Bounded Reasoning: Rationality or Cognition [*]

Amanda Friedenberg[†] Willemien Kets[‡] Terri Kneeland[§]

First Draft: January 31, 2016
This Draft: October 5, 2016

Abstract

Limited reasoning about rationality is well understood to lead to important behavioral consequences. The literature has typically viewed such limited reasoning as an artifact of cognitive bounds—loosely, an inability to reason about how the other player reasons, etc. However, in principle, subjects may not be willing to believe their opponent is "rational, believes rationality, etc..." even if the subject is capable of reasoning in this way. In this paper, we allow for the possibility that a subject's rationality bound may be lower than her cognitive bound: so two subjects with the same cognitive bound may have different rationality bounds. We develop an identification strategy that allows us to disentangle the cognitive bound from the rationality bound and to investigate the extent to which bounds on rationality are driven by bounds on cognition. Using the experimental data from Kneeland (2015), we show that rationality bounds are tighter than cognitive bounds. Rationality bounds are an important determinant of behavior, especially for subjects with high cognitive bound.

1 Introduction

The standard approach to game theory implicitly takes as given that players are strategically sophisticated. In particular, it is often assumed that players are rational and there is common reasoning about rationality: players choose an action that is a best response given their belief about the play of the game, they believe others do the same, etc. However, experimental game theory has suggested that players' behavior may instead reflect bounded reasoning about rationality. (See e.g., Nagel, 1995; Stahl and Wilson, 1995; Costa-Gomes, Crawford, and Broseta, 2001; Camerer,

---

[*] We are especially indebted to Adam Brandenburger for many helpful conversations. We also thank Marina Agranov, Martin Dufwenberg, Guillaume Fréchette, P.J. Healy, Toru Kitagawa, Brian Rogers, and participants of seminars at Barcelona GSE Summer Forum, BRIC, Caltech, and the Game Theory Society World Congress. Amanda Friedenberg thanks NSF grant SES-1358008 for financial support and Caltech, Kellogg School of Management and University College London for their hospitality.

[†] W.P. Carey School of Business, Arizona State University, Tempe, AZ 85287, U.S.A., email: amanda.friedenberg@asu.edu, website: www.public.asu.edu/∼afrieden

[‡] Kellogg School of Management, Northwestern University, Evanston, IL 60208, U.S.A., email: wkets@kellogg.northwestern.edu, website: http://wkets.org

[§] Department of Economics, University College London, Gower Street, London WC1E 6BT, U.K., e-mail: t.kneeland@ucl.ac.uk, website: http://terri.microeconomics.ca

Ho, and Chong, 2004; Crawford, Costa-Gomes, and Iriberri, 2013; Kneeland, 2015, amongst many others.) For example, a player may be rational, believe that her opponent is rational, but not that her opponent believes that she is rational.

Common reasoning about rationality requires that players have an unlimited ability to engage in interactive reasoning—i.e., to reason through sentences of the form "I think that you think that I think…" There is evidence from cognitive psychology that subjects are limited in their ability to engage in such interactive reasoning. (See e.g., Perner and Wimmer, 1985; Kinderman, Dunbar, and Bentall, 1998; Stiller and Dunbar, 2007, amongst many others.) To the extent that such ability limitations are binding in games, they may lead players to engage in bounded reasoning about rationality. But, at least in principle, there can be bounded reasoning about rationality, even if players do not face limitations in their ability to engage in interactive reasoning. For instance, given her past experiences, Ann may simply not be prepared to believe that Bob is rational. Or, she may believe that Bob is rational, but may not be prepared to believe that Bob believes she is rational. And so on.

This paper asks: Is bounded reasoning about rationality driven by limitations on players' ability to engage in interactive reasoning? Or, are there systematic bounds on reasoning about rationality that cannot be explained by such ability limitations?

Answering this question is of first-order importance. While we observe bounded reasoning about rationality in the laboratory setting, there is the question of whether those bounds are behaviorally important, when it comes to important economic and social decisions (i.e., outside of the laboratory): When players face "more important" problems, they may be prepared to "think harder." If so, their ability to engage in interactive reasoning may be endogenous to the nature of the problem. (See Alaoui and Penta, 2016.[1]) As a consequence, limitations on ability may not be binding on important decisions. If bounded reasoning about rationality is driven by limitations to engage in interactive reasoning, then such reasoning may not be significant when it comes to important decisions. That said, if bounds on reasoning about rationality arise from other sources, those bounds may well persist, even when it comes to important decisions.

To better understand this last point, consider two executives engaged in an important business decision. The executives may each be prepared to devote a high level of resources to the problem; they may reason that the other does the same, etc. That is, they may face no limitations on their ability to engage in interactive reasoning. Nonetheless, they may exhibit bounded reasoning about rationality. If the executives have previously interacted—either with each other with a population of executives—they may have observed past behavior that could not be rationalized: Based on Ann's past behavior, Bob may not be prepared to bet on the fact that she is rational. Even if Ann is rational and Bob is prepared to bet on this fact, Ann may consider the possibility that Bob considers the possibility that Ann is irrational. For instance, in the past, Ann may have chosen rationally but, given Bob's past behavior, she may have concluded that Bob did not understand important parameters of her problem (e.g., her payoffs). If so, she may reason, that Bob now is

---

[1] Brandenburger and Li (2015) instead argue such ability is game-independent.

unsure of her rationality. As such, bounded reasoning about rationality may well be important for understanding how the executives act.

Our Approach Above, we pointed out that limited ability to engage in interactive reasoning will limit the players ability to reason about their opponent's rationality. But, more broadly, it would also limit their ability to reason about how their opponent plays the game. The goal then is to identify behavior that is consistent with a player engaging in interactive reasoning about how her opponent plays the game, but inconsistent with interactive reasoning about rationality. This would indicate that bounded reasoning about rationality is not determined (entirely) by the players' ability to engage in interactive reasoning.

To better understand the idea, it will be useful to distinguish between, what we will call, cognition and rationality. Say that a player is *cognitive* if she has a theory about how to play the game—put differently, if she has a method for playing the game. Say that she is *rational* if she plays a best response given her subjective belief about how the game is played—put differently, if she maximizes her expected utility given her subjective belief about how the game is played.[2] So, a player who is rational has a method for playing the game; that is, if a player is rational, then she is also cognitive. However, a player may be cognitive and irrational; that is, a player may have a decision criteria for playing the game that departs from subjective expected utility. For instance, she may instead adopt a decision criteria based on a rule-of-thumb.

Consider a player who engages in full interactive reasoning. That is, consider a player who faces no limitations in her ability to engage in interactive reasoning. Provided we take the notion of cognition to be sufficiently broad, the player will reason that the other player is cognitive, reason that the other player reasons she is cognitive, and so on, ad infinitum. However, the player may still exhibit bounded reasoning about rationality. For instance, she may assign probability one to the other player having *some* method for playing the game, but that method may not involve maximizing his subjective expected utility.

With this in mind, we focus on subjects who are rational, i.e., play a best response given their subjective belief about the play of the game. We distinguish between reasoning about rationality and reasoning about cognition:

- *Reasoning About Rationality:* Say that Ann has a *rationality bound* of level m if she is rational, believes that Bob is rational, believes that Bob believes she is rational, and so on, up to the statement that includes the word "rational" m times, but no further.
- *Reasoning About Cognition:* Say that Ann has a *cognitive bound* of level m if she is cognitive, believes that Bob is cognitive, believes that Bob believes she is cognitive and so on, up to the

---

[2] Our use of the term rationality is consistent with the epistemic game theory literature. In other literatures, the phrase "rationality" may incorporate normative requirements about the decision-maker's preferences. There are no such normative requirements here. Moreover, we use the phrases rationality in a way that does not incorporate a restriction on what subjective beliefs a player may or may not hold. (Contrast this with the level-k and cognitive hierarchy literatures, which impose a restriction on the beliefs of level-1 reasoners. See Appendix A for a further discussion.)

statement that includes the word "cognitive" m times, but no further.

If a rational subject's rationality bound is lower than her cognitive bound, then bounded reasoning about rationality is not entirely determined by limited ability to engage in interactive reasoning. Thus, we seek to identify a gap between a subject's rationality bound and her cognitive bound. If we identify such a gap, then bounded reasoning about rationality is not (entirely) driven by a limited ability to engage in interactive reasoning.

To address the question, we take a broad stance on what we mean by *cognition*, i.e., on what we mean by a theory for how to play the game. We allow the players' theory to depend on the payoffs of the game. But it cannot depend on certain fine presentation effects. (Sections 2-3 will clarify the precise presentation effects that are ruled out.)

Preview of Results To identify a gap between reasoning about rationality vs. reasoning about cognition, we must identify rational behavior that is consistent with reasoning about cognition, but inconsistent with reasoning about rationality. We provide a novel identification strategy based on Kneeland's (2015) ring game experiment. Section 2 previews the identification strategy.

We apply this identification strategy to Kneeland's (2015) experimental dataset. We find that 12% of our subjects have a cognitive bound of 1, 22% have a cognitive bound of 2, 28% have a cognitive bound of 3, and 38% have a cognitive bound of at least 4. Moreover, there is a nontrivial gap between subjects' rationality and cognitive bounds. We find that 47% of the subjects identified as having a low rationality bound (i.e., either 1 or 2) have a higher cognitive bound. The gap between the rationality and cognitive bounds is most pronounced for subjects that have the highest level of cognition—that is, it is most pronounced for subjects whose cognitive bound is at least 4. Just over half of these subjects, 54%, have rationality bounds that are strictly less than their cognitive bound. Only 46% of subjects with cognitive bounds of at least 4 have rationality bounds that coincide with their cognitive bounds. The gap is less pronounced for lower levels of cognition: 76% of subjects with cognitive bounds of 3 have rationality bounds of 3; 81% of subjects with cognitive bounds of 2 have rationality bounds of 2.

Our identification strategy presumes that observed behavior is the result of deliberate choices on the part of subjects. An alternate hypothesis is that there is no gap between rationality and cognitive bounds and, instead, certain observed behavior is an artifact of noise. We rule out this hypothesis by estimating and simulating a noisy decision-making model. The simulations do not replicate the pattern of behavior we observe in the data, suggesting that our identified gap between rationality and cognition is not simply due to noise.

The Identified Gap The gap we identify takes a particular form: It characterizes a situation in which the subjects have non-degenerate beliefs in reasoning about rationality. This has important implications for using existing analyses to make out-of-sample predictions.

To make the points more concrete, refer to Figure 1.1. The figure reflects a class of games that depends on a parameter x, which is restricted to be in $(-\infty, 10)$. If Ann plays a strategy

		Bob	
		L	R
Ann	U	10,0	0,5
	D	x,0	10,5

Figure 1.1: $x < 10$

that survives two rounds of iterated dominance, then she must play D. If, instead, she plays a strategy that survives one round—but not two rounds—of iterated dominance, then she must play U. Importantly, these conclusions holds irrespective of the parameter x. The first case corresponds to the scenario where Ann plays a best response given a belief that assigns probability $p = 1$ to Bob's rationality. The second case corresponds to the scenario where Ann plays a best response given a belief that assigns probability $p = 0$ to Bob's rationality. Thus, if she has a degenerate belief about Bob's rationality (i.e., a belief that assigns $p \in \{0, 1\}$ to Bob's rationality), her best response would not depend on the parameter x and so our predictions would be the same across games parameterized by $x \in (-\infty, 10)$. If however, Ann assigns probability $p \in (0, 1)$ to Bob's rationality, then her best response will depend on the parameter x. That is, if her behavior is driven by non-degenerate beliefs about Bob's rationality, then our predicted behavior should vary across this class of games.

Our analysis points to the fact that, when there is a gap between the rationality and cognitive bounds, players have non-degenerate beliefs about the rationality of their opponents. For instance, consider a subject who is identified has having a rationality bound of 1 and a cognitive bound of $k \geq 2$. We observe this subject play a strategy that survives one but not two rounds of iterated dominance. However, importantly, we will be able to use her behavior to provide insight into how she reasons about her opponent's rationality. Section 5 shows that her beliefs about her opponents rationality are not degenerate. Thus, we cannot simply assume that, in other games, she will also play a strategy that survives one but not two rounds of iterated dominance. Our analysis in Section 6 will provide insight into these non-degenerate beliefs.

Related Literature There is a long history of studying iterative reasoning in games. Bernheim (1984) and Pearce (1984) defined iterative reasoning as *rationalizability*; subsequent work has drawn a relationship between rationalizability and *reasoning about rationality* (as used in this paper). A prominent and influential literature sought to study limitations on such iterative reasoning. Toward that end, the literature introduced level-k thinking (see, e.g., Nagel, 1995, Stahl and Wilson, 1995, Costa-Gomes, Crawford, and Broseta, 2001) and the cognitive hierarchy model (see, e.g., Camerer, Ho, and Chong, 2004). There is a subtle relationship between rationalizability, level-k thinking, and the cognitive hierarchy model. See Appendix A for a discussion.

The level-k and cognitive hierarchy literatures are often motivated by limitations on the players' ability to engage in interactive reasoning. (See, e.g., pg. 1313 in Nagel, 1995 or pg. 864 in Camerer, Ho, and Chong, 2004.) That said, typically, the literature identifies subjects' ability to engage in interactive reasoning based on the extent to which they iterate over best responses. For instance, a subject is identified as a level-k thinker if she performs exactly k rounds of iterated best responses. Similarly, Kneeland (2015) identifies a subject's reasoning based on the number of rounds of iterated dominance she performs. So, while the literature is often motivated by limited ability, it identifies levels of reasoning with rationality bounds. Evidence of ability limitations is typically based on auxiliary data.[3]

To the best of our knowledge, this is the first paper that can directly address whether the rationality bounds are determined by limited ability. That said, there are results in the literature that appear to speak to this question. In the context of level-k reasoning, Agranov, Potamites, Schotter, and Tergiman (2012), Georganas, Healy, and Weber (2015), Alaoui and Penta (2016), and Gill and Prowse (2016) show that subjects' rationality bounds may vary based on whether they are playing against more versus less "sophisticated" players.[4] At first glance, this variation might suggest that reasoning about rationality is not driven by limited ability: If the subjects' ability to engage in interactive reasoning is fixed—i.e., does not depend on 'who' their opponents are—then variation in the rationality bounds indicates that the bounds cannot be entirely determined by the difficulties of interactive reasoning. However, this conclusion is premature. The rationality bounds might vary, even if they are solely determined by the difficulties of interactive reasoning. This would, in particular, be the case if subjects vary their effort in interactive reasoning (i.e., how much effort they exert on "i think, you think, i think ..." sentences), based on 'who' their opponents are. Thus, absent directly identifying limitations on interactive reasoning—i.e., separate from identifying the rationality bounds—these results cannot address whether the rationality bounds are driven by limited ability.

Our analysis shows that players have non-degenerate beliefs about their opponents' rationality and, more generally, how their opponents reason about rationality. As a consequence, subjects who play k-undominated but not $(k+1)$-undominated strategies do not view all k-undominated strategies equally. Some k-undominated strategies are preferred to others. This is a point also raised by Fragiadakis, Knoepfle, and Niederle (2013). It suggests that there may be a role for bringing a richer set of non-degenerate beliefs to bear on the literature.

The remainder of this paper is organized as follows. Section 2 gives an example, which highlights the key ingredients of the identification strategy. Section 3 describes the identification strategy. Section 4 presents the main result: a gap between the rationality and cognitive bounds. Section 5 explains that, when there is a gap, subjects have non-degenerate beliefs about rationality. Finally, Section 6 argues that the observed gap is not an artifact of noisy decision-making.

---

[3] For instance, Costa-Gomes, Crawford, and Broseta (2001) and Costa-Gomes and Crawford (2006) use look-up patterns, Rubinstein (2007) uses response times, Chen, Huang, and Wang (2009) and Wang, Spezio, and Camerer (2009) use eye-tracking data, Burchardi and Penczynski (2014) use incentivized communication, Bhatt and Camerer (2005) and Coricelli and Nagel (2009) measure brain activity, etc.

[4] There are also papers that investigate the extent to which reasoning varies, based on whether a subject plays against another subject versus her own self. See, e.g., Blume and Gneezy (2010) and Fragiadakis, Knoepfle, and Niederle (2013). An analogous argument applies to that experimental design.

2 An Illustrative Example

Figures 2.1a-2.1b describe two games, G and G_*. We will write $(\mathrm{d}, \mathrm{e}_*)$ to denote the fact that a player chooses action d in G and action e_* in G_*. We often refer to such an action profile a *strategy*. Notice three features of these games. First, for Player 1 (P1, she), the payoff matrix given by G_* is a relabeling of the payoff matrix given by G. Specifically, the row a in G is labelled c_* in G_*, the row b in G is labelled a_* in G_*, and the row c in G is labelled b_* in G_*. Second, in each game, P1 has a dominant action; it is a in G and c_* in G_*. Third, in the two games, Player 2 (P2, he) has the same payoff matrix.

