1	2	3	3	3
60	3	1	1	2	1	1	1	1	2	3	1	3
61	3	1	1	2	1	1	1	1	2	3	1	2
62	3	1	1	2	1	1	1	1	2	3	3	3
63	3	1	3	2	1	1	3	1	2	3	1	3
64	3	1	3	2	1	1	3	1	2	3	3	3
65	4	0	1	2	1	1	3	3	2	3	2	3
66	4	0	1	2	1	1	3	3	2	3	3	3
67	4	0	1	2	1	1	3	1	1	3	2	2
68	4	1	1	2	1	1	3	1	2	3	3	2
69	4	1	1	2	1	1	3	1	2	3	1	3
70	4	1	1	2	1	1	3	1	2	3	2	2
71	4	1	1	2	1	1	3	1	2	3	3	3
72	4	1	1	2	1	1	3	1	2	3	3	3
73	4	1	1	2	1	1	3	1	2	3	2	3
74	4	1	1	2	1	1	3	1	2	3	3	2
75	4	1	1	2	1	1	3	1	2	3	3	3
76	4	1	1	2	1	1	3	1	2	3	3	2
77	4	1	1	2	1	1	3	1	2	3	1	3
78	4	1	1	2	1	1	3	1	2	3	3	3
79	4	1	1	2	1	1	3	1	2	3	1	3
80	4	1	1	2	1	1	3	1	2	3	3	3
Nash Equilibrium		**1**	**1**	**2**	**1**	**1**	**3**	**1**	**2**	**3**	**1**	**1**

Notes: ExactMatch=1 means that the action profile matches one of type R1-R4's predicted profiles exactly in games G1 and G2.
Actions 1='a' , 2='b', 3='c'

Table 6: Raw Data