P2

		a	b	c
	a	12	16	14
P1	b	8	12	10
	c	6	10	8

P1

		a	b	c
	a	20	14	8
P2	b	16	2	18
	c	0	16	16

(a) Figure G

P2

		a_*	b_*	c_*
	a_*	8	12	10
P1	b_*	6	10	8
	c_*	12	16	14

P1

		a_*	b_*	c_*
	a_*	20	14	8
P2	b_*	16	2	18
	c_*	0	16	16

(b) Figure G_*

Figure 2.1: A Two-Player Example

Rationality versus Cognition: To illustrate the relationship between rationality and cognition, we focus on P1. Suppose that P1 is *rational*, in the sense that she maximizes her expected utility given her subjective belief about how P2 plays the game. Then, she would play the strategy $(\mathrm{a}, \mathrm{c}_*)$. Notice that, if she is rational, then she has a specific theory about how to play the game. In this sense, she is also *cognitive*.

But, at least in principle, P1 may be cognitive and irrational. For instance, suppose that P1 instead adopts a rule-of-thumb, in which she plays an action that could potentially lead to a payoff of 6, provided that such an action exists. She does so, even if such an action does not maximize her expected utility given her subjective belief about how to play the game. For the purpose of illustration, she adopts such a method for playing the game, because 6 is her lucky number. In this case, she would choose the "lucky-6" strategy profile $(\mathrm{c}, \mathrm{b}_*)$.

Reasoning about Rationality versus Reasoning about Cognition: To illustrate the relationship between reasoning about rationality and reasoning about cognition, we focus on P2. Throughout the discussion, we suppose that P2 is rational (and so cognitive). We will distinguish between three scenarios.

First, suppose that P2 reasons about rationality. By this we mean, P2 *believes*—i.e., assigns probability 1 to the event—that P1 is rational. In this case, he must assign probability 1 to P1 playing $(\mathrm{a}, \mathrm{c}_*)$. Because P2 also maximizes his expected utility given his belief about how P1 plays the game, he plays $(\mathrm{a}, \mathrm{b}_*)$. Notice, because P2 believes that P1 is rational, P2 also believes that P1 is cognitive. Put differently, the fact that P2 reasons about rationality implies P2 also reasons about cognition.

Second, suppose that P2 does not assign probability 1 to P1's rationality but does assign probability 1 to P1's cognition. For instance, he may assign probability .8 to the rational strategy $(\mathrm{a}, \mathrm{c}_*)$ and probability .2 to the lucky-6 strategy $(\mathrm{b}, \mathrm{c}_*)$. In this case, he maximizes his expected utility by playing $(\mathrm{a}, \mathrm{c}_*)$.

Third, suppose that, unlike the two scenarios above, P2 reasons that P1 lacks cognition. In this case, he reasons that P1 does not have a theory about how to play the game. As a consequence, he thinks that P1's behavior does not depend on specific parameters of the game—including P1's payoffs. Thus, P2 has the same belief about how P1 plays the game, in both G and G_*. That is, if he assigns probability p to P1 playing a, then he also assigns probability p to P1 playing a_*. This has important implications for how P2 plays the game. In particular, since P2 has the same payoff matrix in G and G_*, this implies that P2 plays a *constant strategy*—i.e., $(\mathrm{a}, \mathrm{a}_*)$, $(\mathrm{b}, \mathrm{b}_*)$, or $(\mathrm{c}, \mathrm{c}_*)$.

Observe that both the first and third scenarios involve no gap between reasoning about rationality and reasoning about cognition. In the first case, P2 reasons both that P1 is rational and that P1 is cognitive. In that case, he rationally plays the only strategy that survives two rounds of iterated dominance. In the third case, P2 reasons that P1 lacks cognition and, so, he does not reason that P1 is rational. In that case, he rationally plays a constant strategy. By contrast, the second scenario is an example where there is a gap between reasoning about cognition and reasoning about rationality: P2 assigns probability 1 to P1's cognition but not to P1's rationality. He, then, rationally plays a non-constant strategy—one that does not survive iterated dominance.

Identification: A player's cognitive bound must be at least as high as her rationality bound: If she lacks cognition, then she cannot be rational. So, if she reasons that the other player lacks cognition, then she cannot reason that the other player is rational.

Our question is: Does there exist a gap between the cognitive and rationality bounds? We seek a conservative estimate of the gap. With this in mind, we seek to identify:

(i) the *maximum* level of reasoning about rationality consistent with observed behavior, and

(ii) the *minimum* level of cognition consistent with observe behavior.

The example illustrates how we identify these bounds.

To identify these bounds, we assume that the observed behavior is rational, in the sense that it is consistent with maximizing a player's subjective expected utility. (Most of the observations in our dataset will be consistent with rational behavior. We will restrict attention to those observations.) As a consequence, we assume that all behavior is also cognitive. That is, we do not attempt to distinguish rational behavior from cognitive behavior. Instead, our identification focuses on reasoning about rationality and reasoning about cognition. In light of this, we focus on the observed behavior of P2. It will be useful to distinguish between three observations.

First, if we observe P2 play $(\mathrm{a}, \mathrm{b}_*)$, then we identify the subject as being 2-rational. This is because the strategy $(\mathrm{a}, \mathrm{b}_*)$ is consistent with P2 being rational and believing (i.e., assigning probability 1 to the event) that P1 is rational. Notice that this behavior is also consistent with P2 being rational and assigning (only) probability .9 to P1's rationality.[5] Despite this, we identify the subject as 2-rational. This is because we seek the maximum level of reasoning about rationality consistent with observed behavior.

Notice, because the behavior is consistent with 2-rationality it is also consistent with reasoning about cognition. Importantly, this behavior is inconsistent with a rational P2 reasoning that P1 lacks cognition: Recall, if P2 is rational and reasons that P1 lacks cognition, then P2's behavior cannot vary across G and G_*. Thus, if we observe P2 play the non-constant action profile $(\mathrm{a}, \mathrm{b}_*)$, we must conclude that P2 reasons that P1 is cognitive. As such, the minimum cognitive bound consistent with observed behavior is 2 and, so, we identify this behavior as being 2-cognitive.

Second, if we observe P2 play $(\mathrm{a}, \mathrm{a}_*)$, then we identify the subject as being 1-cognitive. In particular, this behavior is consistent with a rational P2 reasoning that P1 lacks cognition. (For instance, it is a best response for P2 to play $(\mathrm{a}, \mathrm{a}_*)$, if he assigns probability 1 to P1 also playing $(\mathrm{a}, \mathrm{a}_*)$. Such a constant belief is consistent with believing that P1 lacks cognition.) That said, the behavior is also consistent with P2 being rational and reasoning that P1 is cognitive. (For instance, it is a best response for P2 to play $(\mathrm{a}, \mathrm{a}_*)$ if he assigns probabilities $\Pr(\mathrm{a}) = \Pr(\mathrm{c}_*) = .4$ and $\Pr(\mathrm{b}) = \Pr(\mathrm{c}) = \Pr(\mathrm{a}_*) = \Pr(\mathrm{b}_*) = .3$.) Despite this, we identify this subject as being 1-cognitive. This is because we seek the minimum level of reasoning about cognition consistent with observed behavior.

Notice, this observation is inconsistent with P2 being rational and assigning probability 1 to P1's rationality. Thus, we also identify this subject as being 1-rational.

For both of these observations—i.e., behavior of $(\mathrm{a}, \mathrm{b}_*)$ and $(\mathrm{a}, \mathrm{a}_*)$—we would not identify a gap between the rationality and cognitive bounds. (In the former case, the bounds are both 2 and, in the latter case, the bounds are both 1.) If, instead, we observe P2 play $(\mathrm{a}, \mathrm{c}_*)$, we would identify a gap between the rationality and cognitive bounds: This behavior is inconsistent with rationality and belief of rationality. Thus, we identify the subject as 1-rational. But, at the same time, because this observation is not a constant action profile, it is inconsistent with rationality and reasoning that P1 lacks cognition. Moreover, we have also seen that it is consistent with rationality and belief that P1 is cognitive. Thus, we identify the subject as 2-cognitive.

Identifying the Bounds: A Comment Suppose that P2 reasons about P1's rationality. In the above discussion (and, indeed, throughout the paper) we think of this scenario as one in which

---

[5] For instance, it is a best response for P2 to play $(\mathrm{a}, \mathrm{b}_*)$, if he assigns probabilities $\Pr(\mathrm{a}) = \Pr(\mathrm{c}_*) = .9$ and $\Pr(\mathrm{b}) = \Pr(\mathrm{c}) = \Pr(\mathrm{a}_*) = \Pr(\mathrm{b}_*) = .05$.

P2 believes—i.e., assigns probability 1 to the event—that P1 is rational. If P2 assigns probability $p = .8$ to P1's rationality, we think of this as a departure from reasoning about rationality.

In a similar fashion, we think of P2 reasoning about P1's cognition as believing–i.e., assigning probability 1 to the event—that P1 is cognitive. But, if P2 assigns probability $p = .8$ to P1's cognition, we do not think of this as a departure from reasoning about cognition. Such a belief would exhibit an ability to engage in interactive reasoning. Thus, we only identify P2's cognitive bound as 1 if he assigns probability 0 to the event that P1 is cognitive.

With this in mind, we only identify P2's level of cognition if his behavior is consistent with rationality and assigning probability $p \in \{0, 1\}$ to the even that P1 is cognitive. We identify P2 as 1-cognitive if we can take $p = 0$, i.e., if P2's behavior is consistent with rationality and belief of P1 lacks cognition. We identify P2 as 2-cognitive if we cannot take $p = 0$ but can take $p = 1$, i.e., if P2's behavior is consistent with rationality and belief of P1's cognition, but inconsistent with rationality and belief that P1 lacks cognition. This choice limits our ability to rationalize the data.

To sum up, we have used this example to illustrate how we can separately identify the cognitive and rationality bounds. We identify these bounds in a way that gives a conservative estimate of the gap. In this two-player example, we can only identify the gap up to two levels of reasoning. The main paper studies a four-player game and experiment. This allows us to identify the gap up to four levels of reasoning. The next section elaborates on the more general identification strategy.

3 Identification

Figures 3.1a-3.1b describe two games, G and G_*, from Kneeland (2015). Notice that P1's and P2's payoff matrices are as in the example. However, now there is a *ring structure* to the game: Player 1's (P1's) payoffs depend only on the behavior of Player 4 (P4). Player 2's (P2's) payoffs depend only on the behavior of P1. Likewise, Player 3's (P3's) payoffs depend only on the behavior of P2 and Player 4's (P4's) payoffs depend only on P3's behavior.

Each of these games is dominance solvable. This will allow us to identify reasoning about rationality. In particular, we argued that $(\mathrm{a}, \mathrm{c}_*)$ is the unique rational strategy profile for P1. This corresponds to the fact that P1 has a dominant strategy to play $(\mathrm{a}, \mathrm{c}_*)$. Under iterated dominance, P2 would play $(\mathrm{a}, \mathrm{b}_*)$, P3 would play $(\mathrm{b}, \mathrm{a}_*)$, and P4 would play $(\mathrm{a}, \mathrm{c}_*)$.

Each subject plays both games (G and G_*) in each of the player roles (P1, P2, P3, and P4). As such, we observe each subject's behavior across eight games. We assume that each subject is rational (and cognitive). Thus, we can use the subjects' behavior across both the games and the player roles to provide a lower-bound on reasoning about cognition and an upper-bound on reasoning about rationality. This provides us with a conservative estimate (i.e., an underestimate) of the gap between cognition and rationality.

For the purpose of illustration, we begin our discussion of the identification by assuming that we only observe the subjects play the games G and G_* in a single player role. See Sections 3.1-3.2

		P4		
		a	b	c
P1	a	12	16	14
	b	8	12	10
	c	6	10	8

		P1		
		a	b	c
P2	a	20	14	8
	b	16	2	18
	c	0	16	16

		P2		
		a	b	c
P3	a	14	18	4
	b	20	8	14
	c	0	16	18

		P3		
		a	b	c
P4	a	8	20	12
	b	0	8	16
	c	18	12	6

(a) Figure G

		P4		
		a_*	b_*	c_*
P1	a_*	8	12	10
	b_*	6	10	8
	c_*	12	16	14

		P1		
		a_*	b_*	c_*
P2	a_*	20	14	8
	b_*	16	2	18
	c_*	0	16	16

		P2		
		a_*	b_*	c_*
P3	a_*	14	18	4
	b_*	20	8	14
	c_*	0	16	18

		P3		
		a_*	b_*	c_*
P4	a_*	8	20	12
	b_*	0	8	16
	c_*	18	12	6

(b) Figure G_*

Figure 3.1: Kneeland's (2015) Ring Game

below. Then, in Section 3.3, we exploit the fact that we observe behavior across all player roles to provide a tighter identification of the bounds.

3.1 Behavior in the Role of P2

We identify the cognitive bound as the minimum level of cognition consistent with observed behavior. With this in mind, we identify a subject as having a *1-cognitive bound* (or a *cognitive bound of 1*) if his behavior is consistent with rationality and belief that P1 lacks cognition. To understand what this involves, return to Section 2. There, we explained that, if P1 lacks cognition, then P1's behavior cannot depend on details of the game. Thus, if P2 believes that P1 lacks cognition, then P2 has the same belief about P1's play in both G and G_*. With this, if an action is a best response in G, then the associated action is also a best response in G_*. So, any constant action profile—i.e., $(\mathrm{a}, \mathrm{a}_*)$, $(\mathrm{b}, \mathrm{b}_*)$, or $(\mathrm{c}, \mathrm{c}_*)$—can be identified as having a cognitive bound of 1.

We assume that, within a given game, P2 is not indifferent between any two actions.[6] This implies that, if P2 has the same belief in both G and G_*, then his behavior in the two games must be the same. Thus, the non-constant action profiles cannot be identified as having a 1-cognitive bound. This is the basic idea of how we identify behavior consistent with a cognitive bound of at least 2 (but potentially only 1-rational): Non-constant action profiles indicate that the subject reasons that P1 is cognitive. As such, the non-constant action profiles will correspond to belief about cognition. But at the same time, they will be inconsistent with rationality and belief of rationality, provide they are

To be more precise about this idea, suppose that we observe the subject play $(\mathrm{a}, \mathrm{c}_*)$. As we saw in Section 2, this behavior is inconsistent with rationality and belief of rationality, but consistent with rationality and belief of cognition. In particular, it is a best response for this subject, if he assigns probability .8 to the rational strategy $(\mathrm{a}, \mathrm{c}_*)$ and probability .2 to lucky-6 strategy $(\mathrm{b}, \mathrm{c}_*)$.

---

[6] Kneeland's (2015) Footnote 20 points out that the data from this experiment supports this assumption.

Notice that, under this belief, he assigns probability .8 to a cognitive and rational strategy profile and probability .2 to a cognitive but irrational strategy profile. Thus, we think of this belief as one that exhibits "belief of cognition."

In this example, P2's belief can be decomposed into two parts: a cognitively-rational part and a cognitively-irrational part. The cognitively-rational part is determined by reasoning about rationality. The cognitively-irrational part is a belief that satisfies a certain invariance property, which we now describe.

To understand the invariance property, recall that, if P1 is cognitive, then she has a theory about how to play the game. That theory ought to translate across G and G_*—after all, the two games only differ in their labeling of strategies. So, she should view the strategy a in G as she views the strategy c_* in G_*; she should view the strategy b as she views the strategy a_*; and she should view the strategy c as she views the strategy b_*. Now consider a belief about cognitively-irrational behavior. This will consist of probability distributions $(\rho_{\mathrm{a}}, \rho_{\mathrm{b}}, \rho_{\mathrm{c}})$ on $(\mathrm{a}, \mathrm{b}, \mathrm{c})$ and $(\rho_{\mathrm{a}_*}, \rho_{\mathrm{b}_*}, \rho_{\mathrm{c}_*})$ on $(\mathrm{a}_*, \mathrm{b}_*, \mathrm{c}_*)$ that satisfy the following invariance property:

$$\rho_{\mathrm{a}} = \rho_{\mathrm{c}_*}, \qquad \rho_{\mathrm{b}} = \rho_{\mathrm{a}_*}, \qquad \text{and,} \qquad \rho_{\mathrm{c}} = \rho_{\mathrm{b}_*}.$$

We refer to such probabilities as *invariant distributions.*

At a broader level, if P2 believes that P1 is cognitive, then his belief can be decomposed into a cognitively-rational part and a cognitively-irrational part. Thus, he has beliefs $(\Pr(\mathrm{a}), \Pr(\mathrm{b}), \Pr(\mathrm{c}))$ in G and beliefs $(\Pr(\mathrm{a}_*), \Pr(\mathrm{b}_*), \Pr(\mathrm{c}_*))$ in G_* so that, there exists $p \in [0, 1]$ and invariant distributions $(\rho_{\mathrm{a}}, \rho_{\mathrm{b}}, \rho_{\mathrm{c}})$ and $(\rho_{\mathrm{a}_*}, \rho_{\mathrm{b}_*}, \rho_{\mathrm{c}_*})$ so that the following property is satisfied:

$$\begin{array}{ll} \Pr(\mathrm{a}) = p + (1-p)\rho_{\mathrm{a}} & \qquad \Pr(\mathrm{a}_*) = (1-p)\rho_{\mathrm{a}_*} = (1-p)\rho_{\mathrm{b}} \\ \Pr(\mathrm{b}) = (1-p)\rho_{\mathrm{b}} & \qquad \Pr(\mathrm{b}_*) = (1-p)\rho_{\mathrm{b}_*} = (1-p)\rho_{\mathrm{c}} \\ \Pr(\mathrm{c}) = (1-p)\rho_{\mathrm{c}} & \qquad \Pr(\mathrm{c}_*) = p + (1-p)\rho_{\mathrm{c}_*} = p + (1-p)\rho_{\mathrm{a}}. \end{array}$$

We call a belief that satisfies this property a *1-cognitive belief.* If an action profile is a best response to a 1-cognitive belief, then the behavior is consistent with rationality and belief that P1 is cognitive. Thus, the behavior is consistent with a *2-cognitive bound* (or, potentially, a higher cognitive bound).

The middle column of Table 3.1 describes the strategy profiles for P2 that are consistent with rationality and belief of cognition. There are six such strategy profiles. Five of those profiles are non-constant. This fits with the earlier observation that non-constant behavior is indicative of rationality and belief of cognition. However, notice, it is not the case that *all* non-constant action profiles are consistent with rationality and belief of cognition. In particular, $(\mathrm{c}, \mathrm{b}_*)$ is not a best response to any 1-cognitive belief.

Table 3.1 also highlights the fact that the constant action profile $(\mathrm{a}, \mathrm{a}_*)$ is consistent with both (i) rationality and belief of cognition, and (ii) rationality and belief of lack of cognition. To see why, observe that a is a best response if the subject assigns equal probability to each of P1's actions. This equal probability belief can reflect both belief in cognition (in so far as it satisfies the invariance property) and belief in lack of cognition (in so far as the beliefs are constant across games). That

Rationality and		
Belief of Rationality	Belief of Cognition	Belief of Lack of Cognition
$(\mathrm{a}, \mathrm{b}_*)$	$(\mathrm{a}, \mathrm{a}_*)$, $(\mathrm{a}, \mathrm{b}_*)$, $(\mathrm{a}, \mathrm{c}_*)$	$(\mathrm{a}, \mathrm{a}_*)$, $(\mathrm{b}, \mathrm{b}_*)$, $(\mathrm{c}, \mathrm{c}_*)$
	$(\mathrm{b}, \mathrm{a}_*)$, $(\mathrm{b}, \mathrm{c}_*)$, $(\mathrm{c}, \mathrm{a}_*)$	

Table 3.1: P2's Behavior

said, the constant action profiles $(\mathrm{b}, \mathrm{b}_*)$ and $(\mathrm{c}, \mathrm{c}_*)$ are *only* consistent with rationality and belief in lack of cognition.

Finally, Table 3.1 points out that the only non-constant action profile consistent with rationality and belief of rationality is $(\mathrm{a}, \mathrm{b}_*)$. This follows from our discussion in Section 2. (Observe that a belief that assigns probability 1 to rationality is a 1-cognitive belief.)

To summarize: There are nine potentially observed action profiles. The strategy $(\mathrm{a}, \mathrm{b}_*)$ is identified as both having a cognitive bound and a rationality bound of (at least) 2. The three constant action profiles are identified as both having a cognitive bound and a rationality bound of 1. These four action profiles exhibit no gap between the cognitive and rationality bounds. By contrast, there are four non-constant action profiles that are identified with a cognitive bound of (at least) 2 and a rationality bound of 1. Those four profiles exhibit a gap between the cognitive and rationality bounds. There is one profile that is inconsistent with our classification.

3.2 Behavior in the Role of P3

The behavior of P3 will allow us to distinguish the third level of reasoning (about cognition and/or rationality) from the second level of reasoning. We begin by focusing on two extreme cases, where there is no identified gap between the cognitive and rationality bounds.

First, suppose that the subject forms her belief by reasoning that P2 is "rational and believes that P1 is rational." In this case, the subject believes that P2 plays the action profile $(\mathrm{a}, \mathrm{b}_*)$. Thus, she rationally plays $(\mathrm{b}, \mathrm{a}_*)$.

Second, suppose that the subject instead forms her belief by reasoning that P2 is "rational and believes that P1 lacks cognition." In that case, she reasons that P2 plays a constant action profile—i.e., P2 plays either $(\mathrm{a}, \mathrm{a}_*)$, $(\mathrm{b}, \mathrm{b}_*)$, or $(\mathrm{c}, \mathrm{c}_*)$. This implies that she has the same belief in G and G_* about P2's play. With this, if an action is a best response in G, then the associated action is also a best response in G_*. So, any constant action profile can be identified as having a *2-cognitive bound* (or a *cognitive bound of 2*). Because we assume that, within a given game, a subject is not indifferent between any two actions, this is the only behavior that will be identified as having a 2-cognitive bound.

In these two cases, there is no gap between the rationality and the cognitive bounds. If we observe $(\mathrm{b}, \mathrm{a}_*)$, then the subject is identified as having a both a rationality and cognitive bound of 3. If we observe a constant action profile, then the subject is identified as having both a rationality and cognitive bound of 2. We will, again, be able to identify behavior that is 3-cognitive but is not 3-rational, by observing non-constant action profiles distinct from $(\mathrm{b}, \mathrm{a}_*)$. We now explain how we do so.

If P3 is (at least) 3-rational, then she is rational and believes

$$[\text{Rat+B(Rat)}] \equiv \text{“P2 is rational and believes that P1 is rational.”}$$

We will be interested in departures from this belief—departures that still maintain cognition and belief of cognition. One possibility is the case where P3 believes

$$[\text{Rat+B(Cog)}] \equiv \text{“P2 is rational and believes that P1 is cognitive.”}$$

A second possibility is the case where P3 believes

$$[\text{Cog+B(Cog)}] \equiv \text{“P2 is cognitively irrational and believes that P1 is cognitive.”}$$

We will consider both types of departures. We first explain what each individually involves.

If P3 believes [Rat+B(Cog)], then P3 must assign probability 1 to the six strategy profiles consistent with P2 being rational and having a cognitive bound of 1. These are the profiles in the middle column of Table 3.1. Write μ for a distribution on $\{\mathrm{a}, \mathrm{b}, \mathrm{c}\} \times \{\mathrm{a}_*, \mathrm{b}_*, \mathrm{c}_*\}$ that assigns probability one to these six strategy profiles. We call any such μ a *cognitively rational distribution.*

If, on the other hand, P3 believes [Cog+B(Cog)], then she assigns probability one to cognitively irrational behavior. As such, her beliefs satisfy the invariance property. However, now the invariance property involves a different requirement on beliefs, since the labeling of P2's actions across the two games are the same.[7] Thus, for P3, the invariance property consists of probability distributions $(\rho_{\mathrm{a}}, \rho_{\mathrm{b}}, \rho_{\mathrm{c}})$ on $(\mathrm{a}, \mathrm{b}, \mathrm{c})$ and $(\rho_{\mathrm{a}_*}, \rho_{\mathrm{b}_*}, \rho_{\mathrm{c}_*})$ on $(\mathrm{a}_*, \mathrm{b}_*, \mathrm{c}_*)$ satisfying $\rho_{\mathrm{a}} = \rho_{\mathrm{a}_*}$, $\rho_{\mathrm{b}} = \rho_{\mathrm{b}_*}$, and $\rho_{\mathrm{c}} = \rho_{\mathrm{c}_*}$. Write ρ for the product of these invariant distribution (i.e., the associated distribution on $\{\mathrm{a}, \mathrm{b}, \mathrm{c}\} \times \{\mathrm{a}_*, \mathrm{b}_*, \mathrm{c}_*\}$). We call any such ρ a *cognitively irrational distribution.*

At a broader level, departures from belief in [Rat+B(Rat)] will involve beliefs that are a convex combinations of a cognitively rational distribution and a cognitively irrational distribution. Specifically, there will be some $p \in [0, 1]$ so that P3's beliefs on $\{\mathrm{a}, \mathrm{b}, \mathrm{c}\} \times \{\mathrm{a}_*, \mathrm{b}_*, \mathrm{c}_*\}$ can be characterized by $p\mu + (1 - p)\rho$ for some cognitively rational distribution μ and some cognitively irrational distribution ρ. We call such a belief a *2-cognitive belief.* Any action profile that is a best response to a 2-cognitive belief is consistent with rationality and belief that “P2 is cognitive and that P2 believes that P1 is cognitive.” Such an action profile is consistent with a *3-cognitive bound* (or, potentially, a higher cognitive bound).

To make this more concrete, return to the example from Section 2. We saw a specific example in which P2 was rational and believed P1 is cognitive, and played $(\mathrm{a}, \mathrm{c}_*)$. Thus, there is a cognitively rational distribution μ that assigns probability 1 to $(\mathrm{a}, \mathrm{c}_*)$. If P3 has a 2-cognitive belief that assigns probability $p = 1$ to this distribution, then P3 rationally plays $(\mathrm{b}, \mathrm{c}_*)$. Suppose instead that P3 has a 2-cognitive belief that assigns probability $p = .35$ to this cognitively rational μ and

---

[7] Recall, this was not the case for the labeling of P1's actions.

assigns probability .65 to the cognitively irrational distribution ρ with $\rho((\mathrm{a}, \mathrm{a}_*)) = 1$. In that case, P3 rationally plays $(\mathrm{a}, \mathrm{b}_*)$. In both cases, she plays a non-constant action profile that is distinct from the 3-rational $(\mathrm{b}, \mathrm{a}_*)$. Thus, if we observe P3 play either of $(\mathrm{b}, \mathrm{c}_*)$ or $(\mathrm{a}, \mathrm{b}_*)$, we conclude that there is a gap between the cognitive and the rationality bounds.

Let us point to a detail of this example: The cognitively rational distribution μ assigned probability 1 to $(\mathrm{a}, \mathrm{c}_*)$. Earlier, we said that P2 would rationally play this profile if he had a cognitive belief that assigned probability .8 to P1's rationality. (There, he assigned probability .2 to the cognitively irrational lucky-6 strategy; he could not assign probability 1 to the rational strategy and, nonetheless, play $(\mathrm{a}, \mathrm{c}_*)$.) Thus, we can conceptualize this cognitively rational distribution as one that assigns probability 1 to "Rationality and .8-belief of rationality."

This points to a more general way, in which we can reconceptualize the cognitively rational distribution (and so the 2-cognitive belief). In particular, say that an event is *q-believed* if the event gets at least probability q. With this in mind, any cognitively rational distribution μ can be thought of as a distribution that assigns probability 1 to

"Rationality and q-belief of Rationality"

for some $q \in [0, 1]$. Thus, a 2-cognitive belief can be reconceptualized as a belief assigns

- probability p to "P2 is rational and q-believes P1's rationality," and
- probability $1 - p$ to "P2 is cognitively irrational,"

for some $p, q \in [0, 1]$.[8] (Table C.2 in Appendix C points out which strategy profiles are rational, as a function of p and q.) We will make use of this reconceptualization in Section 5.

To summarize: The action profile $(\mathrm{b}, \mathrm{a}_*)$ is identified as both having a cognitive and rationality bound of (at least) 3. The three constant action profiles are both identified as having a cognitive and rationality bound of 2. Other action profiles that are a best response to a 2-cognitive belief will indicate a gap between the cognitive and rationality bounds. Some of these action profiles will indicate a larger departure from 3-rationality than other profiles. For instance, in the example, the action profile $(\mathrm{a}, \mathrm{b}_*)$ was associated with P2 assigning probability .5 to "Rationality and .8-belief of rationality," whereas the action profile $(\mathrm{b}, \mathrm{c}_*)$ was associated with P2 assigning probability 1 to "Rationality and .8-belief of rationality." The former belief involves a larger departure from 3-rationality.[9]

3.3 Behavior in the Role of P4

The behavior of P4 will allow us to distinguish the fourth level of reasoning (about cognition and/or rationality) from the third level of reasoning. The approach to identification is similar to that of P3.

---

[8] As we will later see, the choice of p, q here is not unique.

[9] Of course, in this example, we have not show that these choices of p, q are tight. We will do so, when we make use of the reconceptualization.

We begin with the two extreme cases, where the cognitive and rationality bounds coincide. First, if P4 forms his belief by reasoning that P3 is

> "rational and believes that 'P2 is rational and believes that P1 is rational,' "

then P4 believes that P3 plays $(\mathrm{b}, \mathrm{a}_*)$. In that case, he rationally plays $(\mathrm{a}, \mathrm{c}_*)$. Thus, if we observe $(\mathrm{a}, \mathrm{c}_*)$, we identify the subject as having both a rationality and cognitive bound of (at least) 4. Second, if P4 forms his belief by reasoning that P3 is

> "rational and believes that 'P2 is rational and believes that P1 lacks cognition,' "

then P4 believes that P3 plays a constant action profile. So, any constant action profile can be identified as having a *3-cognitive bound.* Thus, if we observe any constant action profile, we identify the subject as having both a cognitive and rationality bound of 3. Moreover, since we assume that no subject is indifferent between any two action profiles, the constant action profiles are the only ones that are identified as having a cognitive (and rationality) bound of 3.

We will again be able to identify behavior that is 4-cognitive but not 4-rational by observing non-constant action profiles distinct from $(\mathrm{a}, \mathrm{c}_*)$. Similar to the approach taken above, there are two departures from 4-rationality, which also maintain 4-cognition. One possibility is that P4 believes

> "P3 is rational and believes that 'P2 is cognitive and believes that P1 is cognitive.' "

A second possibility is the case where P4 believes

> "P3 is cognitively irrational and believes that 'P2 is cognitive and believes that P1 is cognitive.' "

The first of these corresponds to a cognitively rational distribution μ on $\{\mathrm{a}, \mathrm{b}, \mathrm{c}\} \times \{\mathrm{a}_*, \mathrm{b}_*, \mathrm{c}_*\}$. This is now a distribution over all the strategies consistent with P3 being rational and having a cognitive bound of (at least) 3. The second of these corresponds to a cognitively irrational distribution ρ on $\{\mathrm{a}, \mathrm{b}, \mathrm{c}\} \times \{\mathrm{a}_*, \mathrm{b}_*, \mathrm{c}_*\}$. The cognitively irrational distribution satisfies the same invariance property as required for P3.

Any *3-cognitive belief* can be characterized by a distribution $p\mu + (1-p)\rho$, for some cognitively rational distribution μ, some cognitively irrational distribution ρ, and some $p \in [0, 1]$. Any action profile that is a best response to a 3-cognitive belief is consistent with rationality and belief that "P3 is cognitive and believes 'P2 is cognitive and believes that P1 is cognitive.' " Thus, it is consistent with a 4-*cognitive bound.*

Much as in the case of a 2-cognitive belief, we can reconceptualize a 3-cognitive belief as follows: P3 assigns probability

- p to "P3 is rational and q-believes 'P2 is rational and r-believes that P1 is rational,' " and
- $1-p$ to "P3 is cognitively irrational,"

for some $p, q, r \in [0, 1]$. We make use of this reconceptualizaiton in Section 5.

3.4 Identification Strategy

In the experiment, we observe each subject play in each of the player roles. Thus, an observation will consist of some $x = (x(1), \ldots, x(4))$ where $x(j) \in \{\mathrm{a}, \mathrm{b}, \mathrm{c}\} \times \{\mathrm{a}_*, \mathrm{b}_*, \mathrm{c}_*\}$ is the behavior observed in the role of Pj. We assume that the subjects' cognitive bound and rationality bound are constant across the player roles. (We had previously implicitly assumed they were constant across G and G_*.) This will allow us to use behavior across player roles to identify the cognitive bound and the rationality bound. We next explain how this is done.

Identifying the Cognitive Bound To illustrate how we identify the cognitive bound, refer to Table 3.2. It provides four examples of (potential) observed behavior.

Begin with Observation 1. In this case, the subject's behavior is consistent with rationality and, so, cognition. Notice that $x(2) = (\mathrm{a}, \mathrm{a}_*)$ is consistent with both P2 believing in "cognition" and in "lack of cognition." Since we are interested in identifying the minimum bound consistent with observed behavior, we might be tempted to identify this subject as having a cognitive bound of level 1. However, $x(3) = (\mathrm{b}, \mathrm{c}_*)$ is inconsistent with belief in lack of cognition. It is instead consistent with 3-cognition and so we would identify the subject as being 3-cognitive. Constrast this with Observation 2. It differs from Observation 1 only in $x(2)$. Now, the observed $x(2) = (\mathrm{b}, \mathrm{b}_*)$ is inconsistent with belief of cognition—it is only consistent with belief in lack of cognition. On the other hand, $x(3)$ is still only consistent with at least 3 levels of cognition. Thus, we would fail to classify this subject.

	P1	P2	P3	P4
Observation 1	$(\mathrm{a}, \mathrm{c}_*)$	$(\mathrm{a}, \mathrm{a}_*)$	$(\mathrm{b}, \mathrm{c}_*)$	$(\mathrm{b}, \mathrm{b}_*)$
Observation 2	$(\mathrm{a}, \mathrm{c}_*)$	$(\mathrm{b}, \mathrm{b}_*)$	$(\mathrm{b}, \mathrm{c}_*)$	$(\mathrm{b}, \mathrm{b}_*)$
Observation 3	$(\mathrm{a}, \mathrm{c}_*)$	$(\mathrm{b}, \mathrm{b}_*)$	$(\mathrm{c}, \mathrm{c}_*)$	$(\mathrm{b}, \mathrm{b}_*)$
Observation 4	$(\mathrm{a}, \mathrm{c}_*)$	$(\mathrm{b}, \mathrm{b}_*)$	$(\mathrm{a}, \mathrm{a}_*)$	$(\mathrm{b}, \mathrm{b}_*)$

Table 3.2: Identifying the Cognitive Bound

Consider instead Observation 3. This differs from Observation 2, in that now $x(3) = (\mathrm{c}, \mathrm{c}_*)$. Notice, now we would classify the subject as being 1-cognitive. In particular, notice that $x(2) = (\mathrm{b}, \mathrm{b}_*)$ is a best response given a belief that assigns probabilities $(.05, .05, .9)$ to both $(\mathrm{a}, \mathrm{b}, \mathrm{c})$ and $(\mathrm{a}_*, \mathrm{b}_*, \mathrm{c}_*)$. Moreover, $x(1)$, $x(3)$, and $x(4)$ are also a best response given these same beliefs. This belief is consistent with belief in lack of cognition.

Notice a key feature of this example. The beliefs of the subject were constant—not only across G and G_*—but also across player roles. This is important. In particular, if a subject reasons that other subjects lack cognition, then the subject's reasons that the behavior of others does not depend on the parameters of the game. Thus, she must also have the same beliefs across player roles. This restriction has bite. To see this, refer to Observation 4, which differs from Observation 3 only in $x(3)$. It is readily verified that there are no beliefs that are constant across both games and player roles, so that $x(2)$ and $x(3)$ are best responses.

More generality, we identify the cognitive bound as follows:

Identification. *Given an observation* $(x(1), x(2), x(3), x(4))$*, we assign the cognitive bound of* k *if the following hold:*

(i) *There exist distributions* $\mathrm{Pr} = (\mathrm{Pr(a)}, \mathrm{Pr(b)}, \mathrm{Pr(c)})$ *and* $\mathrm{Pr}_* = (\mathrm{Pr}_*(\mathrm{a}_*), \mathrm{Pr}_*(\mathrm{b}_*), \mathrm{Pr}_*(\mathrm{c}_*))$ *(on* $(\mathrm{a}, \mathrm{b}, \mathrm{c})$ *and* $(\mathrm{a}_*, \mathrm{b}_*, \mathrm{c}_*)$*) so that*

  (a) $\mathrm{Pr} = \mathrm{Pr}_*$ *and,*

  (b) *for each* $j = k+1, \ldots, 4$*,* $x(j)$ *is a best response under* $(\mathrm{Pr}, \mathrm{Pr}_*)$*.*

(ii) *For each* $j = 1, \ldots, k$*,* $x(j)$ *is a best response under some* $(j-1)$*-cognitive belief, where the* 0*-cognitive belief is any belief.*

(iii) *The strategy* $x(k)$ *is not constant across* G *and* G_**.*

An implication of criterion (i), for each $j = k+1, \ldots, 4$, $x(j)$ is constant across G and G_*. Criterion (ii) does not imply that $x(k)$ is not constant. If it were constant, we would seek to identify a lower cognitive bound.

Identifying the Rationality Bound We identify the rationality bound based off of iterated dominance. However, again, we will use the subject's behavior across player roles to provide this bound. For instance, refer to Observation 5 in Table 3.3. If we focus on the observed subject's behavior in the role of P4, then we would use the fact that $x(4)$ survives four rounds of iterated dominance to conclude that the subject's rationality bound is 4. This would, in particular, imply that the subject assigns probability one to the event that "P3 is rational." However, the subject's behavior in the role of P2, namely $x(2)$, does not survive two rounds of iterated dominance. As such, it is inconsistent with a rational subject that assigns probability one to the event that "P1 is rational." Thus, we identify the subject's rationality bound as 1.

	P1	P2	P3	P4
Observation 5	$(\mathrm{a}, \mathrm{c}_*)$	$(\mathrm{b}, \mathrm{a}_*)$	$(\mathrm{b}, \mathrm{c}_*)$	$(\mathrm{a}, \mathrm{c}_*)$

Table 3.3: Identifying the Rationality Bound

With this in mind, we identify the rationality bound as follows:

Identification. *Given an observation* $(x(1), x(2), x(3), x(4))$*, we assign the rationality bound of* k *if the following hold:*

(i) *For all* $j = 1, \ldots, k$*,* $x(j)$ *is the strategy profile of Pj that survives* j *rounds of iterated dominance.*

(ii) *If* $k = 1, \ldots, 3$*,* $x(k+1)$ *does not survive* $k+1$ *rounds of iterated dominance.*

Cognitive Bound	Potential Observations	Subjects
1	13	**9**
2	35	**16**
3	108	**21**
4	324	**28**
Not Classified	249	**1**
Total	729	**75**

Table 4.1: Inferring the Minimum Cognition from Observed Behavior

If a subject is identified as having a rationality bound of k, then her behavior is consistent with "rationality and $(k-1)^{th}$-order belief of rationality." (See e.g., Tan and da Costa Werlang (1988), Battigalli and Siniscalchi (2002), and Brandenburger and Friedenberg (2014).) If a subject is identified as having a rationality bound of $k = 1, 2, 3$, then her behavior is inconsistent with "rationality and k^{th}-order belief of rationality." But, if a subject is identified as having a rationality bound of 4, her behavior is consistent with "rationality and common belief of rationality." This feature is an artifact of the four-player ring structure of the game.

In light of the above, we will say that an observation $(x(1), x(2), x(3), x(4))$ exhibits a *gap between cognition and rationality* if the subject is assigned a cognitive bound of k, a rationality bound of n, and $k > n$. (Note, the assignment is such that $k \geq n$ as it should be, at the conceptual level.)

4 Results: Gap Between Cognition and Rationality

We analyze the data from Kneeland's (2015) experiment. In the experiment, subjects are randomly assigned an order by which they each play the eight games in Figure 3.1. The games are presented to the subjects so that the actions across G and G_* have identical labeling. (So, for instance, in the experiment, the actions a and a_* receive the same label. This paper changes the labels only for expositional purposes.) After the subjects play all 8 games and before they have observed any behavior or outcomes, the subjects are given the opportunity to revise their earlier choices. This mitigates potential learning considerations.

In the data, 75 subjects (out of 80) choose the dominant action, viz. $(\mathrm{a}, \mathrm{c}_*)$, in the role of P1. We focus on the behavior of those subjects. The five that play a dominated action fall outside the purview of our analysis.[10]

Table 4.1 shows the identified cognitive bounds. (Table B.1 in Appendix B provides a detailed list of how the cognitive bound would be assigned to potential observations.) There are 249 observations that could potentially lead to a non-classification. Of our 75 subjects, only 1 is not

---

[10] By contrast, a subject who plays $(\mathrm{c}, \mathrm{b}_*)$ in the role of P2 is irrational, but only irrational because of the assumptions we have made about the players' beliefs. Thus, in principle, we would allow such subjects in our analysis. However, there are no subjects that behave in that way.

	Rationality Bound				
Cognitive Bound	1	2	3	≥ 4	Total
1	9	–	–	–	9
2	3	13	–	–	16
3	3	2	16	–	21
4	2	10	3	13	28
NC					1

Table 4.2: Gap Between Cognition and Rationality

classified.[11] More than 37% are classified as having a cognitive bound of 4 and 12% are classified as having a cognitive bound of 1. Recall, we identify the cognitive bound as the minimum cognitive level consistent with the observed data. The fact that we focus on the minimum cognitive bound consistent with observed data has two implications. First, subjects identified as having a cognitive bound of 4 may in fact have no cognitive bound. (Given the nature of the 4-player ring game, we cannot distinguish a bound of four from higher levels of cognition.) Second, subjects identified as having a cognitive bound of 1 may actually have a higher level of cognition. If so, their behavior should indicate a gap between reasoning about cognition versus reasoning about rationality. (After all, subjects identified with a cognitive bound of 1 do not behave in accordance with "rationality and belief of rationality," etc.) However, if their cognitive bound is identified as exactly 1, then our identification strategy does not allow us to identify a lower rationality bound.

Table 4.2 provides information about the gap between the cognitive and the rationality bounds. If there were no gap between the cognitive and rationality bounds, then all subjects would fall along the diagonal. However, we do observe off-diagonal behavior implying that there is a gap. In particular, there are 65 subjects with cognitive bound of at least 2 and, of those subjects, 23 are identified as having a gap between their cognitive and rationality bounds. Almost half of the subjects who have a rationality bound of 1 or 2 also have a higher cognitive bound. (The same is not true for those with a rationality bound of 3.) The gap appears more pronounced at higher levels of cognition.

5 Minding the Gap

Section 4 used observed behavior to show that there is a gap between the cognitive and rationality bounds. This section provides insight into why the gap occurs. Toward that end, we focus on subjects whose (identified) cognitive bound is higher than their rationality bound. We argue that such subjects can be seen as having non-degenerate beliefs in how they reason about rationality. As discussed in the introduction, this has important implications for out-of-sample predictions of observed behavior.

---

[11] That subject plays $(\mathrm{b}, \mathrm{b}_*)$ in player position P2 but plays the non-constant profile $(\mathrm{a}, \mathrm{c}_*)$ in position P4.

	$(\mathrm{a}, \mathrm{a}_*)$	$(\mathrm{a}, \mathrm{b}_*)$	$(\mathrm{a}, \mathrm{c}_*)$	$(\mathrm{b}, \mathrm{a}_*)$	$(\mathrm{b}, \mathrm{c}_*)$	$(\mathrm{c}, \mathrm{a}_*)$
Bounds	$\left(\frac{1}{11}, \frac{62}{133}\right)$	$\left(\frac{2}{7}, 1\right]$	$\left(\frac{2}{5}, \frac{7}{8}\right)$	$\left[0, \frac{2}{5}\right)$	$\left[0, \frac{5}{7}\right)$	$\left[0, \frac{29}{133}\right)$

Table 5.1: Bounds on P2's Belief of Rationality Given Cognition Level ≥ 2

5.1 Behavior in Individual Player Roles

We begin by supposing that we only observe behavior in a single player role and focus on what the behavior tells us about departures from reasoning about rationality.

Role of P2 Suppose that we observe a P2 subject play the non-constant strategy $(\mathrm{d}, \mathrm{e}_*)$. We identify the subject as having a cognitive bound of at least 2, provided this strategy profile is optimal under a 1-cognitive belief. That is, the strategy profile is optimal under a belief that satisfies the following *P2 Requirement*: there is some $p_2 \in [0,1]$ so that the belief assigns probability

(i) p_2 to P1 playing the rational strategy profile $(\mathrm{a}, \mathrm{c}_*)$, and

(ii) $(1 - p_2)$ to an invariant distribution.

So, for any given strategy profile associated with a cognitive bound of $m \geq 2$, we can provide upper- and lower-bounds on the probability P2 assigns to P1's rationality. These are given in Table 5.1. (Appendix C provides more detail on how this and subsequent tables were computed.)

As an illustration, suppose we observe P2 play $(\mathrm{a}, \mathrm{c}_*)$. We previously pointed out that this is rational for P2 if he assigns probability $p_2 = .8$ to the rational strategy $(\mathrm{a}, \mathrm{c}_*)$ and probability $(1 - p_2) = .2$ to the cognitively irrational lucky-6 strategy $(\mathrm{c}, \mathrm{b}_*)$. (So, he assigns probability $1 - p_2 = .2$ to an invariant distribution, which is concentrated on $(\mathrm{c}, \mathrm{b}_*)$.) While this is an example of *some* 1-cognitive belief under which $(\mathrm{a}, \mathrm{c}_*)$ is a best response, it is not the only 1-cognitive belief under which it is a best response. For each $p_2 \in (\frac{2}{5}, \frac{7}{8})$, there exists some 1-cognitive belief—associated with p and some invariant distribution—under which $(\mathrm{a}, \mathrm{c}_*)$ is a the unique best response. And, if $p_2 \notin (\frac{2}{5}, \frac{7}{8})$, then $(\mathrm{a}, \mathrm{c}_*)$ is not the unique best response under any associated 1-cognitive belief.

We can use Table 5.1 to draw inferences on a player's reasoning about rationality. For instance, if we observe P2 play $(\mathrm{a}, \mathrm{c}_*)$, we would infer that he assigns probability $p_2 > 2/5$ to P1's rationality. However, if we observe P2 play $(\mathrm{c}, \mathrm{a}_*)$, we would infer that he assigns probability $p_2 < 2/5$ to P1's rationality. Thus, the latter can be viewed as a larger departure "reasoning about rationality."

Role of P3 We can engage in an analogous exercise for observed behavior in the role of P3. In particular, suppose we observe the subject play the non-constant strategy $(\mathrm{d}, \mathrm{e}_*)$. Recall, we identify the subject as having a cognitive bound of at least 3, provided this strategy profile is optimal under a 2-cognitive belief. We pointed out that we can characterize this 2-cognitive belief

	$(\mathrm{a}, \mathrm{a}_*)$	$(\mathrm{a}, \mathrm{b}_*)$	$(\mathrm{a}, \mathrm{c}_*)$	$(\mathrm{b}, \mathrm{a}_*)$	$(\mathrm{b}, \mathrm{b}_*)$	$(\mathrm{b}, \mathrm{c}_*)$	$(\mathrm{c}, \mathrm{a}_*)$	$(\mathrm{c}, \mathrm{b}_*)$	$(\mathrm{c}, \mathrm{c}_*)$
$q_2 \in [\frac{7}{8}, 1]$	$[0, \frac{5}{8})$	NA	$(0, \frac{15}{32})$	$(0, 1]$	$[0, \frac{14}{31})$	$(0, \frac{7}{8})$	NA	NA	$[0, \frac{15}{62})$
$q_2 \in [\frac{5}{7}, \frac{7}{8})$	$[0, \frac{5}{8})$	$(0, \frac{19}{48})$	$(0, \frac{5}{8})$	$(0, 1]$	$[0, \frac{5}{6})$	$(0, 1]$	NA	NA	$[0, \frac{15}{62})$

Table 5.2: Bounds on Assigning Probability p_3 to "Rationality and q-Belief of Rationality"

as one that satisfies the following *P3 Requirement*: there is some $(p_3, q_2) \in [0, 1]^2$ so that the belief assigns probability

(i) the subject assigns probability p_3 to "P2 is rational and q_2-believes P1's rationality," and

(ii) the subject assigns probability $1 - p_3$ to "P2 is cognitively irrational."

(See the discussion on page 15.) So, for any given strategy profile associated with a cognitive bound of $m \geq 3$, we can provide an analogue to Table 5.1. Specifically, for any given q_2, we can provide lower- and upper-bound on the probability p_3 that the subject assigns to "P2 is rational and q_2-believes P1's rationality."

As an illustration, Table 5.2 provides these bounds for the cases in which $q_2 \in [\frac{7}{8}, 1]$ and $q_2 \in [\frac{5}{7}, \frac{7}{8})$. (See Table C.2 in Appendix C for the full list of bounds.) The label NA indicates not available—i.e., the associated strategy is inconsistent with the assumptions. Let us highlight one feature of this table.

Suppose we observe P3 play $(\mathrm{b}, \mathrm{c}_*)$. This behavior is inconsistent with 3-rationality. But it is a best response, if the subject assigns probability 1 to P2 playing $(\mathrm{a}, \mathrm{c}_*)$. It is also a best response if the subject assigns probability $3/4$ to P2 playing $(\mathrm{a}, \mathrm{b}_*)$ and probability $1/4$ to P2 playing $(\mathrm{c}, \mathrm{c}_*)$. As we next explain, these two different beliefs represent two different forms of reasoning about rationality—both reflected in Table 5.2.

For the first of these beliefs, observe that if P2 is rational and assigns probability $q_2 \in (\frac{2}{5}, \frac{7}{8})$ to P1's is rationality, then P2 may well play $(\mathrm{a}, \mathrm{c}_*)$. (See Table 5.1.) Thus, a belief that assigns probability 1 to P2 playing $(\mathrm{a}, \mathrm{c}_*)$ can be viewed as a belief that assigns probability $p_3 = 1$ to the event that "P2 is rational and q_2-believes P1 is rational," for $q_2 \in (\frac{2}{5}, \frac{7}{8})$. Indeed, this is what the second row of Table 5.2 indicates. For the second of these beliefs, observe that P2 may well play $(\mathrm{c}, \mathrm{c}_*)$ if he rational and believes that P1 is lacks cognition. Thus, the belief that assigns probability $\frac{3}{4} : \frac{1}{4}$ to $(\mathrm{a}, \mathrm{b}_*) : (\mathrm{c}, \mathrm{c}_*)$ can be seen as one that assigns probability (i) $p_3 = 3/4$ to "P2 is rational and 1-believes P1 is rationality," and (ii) $1 - p_3 = 1/4$ to "P2 is cognitively irrational." Indeed, this is what the first row of Table 5.2 indicates.

These two beliefs represent distinct (i.e., non-comparable) forms of reasoning about rationality. In the second, P3 assigns positive probability to P2's cognitive irrationality. But, conditional upon P2 being rational, she believes P2 is "rational and assigns probability 1 toP1 is rational." By contrast, in the first of these cases, P3 assigns zero probability to P2's cognitive irrationality. However, in this case, she assigns positive probability to "P2 is rational and assigns positive probability to P1 being irrational."

Role of P4 Suppose we observe the subject play the non-constant strategy $(\mathrm{d}, \mathrm{e}_*)$. Again, we identify the subject as having a cognitive bound of at least 4 provided the strategy is optimal under a 3-cognitive belief. Referring back to page 16, we can characterize such a 3-cognitive belief as one that satisfies the following *P4 Requirement*: there is some $(p_4, q_3, r_2) \in [0,1]^3$ so that

(i) the subject assigns probability p_4 to "P3 is rational and q_3-believes 'P2 is rational and r_2-believes that P1 is rational,'" and

(ii) the subject assigns probability $1 - p_4$ to "P3 is cognitively irrational."

Again, for any given pair (q_3, r_2), we can provide a lower- and upper-bound on the probability p_4 that the subject assigns to "P3 is rational and q_3-believes 'P2 is rational and r_2-believes that P1 is rational.'" Table C.3 in Appendix C provides these bounds, for the observed behavior in Kneeland's (2015) experiment.

5.2 Cross-Role Restrictions: The Anonymity Assumption

A subject's behavior across player roles can provide additional information about how her behavior departs from reasoning about rationality. To understand why, return to Observation 5 in Table 3.3. There, we pointed out that $x(4) = (\mathrm{a}, \mathrm{c}_*)$ was consistent with the subject playing a best response to a belief that assigns probability $p_4 = 1$ to

> "P3 is rational and 1-believes that 'P2 is rational and 1-believes that P1 is rational.'"

So, certainly $x(4)$ is consistent with the subject playing a best response to a belief that assings probability 1 to "P3 is rational." However, if the subject assigns probability 1 to the event "P3 is rational," then the same subject should be prepared to assign probability 1 to the event that "P1 is rational." However, for this subject $x(2) = (\mathrm{b}, \mathrm{a}_*)$. Referring to Table 5.1, $x(2)$ is only a best response to a belief that assigns probability $p_2 \in [0, \frac{2}{5})$ to P1's rationality. Thus, the fact that $p_2 \leq \frac{2}{5}$ implies that p_4 cannot be 1.

Implicit in the above discussion is the idea that a player's beliefs about the rationality of the other player should not depend on the label of the player. Specifically, we assume:

Assumption 5.1 (Anonymity Assumption)**.** *For all players i, j, k, ℓ:*

(i) A subject assigns probability p to the event "Pi is rational" if and only if she assigns probability p to the event "Pj is rational."

(ii) A subject assigns probability p to the event "Pi is rational and q-believes Pk is rational" if and only if she assigns probability p to the event "Pj is rational and q-believes Pℓ is rational."

We view the Anonymity Assumption (AA) as natural since, in the experiment, the subjects are anonymous and there is a certain symmetry to the nature of the game.

As suggested above, under the AA, the behavior of the subject in the role of Pj provides a bound on her reasoning about rationality in the role of Pi, for $i > j$. Specifically, the AA has the following implications for the P2, P3, and P4 Requirements above:

Proposition 5.1. *Fix a subject whose cognitive bound is at least* $k = 2, 3, 4$.

(i) $p_3 \geq p_k$.

(ii) If $k = 4$ *and* $q_3 \geq q_2$*, then* $p_3 \geq p_4$.

The Appendix provides a proof.

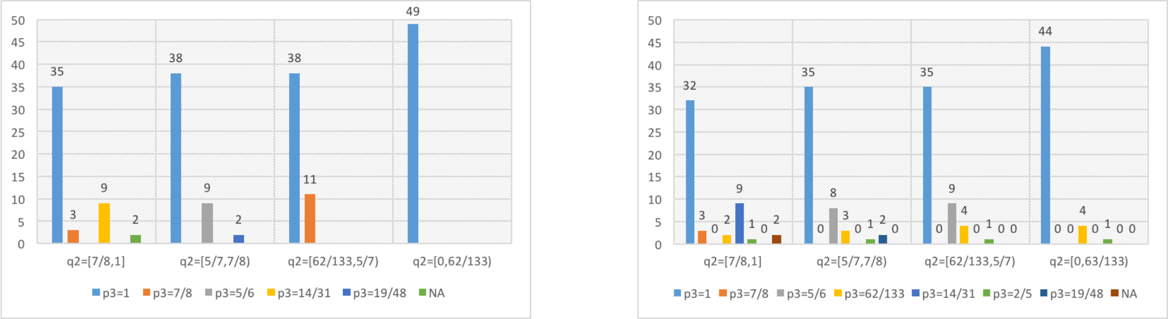

(a) Without AA

(b) With AA

Figure 5.1: Cognitive Bound ≥ 3

Figures 5.1-5.2 illustrate this role of the cross-role restrictions in the data. Figure 5.1 focuses on subjects identified as having a cognitive bound of at least 3 and their behavior in the role of P3. If we were to only observe behavior in the role of P3, we would characterize each subject by a pair (p_3, q_2) satisfying the P3 Requirement. For each (p_3, q_2), Figure 5.1a reports the number of observations that are consistent with the P3 Requirement for (p_3, q_2).[12] Figure 5.1b adds the cross-role restriction that $p_2 \geq p_3$. It points out larger departures from "reasoning about rationality" than suggested by just focusing on observed behavior in the role of P3. Figure 5.2 provides a similar illustration for subjects identified as having a cognitive bound of 4; it focuses on their behavior in the role of P4. The difference is that now the anonymity assumption imposes two cross-role restrictions: $p_2 \geq p_4$ and, if $q_2 \leq q_3$, then $p_3 \geq p_4$. The latter implies that the cross-role restrictions will depend on the specific level of q_2 assigns to a subject. Figures 5.2b-5.2c depict the restrictions for the case of $q_2 \in [{}^7\!/_8, 1]$ and $q_2 \in [{}^5\!/_7, {}^7\!/_8)$.

5.3 Non-Degenerate Beliefs about Rationality

We now argue that, when there is a gap between the cognitive and rationality bounds, subjects have non-degenerate beliefs in "reasoning about rationality." That is, the non-degenerate beliefs on "reasoning about rationality" are an important determinant of behavior.

---

[12] Technically, it reports the supremum of the p_3s so that (p_3, q_2) satisfies the P3 Requirement.

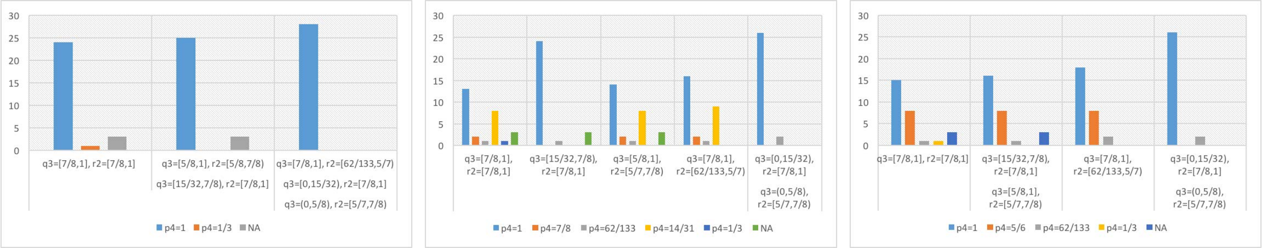

(a) Without AA

(b) With AA and $q_2 \in [\frac{7}{8}, 1]$

(c) With AA and $q_2 \in [\frac{5}{7}, \frac{7}{8})$

Figure 5.2: Cognitive Bound ≥ 4

To understand this point, first consider a subject whose (identified) cognitive bound is $k = 2$. If there is a game between the cognitive and rationality bounds, then $m = 1$. In the data, all such subjects play $(\mathrm{a}, \mathrm{c}_*)$. Thus, referring to Table 5.1, $p_2 \in (2/5, 7/8)$. As such, these subjects assign positive probability to the event that "P1 is rational" but they do not assign probability one to the event that "P1 is rational." Thus, such subjects have a non-degenerate belief about P1's rationality.

With this in mind, we focus on subjects' whose (identified) cognitive bound is $k = 3, 4$ and whose identified rationality bound is $m = 1, \ldots, k-1$—so there is a gap between the cognitive and rationality bounds. If $k = 3$, then the subject can be characterized by a $(p_2; p_3, q_2)$ that satisfies the P2 and P3 requirements. If $k = 4$, then the subject can be characterized by a $(p_2; p_3, q_2; p_4, q_3, r_2)$ that satisfies the P2, P3, and P4 requirements. We will show that, under the AA, such subjects exhibit non-degenerate beliefs about rationality, either in the role of P3 or in the role of P4.

Proposition 5.2. *Suppose that a subject is identified as having a rationality bound $m = 1, \ldots, k-1$ and a cognitive bound of k.*

(i) If $k = 3$, then $p_3 > 0$ and $(p_3, q_2) \neq (1, 1)$.

(ii) If $k = 4$, then either

  (a) $p_4 > 0$ and $(p_4, q_3, r_2) \neq (1, 1, 1)$, or

  (b) $p_3 > 0$ and $(p_3, q_2) \neq (1, 1)$.

If $p_3 > 0$ and $(p_3, q_2) \neq (1, 1)$, then the subject assigns positive probability to the event that "P2 is rational" but does not assign probability one to the event that "P2 is rational and q_2-believes P1's rationality." In this case, the subject is characterized by having non-degenerate beliefs about rationality in the role of P3. If $p_4 > 0$ and $(p_4, q_3) \neq (1, 1)$, then the subject assigns positive probability to the event that "P3 is rational" but does not assign probability one to the event that "P3 is rational and q_3-believes P2 is rational and r_2-believes P1's rationality." In this case, the subject is characterized by having non-degenerate beliefs about rationality in the role of P4.

We describe why this must hold for the case of $k = 3$. (The case of $k = 4$ is more involved and so relegated to the appendix.) First observe that, since the subject's cognitive bound is identified as $k = 3$, she must play a non-constant strategy in the role of P3. Since subjects are assumed to not be indifferent between any two strategies, it follows that $p_3 > 0$.[13] (This can also be verified from Table C.2.) Thus, $p_3 > 0$, i.e., the subject assigns positive probability to P2's rationality.

If, in the role of P3, the subject plays a strategy that does not survive 3 rounds of iterated dominance—i.e., a strategy $(\mathrm{d}, \mathrm{e}_*) \neq (\mathrm{b}, \mathrm{a}_*)$—then it is immediate that $(p_3, q_2) \neq (1, 1)$. So suppose that the subject does play $(\mathrm{b}, \mathrm{a}_*)$ in the role of P3. Then, her rationality bound must be $m = 1$ and so $p_2 < 1$. By the AA, $1 > p_2 \geq p_3$.[14] Thus, in either of these cases we conclude that the subject does not assign probability one to the event that "P2 is rational and 1-believes rationality."

6 Deliberate Choice or Errors?

In Section 4, we argued that there is a gap between the cognitive and rationality bounds. We interpreted the off-diagonal entries in Table 4.2 as evidence of such a gap. To reach this conclusion, we presumed that the off-diagonal entries were a result of deliberate choice on the part of subjects. An alternate hypothesis is that those entries do not reflect deliberate choice, but instead are a result of noise or errors. Under this alternate hypothesis, the subjects' rationality bounds are determined by their cognitive bounds, but subjects are prone to making mistakes. In that case, the off-diagonal entries reflect just those mistakes. In this section, we argue that the alternate hypothesis is incorrect.

To address this question, we estimate three models, corresponding to the two view-points on the off-diagonal entries. The first model corresponds to the identification in the previous sections; we refer to is as the Deliberate Choice model. The latter two models are variants of the alternate hypothesis; we refer to them as the Random Choice and Logistic Choice models. We argue that the data is not well explained by Random or Logistic Choice models.

In what follows, we write $i = 1, \ldots, I$ to indicate a subject in the subject pool. In our dataset, $I = 75$. In both models, we will match subjects to "types of reasoners" by maximizing the log likelihood of observing the subject's action profile.

6.1 Deliberate Choice Model

Recall from Section 5, if a subject is identified as having a cognitive bound of $k = 2, 3, 4$ then, for each $j = 2, \ldots, k$, the subject's behavior in the role of Pj must be optimal under a behavior that satisfies the Pj Requirement. So, if we identify a subject as having a cognitive bound of $k = 2$, then we can characterize the subject by some p_2 satisfying the P2 criterion. And, if we identify a subject as having a cognitive bound of $k = 3$ (resp. $k = 4$), then we can characterize the subject by some $(p_2; p_3, q_2)$ (resp. $(p_2; p_3, q_2; p_4, q_3, r_2)$) satisfying the P2 and P3 Requirements (resp. P2, P3, and P4 Requirements).

---

[13] The same would not hold for P2. This is because invariant distributions are constant for P3 but not constant for P2.

[14] Interestingly, the AA also implies that $p_2 > 0$. However, it allows p_2 to be 1 if, in the role of P2, the subject plays $(\mathrm{a}, \mathrm{b}_*)$.

To simply our analysis in this section, we impose an additional assumption about the subjects' reasoning across player roles. This assumption imposes the requirement that a subject's beliefs correspond to what he believes about the population's beliefs.[15]

Assumption 6.1 (Population Assumption)**.** $p_2 = q_2 = r_2$ *and* $p_3 = q_3$.

To understand what the Population Assumption (PA) delivers, consider a subject who has a cognitive bound of at least 3. Then, in the role of P2, she assigns probability p_2 to P1's rationality (and probability $1 - p_2$ to P1's cognitive irrationality). The PA assumption says that, in the role of P3, she acts as if she assigns probability (i) p_3 to the event that "P2 is rational and p_2-believes 'P1 is rational'" and (ii) $1 - p_3$ to P2's cognitive irrationality. Thus, the subject's beliefs about the populations' beliefs about "P1's rationality," i.e., q_3, are determined by her own beliefs about "P1's rationality," i.e., p_2. And so on.

In light of the PA, the behavior of a subject identified with a cognitive bound of $k = 4$ is determined by three parameters (p_2, p_3, p_4), where p_2 satisfies the P2 requirement, $(p_3, q_2) = (p_3, p_2)$ satisfies the P3 requirement, and $(p_4, q_3, r_2) = (p_4, p_3, p_2)$ satisfies the P4 requirement. Likewise, the behavior of a subject identified with a cognitive bound of $k = 3$ is determined by two parameters (p_2, p_3), where p_2 satisfies the P2 requirement and $(p_3, q_2) = (p_3, p_2)$ satisfies the P3 requirement.

In the Deliberate Choice model, we will think of a *type* as reflecting both a level of cognition and beliefs about rationality. As such, a type can be characterized by a triple in $([0, 1] \cup \{\mathrm{cb}\})^3$. A type with a cognition bound of 4 corresponds to some (p_3, p_2, p_1) that satisfies the P2-P3-P4 Requirements. A type with cognitive bound of 3 corresponds to some (p_3, p_2, cb), where (p_3, p_2) satisfy the P2-P3 Requirements and cb indicates that there is a cognitive bound at level 4. And so on. These triples must satisfy the AA. By Proposition 5.1, this holds if and only if $p_2 \geq \max\{p_3, p_4\}$.

An observation consists of an action profile for Subject i, i.e., some $x_i = (x_i(1), \ldots, x_i(4))$, where $x_i(j) \in \{\mathrm{a, b, c}\} \times \{\mathrm{a}_*, \mathrm{b}_*, \mathrm{c}_*\}$ denotes the behavior in the role of Pj. For each type t^k, write $A^k \subseteq (\{\mathrm{a, b, c}\} \times \{\mathrm{a}_*, \mathrm{b}_*, \mathrm{c}_*\})^4$ for the set of strategy profiles that are a best response for type t^k. Informally, A^k is the set of strategies that can be played by type t^k.

Write $\mathbb{T}^{\mathrm{D}}$ for the set of all types and observe that this set is uncountable. However, many types are essentially equivalent. For instance, the set of strategies that are a best response for $(1, 1, 1)$ is exactly the set of strategies that are a best response for $(.9, .9, .9)$. So, the types in $\mathbb{T}^{\mathrm{D}}$ can be partitioned into a finite number of subsets, so that types t^k, t^l are in the same partition member if and only if $A^k = A^l$. With this in mind, it suffices to restrict attention to finite subsets of $\mathbb{T}^{\mathrm{D}}$. With this in mind, we will use the phrase 'type' as a label for an equivalence class. We look for the minimal set of types that rationalize the data.

Write $\mathcal{M}^{\mathrm{D}} = (T^{\mathrm{D}}, \pi, \boldsymbol{\varepsilon})$ for an (**econometric**) **Deliberate Choice model**. The econometric model has three components. First, $T^{\mathrm{D}} \subseteq \mathbb{T}^{\mathrm{D}}$ is a finite subset of types in $\mathbb{T}^{\mathrm{D}}$. Second, π is a probability distribution over T^{D}; so, $\pi(t^k)$ indicates the probability that a subject is of type t^k. Third, $\boldsymbol{\varepsilon} = (\varepsilon^k)_{i=1}^{|T^{\mathrm{D}}|}$ is type-specific noise.

---

[15] We will show that, under this assumption on beliefs, deliberate behavior outperforms the hypothesis of mistakes. As a consequence, the same will be true absent this assumption on beliefs.

Under Deliberate Choice, the probability of observing x_i depends on the distribution of types (i.e., π) and the likelihood that each type t^k plays the action profile x_i. In turn, this likelihood depends on the set of strategies that can be played by type t^k, A^k; we assume that each strategy in A^k is equally likely to be played by type t^k. The econometric model $\mathcal{M}^D$ allows for type-specific noise. In particular, we assume that type t^k follows the predicted strategies A^k with probability $1-\varepsilon^k$ and deviates from the predicted strategies with probability ε^k. When type t^k deviates from the predicted strategies, type t^k plays each strategy in $A \backslash A^k$ with equal probability.

With this, the probability that observation x_i was generated by type t^k given type-specific noise ε^k is

$$p(x_i, \varepsilon^k | t^k) = \begin{cases} \frac{1-\varepsilon^k}{|A^k|} & \text{if } x_i \in A^k \\ \frac{\varepsilon^k}{|A \backslash A^k|} & \text{if } x_i \notin A^k. \end{cases}$$

The likelihood of observing x_i in model $\mathcal{M}^{\mathrm{D}}$ is then

$$\mathcal{L}_i(x_i; \mathcal{M}^{\mathrm{D}}) = \sum_{t^k \in T} \pi(t_k) p(x_i, \varepsilon^k | t^k).$$

The aggregate log-likelihood of observing the experimental data, $\boldsymbol{x} = (x_i)_{i=1}^I$, is

$$\ln \mathcal{L}(\boldsymbol{x}; \mathcal{M}^{\mathrm{D}}) = \sum_{i \in I} \ln \mathcal{L}_i(x_i, \boldsymbol{\varepsilon}, \mathcal{M}^{\mathrm{D}}).$$

We can always maximize $\ln \mathcal{L}(\boldsymbol{x}; \mathcal{M}^{\mathrm{D}})$ by letting $T^{\mathrm{D}} = \mathbb{T}^{\mathrm{D}}$. However, to limit overfitting of the data we choose amongst models $\mathcal{M}^{\mathrm{D}} = (T^{\mathrm{D}}, \pi, \boldsymbol{\varepsilon})$ by penalizing a model for having more types. We do so by using the Bayesian Information Criterion (BIC). Specifically, for a given model $\mathcal{M}^{\mathrm{D}}$, the BIC is given by

$$\mathrm{BIC}(\mathcal{M}^{\mathrm{D}}) = -2 \ln \hat{\mathcal{L}} + f \ln(I),$$

where $f = 2 \cdot |T^{\mathrm{D}}| - 1$ because each additional type in the model adds an additional 2 parameters, $(\pi(t^k), \varepsilon^k)$. We choose $\hat{\mathcal{M}}^{\mathrm{D}}$ to maximize $\mathrm{BIC}(\mathcal{M}^{\mathrm{D}})$.

6.2 Random and Logistic Choice Models

The Random and Logistic Choice models both take as given that a subject's rationality bound necessarily coincides with her cognitive bound. It interprets what appears to be a gap between cognition and rationality as an artifact of errors. For instance, consider a subject who *actually* has a cognitive bound of 3. Under these models, the subject necessarily also has a rationality bound of 3. Thus, modulo trembles, the subject would play according to iterated dominance in the roles of P1, P2, and P3; the subject would randomize in the role of P4. In these models, the subject can play $(\mathrm{a}, \mathrm{c}_*)$ in the role of P2—but only if she trembles. If these models are correct, the Deliberate Choice model would misidentify this behavior as reflecting a gap between cognition and rationality.

The Random and Logistic Choice models differ in how trembles are modelled. In the Random

Choice model, when a subject trembles in a given game, she plays the remaining (two) actions with equal probability. In the Logistic Choice model, the probability with which each action is played in a given game depends on the logistic best response function. Thus, when a subject makes an error, she is more likely to play the action that gives a higher expected payoff.

Write $\mathcal{M}^{\mathrm{R}} = (T^{\mathrm{R}}, \pi, \boldsymbol{\varepsilon})$ for a **Random Choice model** and $\mathcal{M}^{\mathrm{L}} = (T^{\mathrm{L}}, \pi, \boldsymbol{\lambda})$ for a **Logistic Choice model**. These models have three components. First,

$$T^{\mathrm{R}} = T^{\mathrm{L}} = \{c^1, c^2, c^3, c^4\}$$

is the set of types. In each of the models, type c^k indicates a subject who has cognitive bound k. In each of the models, π is a probability distribution on types $T^{\mathrm{R}} = T^{\mathrm{L}}$. The two models differ in how trembles are modelled. In the Random Choice model, type-specific trembles are given by $\boldsymbol{\varepsilon} = (\varepsilon^k)_{k=1}^4 \in (0,1)^4$. In the Logistic Choice model, type-specific trembles are given by $\boldsymbol{\lambda} = (\lambda^k)_{k=1}^4 \in \mathbb{R}_+^4$. Below we explain how these different trembles affect the likelihood of playing a given action.

In both of these models, the lack of a cognitive bound cannot be inferred by the researcher. (Now the choice of a non-constant action profile can simply reflect trembles.) However, the cognitive bound will influence the likelihood of observing any given strategy profile $x_i = (x_i(1), \ldots, x_i(4))$. Actions that are consistent with iterated dominance are more likely to have resulted from a reasoned choice.

We next proceed to explain how both the cognitive bound and trembles influence the likelihood of observing a given strategy profile, in each of the two models. To do so, it will be convenient to introduce notation. Write $x^{\mathrm{rat}}(j)$ for the iteratively undominated strategy in the role of Pj. (So, $x^{\mathrm{rat}}(1)$ is $(\mathrm{a}, \mathrm{c}_*)$, $x^{\mathrm{rat}}(2)$ is $(\mathrm{a}, \mathrm{b}_*)$, etc.) Let $n[x_i, j]$ be the number of coordinates for which $x_i(j)$ and $x^{\mathrm{rat}}(j)$ agree. So if $x_i(j) = x^{\mathrm{rat}}(j)$, then $n[x_i, j] = 2$; if $x_i(j)$ and $x^{\mathrm{rat}}(j)$ specify the same behavior in G but not G_*, then $n[x_i, j] = 1$. And so on. Write $x_g^{\mathrm{rat}}(j)$ for action specified by $x^{\mathrm{rat}}(j)$ for the game $g \in \{G, G_*\}$.

Random Choice Model Consider a subject whose cognitive bound is $k = 3$. In the role of P4, the subject randomizes equally amongst all actions. Thus, the probability of observing $x_i(4)$ is $(\frac{1}{3})^2$. In the role of P3, the subject plays the iteratively undominated strategy $(\mathrm{b}, \mathrm{a}_*)$ up to trembles. The Random Choice model assumes that trembles are independent of the payoffs, the game, and the player role. So, in the role of P3 in G (resp. G_*), the subject plays the iteratively undominated b (resp. a_*) with some probability $1 - \varepsilon$ and plays a : c (resp. b_* : c_*) with probability $\frac{\varepsilon}{2} : \frac{\varepsilon}{2}$. And similarly for her play in the role of P2 and P1.

With this, the probability of observing x_i in the model $\mathcal{M}^{\mathrm{R}}$ given a subject of type c^k is

$$p(x_i, \varepsilon^k | c^k) = \begin{cases} (\frac{1}{3})^{2(4-k)} & \text{if } \varepsilon^k = 0 \\ (\frac{1}{3})^{2(4-k)} \left(1 - \varepsilon^k\right)^{\sum_{j=1}^k n[x_i, j]} \left(\frac{\varepsilon^k}{2}\right)^{\sum_{j=1}^k (2 - n[x_i, j])} & \text{if } \varepsilon^k > 0. \end{cases}$$

Then, the likelihood of observing behavior x_i in the model $\mathcal{M}^{\mathrm{R}}$ is

$$\mathcal{L}_i(x_i, \mathcal{M}^{\mathrm{R}}) = \sum_{c^k \in T^{\mathrm{R}}} \pi_k(c^k) p(x_i, \varepsilon^k | c^k).$$

And, the aggregate log-likelihood of observing the experimental dataset $\boldsymbol{x} = (x_i)_{i=1}^{I}$ is

$$\ln \mathcal{L}(\boldsymbol{x}, \mathcal{M}^{\mathrm{R}}) = \sum_{i \in I} \ln \mathcal{L}_i(x_i, \mathcal{M}^{\mathrm{R}}).$$

We choose $\hat{\mathcal{M}}^{\mathrm{R}}$ to maximize $\ln \mathcal{L}(\boldsymbol{x}, \mathcal{M}^{\mathrm{R}})$.

Logistic Choice Model Consider again a subject whose cognitive bound is $k = 3$. As in the random choice model, the probability of observing $x_i(4)$ is still $(\frac{1}{3})^2$, since the subject randomizes in the role of P4. But, now, in the role of P3, the subject's randomized play is determined her expected payoffs, when P2 plays according to the iteratively undominated strategy profile. This randomization will be formalized by a logistic best response function.

For each action $\mathrm{d} \in \{\mathrm{a}, \mathrm{b}, \mathrm{c}\}$, write $\mathrm{EP}_G(\mathrm{d}|j)$ for the expected payoff from playing d in the role Pj=P2,P3,P4 (resp. P1), when Pj (resp. P1) expects P$(j-1)$ (resp. P4) to play the iteratively undominated action. The logistic best response function for G is given by $\sigma_G : \mathbb{R}_+ \times \{1, 2, 3, 4\} \to \triangle(\{\mathrm{a}, \mathrm{b}, \mathrm{c}\})$, where

$$\sigma_G(\lambda, j)(\mathrm{d}) = \frac{\mathrm{Exp}(\lambda \cdot \mathrm{EP}_G(\mathrm{d}|j))}{\sum_{\mathrm{e} \in \{\mathrm{a},\mathrm{b},\mathrm{c}\}} \mathrm{Exp}(\lambda \cdot \mathrm{EP}_G(\mathrm{e}|j))},$$

where λ specifies a precision parameter. Define $\mathrm{EP}_{G_*}(\mathrm{d}_*|j)$ and the logistic best response $\sigma_{G_*} : \mathbb{R}_+ \times \{1, 2, 3, 4\} \to \triangle(\{\mathrm{a}_*, \mathrm{b}_*, \mathrm{c}_*\})$ analogously.

For each player role Pj, the logistic best response function specifies a probability distribution over the action space in each game, based on the subject's expected payoff when her opponent plays the iteratively undominated action. The action that gives the highest expected payoff is played with the highest probability; the action that gives the lowest expected payoff is played with the lowest probability. The probability with which any action is played depends on a precision parameter λ. This precision parameter determines the degree of trambles. Increasing λ represents the idea that trembles are less likely than when λ is low. In particular, observe that $\lim_{\lambda \to \infty} \sigma_G(\lambda, j)(\mathrm{d})$ is 1 if $\mathrm{d} = x_G^{\mathrm{rat}}(j)$ and is 0 otherwise.

With this, the probability of observing $x_i = (x_i(1), x_i(2), x_i(3), x_i(4))$ in the model $\mathcal{M}^{\mathrm{L}}$ given a subject of type c^k is

$$p(x_i, \lambda^k | c^k) = \left(\frac{1}{3}\right)^{2(4-k)} \prod_{j=1}^{k} \left(\sigma_G(\lambda^k, j)(x_i^G(j))\right) \prod_{j=1}^{k} \left(\sigma_{G_*}(\lambda^k, j)(x_i^{G_*}(j))\right).$$

Then, the likelihood of observing behavior x_i in the model $\mathcal{M}^{\mathrm{L}}$ is

$$\mathcal{L}_i(x_i, \mathcal{M}^{\mathrm{L}}) = \sum_{c^k \in T^{\mathrm{L}}} \pi_k(c^k) p(x_i, \lambda^k | c^k).$$

And, the aggregate log-likelihood of observing the experimental dataset $\boldsymbol{x} = (x_i)_{i=1}^{I}$ is

$$\ln \mathcal{L}(\boldsymbol{x}, \mathcal{M}^{\mathrm{L}}) = \sum_{i \in I} \ln \mathcal{L}_i(x_i, \mathcal{M}^{\mathrm{L}}).$$

We choose $\hat{\mathcal{M}}^{\mathrm{L}}$ to maximize $\ln \mathcal{L}(\boldsymbol{x}, \mathcal{M}^{\mathrm{L}})$.

6.3 Model Selection

Table 6.1 compares the Deliberate Choice and Random Choice models. The first three columns contain the estimates of the Deliberate Choice model selected by the BIC criteria, the middle three columns correspond to the Random Choice model, and the last three columns correspond to the Logistic Choice model..

Deliberate Choice			**Random Choice**			**Logistic Choice**		
$(\mathbf{p_4}, \mathbf{p_3}, \mathbf{p_2})$	π	ε	$\mathbf{c^k}$	π	ε	$\mathbf{c^k}$	π	λ
$(1, 1, 1)$	.17	.009	c^4	.41	.105	c^4	.16	1.54
$(1, \frac{14}{31}, 1)$	.19	.035	c^3	.18	0	c^3	.24	∞
$(\frac{62}{133}, \frac{62}{133}, \frac{62}{133})$	.03.	.034	c^2	.20	0	c^2	.34	1.90
$(1, 1, \mathrm{cb})$	.21	.005	c^1	.21	.005	c^1	.26	∞
$(\frac{2}{5}, \frac{2}{5}, \mathrm{cb})$	.07	.015						
$(1, \mathrm{cb}, \mathrm{cb})$	.13	.009						
$(\frac{7}{8}, \mathrm{cb}, \mathrm{cb})$	.08	.014						
$(\mathrm{cb}, \mathrm{cb}, \mathrm{cb})$	.12	.003						
Neg. Log-Likelihood	296.79			356.07			364.18	
BIC	658.34			742.36			758.58	
AIC	623.58			726.14			742.36	

Table 6.1: Deliberate Choice model versus Random Choice Mode

The "best fitting" Deliberate Choice model has eight types. We label the types as a particular triple $(p_4, p_3, p_2) \in ([0,1] \cup \{\text{cb}\})^3$, where cb indicates a cognitive bound in a particular player position. As noted above, a type corresponds to a continuum of equivalent types. We provide a labeling that corresponds to the supremum of all labels. For instance, $(1, \frac{14}{31}, 1)$ corresponds to a continuum of types (p_4, p_3, p_2) with $p_4 \in (1/3, 1]$, $p_3 \in (15/62, 14/31)$, $p_2 \in [7/8, 1]$, and $p_2 \geq p_4$. This can, loosely, be thought of as a situation in which P3 assigns probability p_2 to "rationality and 1-belief of rationality," where the upper bound on p_2 is $14/31$. Thus, our choice of labels provides some indication on the size of the gap between the cognitive and rationality bounds.

The Deliberate Choice model explains the data better than both the Random Choice model and the Logistic Choice model according to model fit: The (negative) log-likelihood for the Deliberate Choice model is 296.79, 356.07 for the Random Choice model, and 364.18 for the Logistic Choice model. To account for the fact that the Deliberate Choice model has more parameters than both the Random Choice model and the Logistic Choice model, we can use the BIC and the Akaike Information Criterion (AIC). Under both these criterion, the Deliberate Choice model outperforms the other two models.

In addition, the Vuong Test can be used to test whether one model provides a significantly better fit of the data than a second.[16] When comparing the Deliberate Choice model and the Random Choice model, the Deliberate Choice model provides a significantly better fit of the data than the Random Choice model at all standard significant levels (Vuong test statistic = 2.66; p-value=.009). The same holds true for the comparison between the Deliberate Choice model and the Logistic Choice model (Vuong test statistic = 3.48; p-value=.001). The Random Choice model also provides a significantly better fit of the data than the Logistic Choice model at the 10% level (Vuong test statistic = 1.67; p-value=.095).

6.4 Individual-Level Analysis

Sections 6.1-6.2-6.3 estimated an aggregate econometric model. In particular, it analyzed the distributions of types that best fit behavior in the full dataset. It raises the question of whether any given observation can be explained by a given type (versus a non-degenerate distribution on types). Toward that end, this subsection estimates an individual-level econometric model, in which (in principle) each observation is described by a distribution of types. We show that, each observation is associated with a degenerate distribution on types and so each observation can indeed be associated with a given type. Moreover, the distribution associated with the frequency of estimated types corresponds to the estimated aggregate analysis in Section 6.3. This serves as a consistency check on our earlier results. The next subsection goes on to use these estimates to address an important model-selection question: Can a model generated by Random Choice generate the distribution of Deliberate Choice types observed in the data?

The individual-level analysis will be based on the econometric framework in Sections 6.1-6.2.

---

[16] The Vuong Test can be used to test whether one of two non-nested models provides a significantly better fit of the data. The three models here are all non-nested.

The key difference is that now, the distributions π are subject specific (now π_i) and the noise/errors are also subject specific (now ε_i). We discuss the implications for the Deliberate Choice and Random Choice estimations. (Since the Random Choice model outperforms the Logistic Choice model under the aggregate analysis, we omit discussion of the latter.)

Deliberate Choice Estimation For each subject i, we choose $(\hat{\pi}_i, \hat{\varepsilon}_i)$ to maximize $\ln \mathcal{L}_i(x_i, \mathcal{M}^{\mathrm{D}})$. As a consequence, we will estimate the subject-specific noise to be zero provide we can rationalize the subject's behavior by Deliberate Choice.[17] That is, if there exists some type in the model, viz. t^k, with $x_i \in A^k$, then $\hat{\varepsilon}_i$ must be 0. Moreover, since we have subject specific distributions, each $\hat{\pi}_i$ is degenerate on a particular type. Thus, we view the type that $\hat{\pi}_i$ is concentrated on as categorizing subject i.

We estimate the Deliberate Choice model that corresponds to the model in Section 6.3. Specifically, we focus on models associated with the same eight types that were best fitting (according to BIC) in the analysis of Section 6.3. Most subjects are fit to a type without noise. However, one subject is fit to the type $(\frac{62}{133}, \frac{62}{133}, \frac{62}{133})$ with noise.[18] In light of this, we think of this model as one in which there are effectively eight *nice types*—corresponding to the analysis in Section 6.3—and one *error type.*

Random Choice Estimation For each subject i, we now choose $(\hat{\pi}_i, \hat{\varepsilon}_i)$ to maximize $\ln \mathcal{L}_i(x_i, \mathcal{M}^{\mathrm{R}})$. Again, each $\hat{\pi}_i$ is degenerate on some type; we use that type to classify the subject.

Notice that, if x_i is consistent with 4-rationality, subject i will be assigned to the type c^4 and the error will be estimated as 0 (i.e. $\hat{\pi}_i(c^4) = 1$ and $\hat{\varepsilon}_i = 0$). Similarly, for any subject whose action profile is consistent with m-rationality. If a subject is assigned to a type with an error (i.e., with $\hat{\varepsilon}_i > 0$), then it is more likely that x_i was generated by a type with a higher cognitive bound who made a mistake versus a type with a lower cognitive bound who did not make a mistake.

Estimates Table 6.2 compares the individual analysis for the Deliberate Choice, and Random Choice model. The first four columns contain the estimates of the Deliberate Choice model and the last four columns correspond to the Random Choice model. In both cases, the column "Subject" indicates the number of subjects assigned to the type, according to the estimates $\hat{\pi}_i$. Likewise, in both cases, the column "Distribution" indicates the induced distribution of types.

For the Deliberate Choice model, the individual analysis tells the same story as the aggregate analysis. In particular, the estimated proportions of each type is very close to the distribution in the aggregate analysis. The notable exception occurs for subject classified as having a cognitive bound of 2. The aggregate analysis suggests that a slightly larger fraction of such subjects (i.e., subjects with a cognitive bound of 2) has a gap between the cognitive and rationality bounds.

---

[17] We include the noise to capture the behavior of subjects who cannot be classified according to our approach.

[18] The subject is classified as an error type because his action profile is not consistent with the Deliberate Choice model.

For the Random Choice model, the individual analysis tells a slightly different story than the aggregate analysis. In particular, the estimated proportion of types places relatively less weight on c^4 types and relativel more weight on c^3 types. This is suggestive of the fact that the Random Choice model may not be a robust model.

Deliberate Choice				**Random Choice**			
$(\mathbf{p_4}, \mathbf{p_3}, \mathbf{p_2})$	ε	**Subjects**	**Distribution**	$\mathbf{c^k}$	ε	**Subjects**	**Distribution**
$(1, 1, 1)$	0	13	.17	c^4	0	13	.17
		(3.38)	(.05)			(3.28)	(.04)
$(1, \frac{14}{31}, 1)$	0	13	.17	c^4	$\frac{1}{8}$	11	.15
		(3.27)	(.04)			(3.15)	(.04)
$(\frac{62}{133}, \frac{62}{133}, \frac{62}{133})$	0	2	.03	c^4	$\frac{2}{8}$	1	.01
		(1.38)	(.02)			(1.03)	(.01)
$(1, 1, \text{cb})$	0	16	.21	c^3	0	19	.25
		(3.55)	(.05)			(3.85)	(.05)
$(\frac{2}{5}, \frac{2}{5}, \text{cb})$	0	5	.07	c^3	$\frac{1}{6}$	1	.01
		(2.19)	(.03)			(0.97)	(.01)
$(1, \text{cb}, \text{cb})$	0	13	.17	c^2	0	15	.2
		(3.30)	(.04)			(3.50)	(.05)
$(\frac{7}{8}, \text{cb}, \text{cb})$	0	3	.04	c^1	0	15	.2
		(1.68)	(.02)			(3.63)	(.05)
$(\text{cb}, \text{cb}, \text{cb})$	0	9	.12				
		(2.80)	(.04)				
$(\frac{62}{133}, \frac{62}{133}, \frac{62}{133})$	1	1	.01				
		(0.95)	(.01)				
Neg. Log-Likelihood		153.12				258.80	

Table 6.2: Deliberate Choice model versus Random Choice Model - Individual Analysis

6.5 Simulating the Random Choice Model

To further rule out the possibility that the data was generated by the Random Choice model—and, hence, rule out the possibility that the gap between cognition and rationality was generated by noise—we simulate the (estimated) Random Choice model. We estimate the mean Deliberate Choice distribution from 1000 simulations of the (estimated) Random Choice model. This is shown in Figure 6.1. (Error bars represent standard errors.)

To better understand the approach, refer to Figure 6.1: The estimates from the simulations suggest a gap between reasoning about cognition and reasoning about rationality, with 15% of subjects being assigned to the 'gap type' $(1, \frac{14}{31}, 1)$. This is similar to the 17% of subjects assigned to the 'gap type' $(1, \frac{14}{31}, 1)$ in the actual data. Based on this fact, one might conclude that the gap

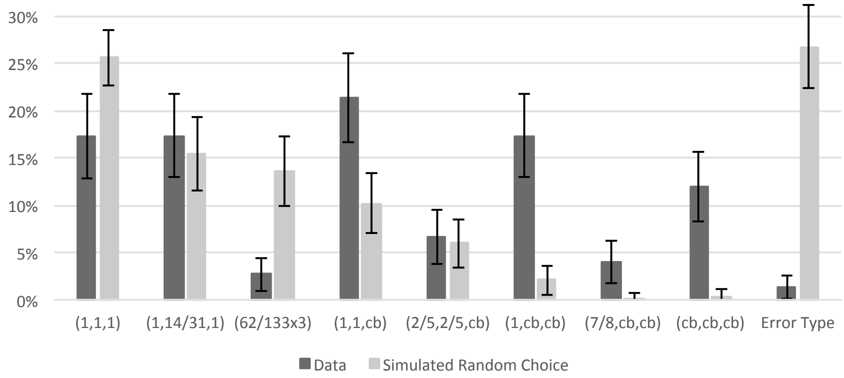

Figure 6.1: Simulated Data with the Random Choice Model

could of been generated by the Random Choice model. However, if the gap between the cognitive and rationality bounds were actually generated by the Random Choice model, there would be other observable implications. Specifically, if the Random Choice model did generate this distribution on types, it would be able to approximate the distribution of types. But, referring to Figure 6.1, it does not. The most striking difference is in the proportion of error types. The simulations predict that we should observe a very large proportion of error types, around 27%. But, in the data, error types are rare, making up only 1% of observed behavior. Second, types with lower cognitive bounds—i.e., $k = 1, 2$—disappear in the simulated data, but make up more than 30% of the actual data. Third, the type $(\frac{62}{133}, \frac{62}{133}, \frac{62}{133})$ accounts for only 3% of the actual data but makes up nearly 15% of the simulated data. These striking differences between the observed type distribution and the simulated distributions give us confidence that the identified gap cannot be driven by noise.

Appendix A Reasoning about Rationality vs. Level-k Reasoning

This paper focuses on bounded reasoning about rationality. In this appendix, we relate that concept to level-k reasoning and cognitive hierarchy reasoning. We begin by doing so in the abstract—i.e., for an arbitrary game—and then discuss how the concepts relate in Kneeland's (2015) ring game.

Bounded Reasoning About Rationality The concept of bounded reasoning about rationality has foundations in the epistemic game theory literature. We do not adopt a formal epistemic framework. Instead, we take "rationality and m^{th}-order belief of rationality" to be $(m+1)$-rationalizability (Bernheim, 1984; Pearce, 1984). This is consistent with results in Tan and da Costa Werlang (1988); Battigalli and Siniscalchi (2002); Brandenburger and Friedenberg (2014).

Say that a strategy of Ann is an R1 strategy if it is rational, i.e., if it maximizes Ann's expected payoffs given *some* belief about Bob's play. And, analogously, a strategy of Bob's is an R1 strategy if it is rational. Then, a strategy of Ann is an R2 strategy if it maximizes Ann's expected payoffs given some belief about Bob's play that assigns probability one to Bob's R1 strategies. More generally,

> A strategy of Ann (resp. Bob) is an **R(m+1) strategy** if it maximizes Ann's (resp. Bob's) expected payoffs given some belief about Bob's (resp. Ann's) play that assigns probability one to Bob's (resp. Ann's) Rm strategies.

In light of the above, we say that Ann is an **Rm reasoner**, if she chooses a strategy that is in Rm but not R(m+1). This is consistent with the definitions adopted in the main text.

Let us make two observations about Rm strategies. First, if a strategy is Rm then it is also Rn strategy for $n \leq m$. Second, a strategy is Rm if and only if it survives m-rounds of iterated strong dominance, where iterated dominance is defined according to maximal simultaneous deletion. (See Pearce, 1984.)

Level-k Reasoning This is the concept of bounded reasoning about (ir)rationality from Costa-Gomes, Crawford, and Broseta (2001). The literature begins by specifying a distribution about Bob's (resp. Ann's) play of the game, written p_b^0 (resp. p_a^0). At times, this distribution is viewed as reflecting the behavior of a player that is not cognitive. Under this interpretation, Ann (resp. Bob) is said to be an L0 reasoner if her play follows the exogenous distribution p_a^0 (resp. p_b^0). A second interpretation is that such non-cognitive L0 reasoners do not exist—rather, the distributions p_b^0 an p_a^0 reflect beliefs of cognitive L1 reasoners.

Ann is said to be an L1 reasoner if she maximizes her expected payoffs given a belief that Bob is an L0 reasoner—i.e., given the belief p_b^0. And similarly for Bob. Let p_a^1 (resp. p_b^1) be a distribution on L1 play for Ann (resp. Bob). Typically an L1 reasoner has a unique best response and so this distribution is degenerate on a particular strategy of Ann (resp. Bob). Ann is said to be an L2 reasoner if she maximizes her expected payoffs given a belief that Bob is an L1 reasoner—i.e., given the belief p_b^1.

More generally, a **k-distribution** for Ann (resp. Bob) is a distribution on the play permissible by an Lk reasoner of Ann (resp. Bob). Write p_a^k (resp. p_b^k) for such an k-distribution—i.e., that assigns probability one to Lk strategies of Ann (resp. Bob). Then:

> A strategy for Ann (resp. Bob) is an **L$(k+1)$ strategy** if it maximizes Ann's (resp. Bob's) expected payoffs given a k-distribution of Bob's (resp. Ann's).

In light of the above, Ann is an **Lk reasoner** if she chooses an Lk strategy.

Because we seek to understand the concepts at an abstract level—with applicability to any game—we have described the concept in generality. In practice, the concept of level-k reasoning is applied to games (and L0 distributions) that satisfy the following property: For each $k \geq 1$, the k-distribution is degenerate—that is, there is a unique k-distribution and that distribution assigns probability one to a particular strategy. (This property would necessarily hold in a "generic" game, provided the distributions p_a^0 and p_b^0 are chosen judiciously.)

Papers that seek to identify level-k reasoning from observed behavior often restrict attention to games (and L0 distributions) that satisfy an additional property: If s_a is an Lk strategy, r_a is an

Ln strategy, and $k \neq n$, then $s_a \neq r_a$. That is, strategies played by Lk reasoners are distinct from strategies played by all lower-order (but cognitive) reasoners.

Cognitive Hierarchy Reasoning This is the concept of bounded reasoning about (ir)rationality from Nagel (1995); Stahl and Wilson (1995); Camerer, Ho, and Chong (2004). The concepts of CH0 reasoners and CH1 reasoners are defined as L0 and L1 reasoners. However, CH2 reasoners are different from L2 reasoners: Specifically, a CH2 reasoner of Ann has a belief about Bob's strategies, viz. q_b^2, that is obtained as a convex combination of $p_a^0 : p_a^1$. The convex combination is determined by the mean level of reasoning of the population.

More generally, begin with a prior distribution on levels $k \in \mathbb{N}$ as determined by a Poisson distribution, $f(k;\tau) = {}^{e^{-\tau}\tau}/_{k!}$. The parameter τ captures the mean level k in the population. Call q_a^k a **k-cognitive-distribution** for Ann if there exists $p_a^0, \ldots, p_a^k$ distributions on Ann's strategies, so that (i) for each $j \leq k$, p_a^j assigns probability one to CHj strategies of Ann, and (ii) $q_a^k = \sum_{l=0}^{k} \left({}^{f(l;\tau)}/_{\sum_{j=0}^{k} f(j;\tau)} \right) p_a^l$. And, analogously for Bob. Then:

> A strategy for Ann (resp. Bob) is an **CH$(k+1)$ strategy** if it maximizes Ann's (resp. Bob's) expected payoffs given a k-cognitive-distribution of Bob's (resp. Ann's).

In light of the above, Ann is a **CHk reasoner** if she chooses an CHk strategy.

Connections Let us draw connections between bounded reasoning about rationality, Lk reasoning, and cognitive hierarchy reasoning.

First, observe that an L1 strategy (or CH1) is rational and, so, an R1 strategy. As a consequence, an L2 strategy is also an R2 strategy: It is a best response under a 1-distribution and a 1-distribution assigns probability one to L1—and, so, R1—strategies. More generally, for any $k \geq 1$, an Lk strategy is also an Rk strategy.[19] However, the converse does not hold. There may be strategies that are Rk but not Lk strategies. This is because there are strategies that may be a best response for Ann, even though they are not a best response to the exogenous belief p_b^0.

Second, while a CH1 strategy is an R1 strategy, a CH2 strategy need not be an R2 strategy. This is because a CH2 strategy is optimal under a distribution that assigns positive probability to strategies in the support of p_b^0 and, in turn, p_b^0 can assign positive probability to irrational strategies of Bob. In fact, for any given τ, there exists some game in which the Rk strategies are disjoint from the CHk strategies, for all $k \geq 2$.

Third, typically, both Lk reasoning and CHk reasoning involve non-degenerate beliefs about rationality (and reasoning about rationality). The 0-belief p_a^0 (resp. p_b^0) often assigns positive probability to *irrational* strategies of Ann (resp. Bob). When that is the case, the L1 (or CH1) reasoner can be interpreted as one that is rational but does not believe rationality. With this, when Ann engages in L2 reasoning, she is rational and assigns probability one to

> Bob is rational and assigns probability p to my rationality,

---

[19] In fact, if $j \geq k \geq 1$, then an Lj strategy is an Rk strategy.

	P1	P2	P3	P4
L1	$(\mathrm{a}, \mathrm{c}_*)$	$(\mathrm{a}, \mathrm{a}_*)$	$(\mathrm{b}, \mathrm{b}_*)$	$(\mathrm{a}, \mathrm{a}_*)$
L2	$(\mathrm{a}, \mathrm{c}_*)$	$(\mathrm{a}, \mathrm{b}_*)$	$(\mathrm{b}, \mathrm{b}_*)$	$(\mathrm{a}, \mathrm{a}_*)$
L3	$(\mathrm{a}, \mathrm{c}_*)$	$(\mathrm{a}, \mathrm{b}_*)$	$(\mathrm{b}, \mathrm{a}_*)$	$(\mathrm{a}, \mathrm{a}_*)$
L4	$(\mathrm{a}, \mathrm{c}_*)$	$(\mathrm{a}, \mathrm{b}_*)$	$(\mathrm{b}, \mathrm{a}_*)$	$(\mathrm{a}, \mathrm{c}_*)$

(a) Level-k Reasoning

	P1	P2	P3	P4
CH1	$(\mathrm{a}, \mathrm{c}_*)$	$(\mathrm{a}, \mathrm{a}_*)$	$(\mathrm{b}, \mathrm{b}_*)$	$(\mathrm{a}, \mathrm{a}_*)$
CH2	$(\mathrm{a}, \mathrm{c}_*)$	$(\mathrm{a}, \mathrm{a}_*)$	$(\mathrm{a}, \mathrm{a}_*)$	$(\mathrm{c}, \mathrm{c}_*)$
CH3	$(\mathrm{a}, \mathrm{c}_*)$	$(\mathrm{a}, \mathrm{a}_*)$	$(\mathrm{a}, \mathrm{a}_*)$	$(\mathrm{c}, \mathrm{c}_*)$
CH4	$(\mathrm{a}, \mathrm{c}_*)$	$(\mathrm{a}, \mathrm{a}_*)$	$(\mathrm{a}, \mathrm{a}_*)$	$(\mathrm{c}, \mathrm{c}_*)$

(b) Cognitive Hierarchy Reasoning

Table A.1: Application to the Ring Game

for some $p \in (0, 1)$. By contrast, if Ann engages in CH2 reasoning, she is rational and assigns probability $q \in (0, 1)$ to this same event. And so on.

In light of the above, we may well have an L1 (resp. CH1) strategy that is an R1 but not an R2 strategy. However, in some 'special' games, an L1 (resp. CH1) strategy may in fact be an R2 strategy. This would occur if there is another distribution $\tilde{p}_b^0 \neq p_b^0$ so that s_a is optimal under $\tilde{p}_b^0$ and $\tilde{p}_b^0$ only assigns probability one to Bob's R1 strategies. If all L1 strategies are R2, then all L2 strategies are R3 strategies. And so on. However, because CH2 reasoning assigns positive probability to p_b^0, the same need not follow for the CH strategies.

Application to the Ring Game The typical L0 distribution (and CH0 distribution) is uniform on the actions of the other player. Under this distribution, an L1 (and CH1) reasoner would play $(\mathrm{a}, \mathrm{c}_*)$ in the role of P1, $(\mathrm{a}, \mathrm{a}_*)$ in the role of P2, $(\mathrm{b}, \mathrm{b}_*)$ in the role of P3, and $(\mathrm{a}, \mathrm{a}_*)$ in the role of P4. Tables A.1a-A.1b give the behavior of the Lk and CHk reasoners (calculated using $\hat{\tau} = 1.61$ which was the median τ found in Camerer, Ho, and Chong, 2004), in the roles of each of the players.

There are two things to take note of. First, if we were to observe the behavior of an Lk $\neq$L0 reasoner, we would conclude that the subject is an Rk reasoner, whose cognitive bound is k. The subjects whose behavior indicates a gap between the cognitive and rationality bound are subjects who would not be classified as Lk reasoners, for any k. Second, if we were to observe the behavior of any CHk $\neq$CH0 player, we would conclude that the subjects' cognitive bound is 1.

Appendix B Identifying Cognitive Bounds

Cognitive Bound	P1	P2	P3	P4	Strategy	Subject
1	(a, c_*)	(a, a_*)	(a, a_*), (b, b_*), (c, c_*)	(a, a_*)	3	**7**
1	(a, c_*)	(a, a_*)	(a, a_*), (b, b_*)	(c, c_*)	2	**1**
1	(a, c_*)	(b, b_*)	(b, b_*)	(a, a_*), (b, b_*), (c, c_*)	3	**1**
1	(a, c_*)	(b, b_*)	(c, c_*)	(a, a_*), (b, b_*)	2	**0**
1	(a, c_*)	(c, c_*)	(a, a_*)	(a, a_*)	1	**0**
1	(a, c_*)	(c, c_*)	(c, c_*)	(a, a_*), (b, b_*)	2	**0**
2	(a, c_*)	(a, b_*), (a, c_*), (b, a_*), (b, c_*), (c, a_*)	(a, a_*)	(a, a_*), (c, c_*)	10	**1**
2	(a, c_*)	(a, b_*), (a, c_*), (b, a_*), (b, c_*), (c, a_*)	(b, b_*)	(a, a_*), (b, b_*), (c, c_*)	15	**15**
2	(a, c_*)	(a, b_*), (a, c_*), (b, a_*), (b, c_*), (c, a_*)	(c, c_*)	(a, a_*), (b, b_*),	10	**0**
3	(a, c_*)	(a, a_*), (a, b_*), (a, c_*), (b, a_*), (b, c_*), (c, a_*)	(a, b_*), (a, c_*), (b, a_*), (b, c_*), (c, a_*), (c, b_*)	(a, a_*), (b, b_*), (c, c_*)	108	**21**
4	(a, c_*)	(a, a_*), (a, b_*), (a, c_*), (b, a_*), (b, c_*), (c, a_*)	(a, a_*), (a, b_*), (a, c_*), (b, a_*), (b, b_*), (b, c_*), (c, a_*), (c, b_*), (c, c_*),	(a, b_*), (a, c_*), (b, a_*), (b, c_*), (c, a_*), (c, b_*)	324	**28**
NC					249	**1**
Total					729	**75**

Table B.1: Inferring the Cognitive Bound from Observed Behavior

Appendix C Intervals for Section 5.1

This appendix first describes how we calculated the intervals described in Section 5.1. It then goes on to provide the intervals for P3 and for the observed strategies of P4.

C.1 Interval Calculations

Fix an action profile $(\mathrm{d}, \mathrm{e}_*) \in \{\mathrm{a}, \mathrm{b}, \mathrm{c}\} \times \{\mathrm{a}_*, \mathrm{b}_*, \mathrm{c}_*\}$ of some player Pi, where $i = 2, 3, 4$. We consider the case where d and e_* are optimal for Pi under beliefs $\Pr = (\Pr(\mathrm{a}), \Pr(\mathrm{b}), \Pr(\mathrm{c}))$ and $\Pr_* = (\Pr(\mathrm{a}_*), \Pr(\mathrm{b}_*), \Pr(\mathrm{c}_*))$, where Pr is a belief on the strategies of P$(i+1)$ in G and $\Pr_*$ is a belief on the strategies of P$(i-1)$ in G_*. We require the beliefs of i to be $(i-1)$-*cognitive.* This corresponds to the following requirement: There are $(\rho_\mathrm{a}, \rho_\mathrm{b}, \rho_\mathrm{c}) \in [0, 1]^3$ and $(\mathrm{q}_{(\mathrm{a},\mathrm{a}_*)}, \mathrm{q}_{(\mathrm{a},\mathrm{b}_*)}, \mathrm{q}_{(\mathrm{a},\mathrm{c}_*)}, \mathrm{q}_{(\mathrm{b},\mathrm{a}_*)}, \mathrm{q}_{(\mathrm{b},\mathrm{b}_*)}, \mathrm{q}_{(\mathrm{b},\mathrm{c}_*)}, \mathrm{q}_{(\mathrm{c},\mathrm{a}_*)}, \mathrm{q}_{(\mathrm{c},\mathrm{b}_*)}, \mathrm{q}_{(\mathrm{c},\mathrm{c}_*)}) \in [0, 1]^9$ where, for each $\mathrm{d} \in \{\mathrm{a}, \mathrm{b}, \mathrm{c}\}$ and each $\mathrm{d}_* \in \{\mathrm{a}_*, \mathrm{b}_*, \mathrm{c}_*\}$ the following hold:

(i) $\Pr(\mathrm{d}) = \rho_\mathrm{d} + \mathrm{q}_{(\mathrm{d},\mathrm{a}_*)} + \mathrm{q}_{(\mathrm{d},\mathrm{b}_*)} + \mathrm{q}_{(\mathrm{d},\mathrm{c}_*)}$ and

(ii) $\Pr(\mathrm{d}_*) = \rho_\mathrm{d} + \mathrm{q}_{(\mathrm{a},\mathrm{d}_*)} + \mathrm{q}_{(\mathrm{b},\mathrm{d}_*)} + \mathrm{q}_{(\mathrm{c},\mathrm{d}_*)}$.

We will consider the case where $(\mathrm{d}, \mathrm{e}_*)$ is optimal for Pi under $(i-1)$-cognitive beliefs $(\Pr, \Pr_*)$, where these beliefs assigns probability p to a "rationality event about P$(i-1)$." Such a rationality event can be viewed as a $E \subseteq \{\mathrm{a}, \mathrm{b}, \mathrm{c}\} \times \{\mathrm{a}_*, \mathrm{b}_*, \mathrm{c}_*\}$. (Note, E need not be a product set.) The fact that the belief assigns probability p to this event has two implications. First, if $(\mathrm{d}, \mathrm{e}_*) \notin E$, then we impose the constraint that $q_{(\mathrm{d},\mathrm{e}_*)} = 0$. Second,

$$p = \mathrm{q}_{(\mathrm{a},\mathrm{a}_*)} + \mathrm{q}_{(\mathrm{a},\mathrm{b}_*)} + \mathrm{q}_{(\mathrm{a},\mathrm{c}_*)} + \mathrm{q}_{(\mathrm{b},\mathrm{a}_*)} + \mathrm{q}_{(\mathrm{b},\mathrm{b}_*)} + \mathrm{q}_{(\mathrm{b},\mathrm{c}_*)} + \mathrm{q}_{(\mathrm{c},\mathrm{a}_*)} + \mathrm{q}_{(\mathrm{c},\mathrm{b}_*)} + \mathrm{q}_{(\mathrm{c},\mathrm{c}_*)}.$$

For each $(\mathrm{d}, \mathrm{e}_*)$, write $X[(\mathrm{d}, \mathrm{e}_*)]$ (resp. $\overline{X}[(\mathrm{d}, \mathrm{e}_*)]$) for the set of p so that there exists $(i-1)$-cognitive beliefs $(\Pr, \Pr_*)$ where (i) d is a strict (resp. weak) best response under Pr; (ii) e_* is a strict (resp weak) best response under $\Pr_*$; (iii) the (relevant) event E gets probability p. Note that $X[(\mathrm{d}, \mathrm{e}_*)] \subseteq \overline{X}[(\mathrm{d}, \mathrm{e}_*)]$.

We can now describe how the lower and upper bounds are computed. We first find the supremum (resp. infimum) of $p \in \overline{X}[(\mathrm{d}, \mathrm{e}_*)]$. Notice, this is a maximization problem subject to the constraint that certain linear inequalities (or equalities) are satisfied. Thus, we can use standard techniques to compute the bounds.

At times, we find that either the supremum or infimum that is also contained in $X[(\mathrm{d}, \mathrm{e}_*)]$. But, at times, they are both not contained in $X[(\mathrm{d}, \mathrm{e}_*)]$. In those instances, we attempt to find some element of $X[(\mathrm{d}, \mathrm{e}_*)]$ (or show that one does not exist). Since $X[(\mathrm{d}, \mathrm{e}_*)] \subseteq \overline{X}[(\mathrm{d}, \mathrm{e}_*)]$, that element must lie between the upper and lower bounds. We then use the following result to complete the argument:

Claim C.1. *Suppose* $(\mathrm{d}, \mathrm{e}_*)$ *is strictly optimal under* $(i-1)$*-cognitive beliefs* $(\Pr, \Pr_*)$ *and optimal under* $(i-1)$*-cognitive beliefs* $(\Pr', \Pr'_*)$*. Further, suppose that* $(\Pr, \Pr_*)$ *assigns probability* p *to*

$E \subseteq \{\mathrm{a}, \mathrm{b}, \mathrm{c}\} \times \{\mathrm{a}_*, \mathrm{b}_*, \mathrm{c}_*\}$ *and* $(\mathrm{Pr}', \mathrm{Pr}'_*)$ *assigns probability* p' *to* $E \subseteq \{\mathrm{a}, \mathrm{b}, \mathrm{c}\} \times \{\mathrm{a}_*, \mathrm{b}_*, \mathrm{c}_*\}$*, where* $p' \geq p$ *(resp.* $p \geq p'$*). Then for any* $p'' \in (p, p')$ *(resp.* $p'' \in (p', p)$*), there exists invariant beliefs* $(\mathrm{Pr}'', \mathrm{Pr}''_*)$ *that assigns probability* p'' *to* $E \subseteq \{\mathrm{a}, \mathrm{b}, \mathrm{c}\} \times \{\mathrm{a}_*, \mathrm{b}_*, \mathrm{c}_*\}$ *and such that* $(\mathrm{d}, \mathrm{e}_*)$ *is strictly optimal under invariant beliefs* $(\mathrm{Pr}'', \mathrm{Pr}''_*)$.

Proof. Without loss of generality, suppose that $p' \geq p$ and fix $p'' \in (p, p')$. Note that, there exists $\alpha \in (0, 1)$ such that $p'' = \alpha p + (1 - \alpha)p''$. Take each $\rho''_{\mathrm{d}} = \alpha \rho_{\mathrm{d}} + (1 - \alpha)\rho'_{\mathrm{d}}$ and each $q''_{(\mathrm{d},\mathrm{e}_*)} = \alpha q_{(\mathrm{d},\mathrm{e}_*)} + (1 - \alpha)q'_{(\mathrm{d},\mathrm{e}_*)}$. This determines invariant probabilities $(\mathrm{Pr}'', \mathrm{Pr}''_*)$ that assigns probability p'' to E. It suffices to show that d is a strict best response under Pr'' and e_* is a strict best response under Pr''_*. We show the claim for Pr''; an analogous argument applies to Pr''_*.

To show that d is a strict best response under Pr'': Construct an auxiliary game. Player Pi chooses again between $\{\mathrm{a}, \mathrm{b}, \mathrm{c}\}$ and player P$(i - 1)$ chooses between $\{\mathrm{a}, \mathrm{b}, \mathrm{c}, \hat{\mathrm{a}}, \hat{\mathrm{b}}, \hat{\mathrm{c}}\}$. P$i$'s payoffs correspond to the original game, when restricting to the strategy profiles in $\{\mathrm{a}, \mathrm{b}, \mathrm{c}\}^2$. From the perspective of Pi, the columns $\hat{\mathrm{d}}$ are payoff copies of columns d. (So, for any action Pi chooses, she gets the same payoff if P$(i - 1)$ plays d versus $\hat{\mathrm{d}}$.) The payoffs of P$(i - 1)$ are irrelevant.

We can use Pr, Pr' and Pr'' to construct belief $\hat{\mathrm{Pr}}$, $\hat{\mathrm{Pr}}'$ and $\hat{\mathrm{Pr}}''$ on $\{\mathrm{a}, \mathrm{b}, \mathrm{c}, \hat{\mathrm{a}}, \hat{\mathrm{b}}, \hat{\mathrm{c}}\}$: Specifically $\hat{\mathrm{Pr}}(\mathrm{d}) = \rho_{\mathrm{d}}$ and $\hat{\mathrm{Pr}}(\hat{\mathrm{d}}) = q_{(\mathrm{d},\mathrm{a}_*)} + q_{(\mathrm{d},\mathrm{b}_*)} + q_{(\mathrm{d},\mathrm{c}_*)}$. And analogously for $\hat{\mathrm{Pr}}'$ and $\hat{\mathrm{Pr}}''$. Observe that an action for Pi is a best response under Pr (resp. Pr', Pr'') in the original game if and only if it is a best response under $\hat{\mathrm{Pr}}$ (resp. $\hat{\mathrm{Pr}}'$, $\hat{\mathrm{Pr}}''$) in the constructed game. Thus, d is a strict best response under $\hat{\mathrm{Pr}}$ and a a best response under $\hat{\mathrm{Pr}}'$. Since $\hat{\mathrm{Pr}}''$ is a non-degenerate convex combination of $\hat{\mathrm{Pr}}$ and $\hat{\mathrm{Pr}}'$, it follows that d is a strict best response under $\hat{\mathrm{Pr}}''$. From this, d is a strict best response under Pr'' in the original game. And, analogously for e_*. ■

C.2 Bound Calculations

The following re-expresses Table 5.1, in a way that easily permits computing the lower- and upper-bounds of P3:

p	Actions
$[\frac{7}{8}, 1]$	$(\mathrm{a}, \mathrm{b}_*)$
$[\frac{5}{7}, \frac{7}{8})$	$(\mathrm{a}, \mathrm{b}_*)$, $(\mathrm{a}, \mathrm{c}_*)$
$[\frac{62}{133}, \frac{5}{7})$	$(\mathrm{a}, \mathrm{b}_*)$, $(\mathrm{a}, \mathrm{c}_*)$, $(\mathrm{b}, \mathrm{c}_*)$
$(\frac{2}{5}, \frac{62}{133})$	$(\mathrm{a}, \mathrm{b}_*)$, $(\mathrm{a}, \mathrm{c}_*)$, $(\mathrm{b}, \mathrm{c}_*)$,$(\mathrm{a}, \mathrm{a}_*)$

p	Actions
$[\frac{2}{5}, \frac{2}{5}]$	$(\mathrm{a}, \mathrm{b}_*)$, $(\mathrm{b}, \mathrm{c}_*)$,$(\mathrm{a}, \mathrm{a}_*)$
$(\frac{2}{7}, \frac{2}{5})$	$(\mathrm{a}, \mathrm{b}_*)$, $(\mathrm{b}, \mathrm{c}_*)$,$(\mathrm{a}, \mathrm{a}_*)$,$(\mathrm{b}, \mathrm{a}_*)$
$[\frac{29}{133}, \frac{2}{7}]$	$(\mathrm{b}, \mathrm{c}_*)$,$(\mathrm{a}, \mathrm{a}_*)$,$(\mathrm{b}, \mathrm{a}_*)$
$(\frac{1}{11}, \frac{29}{133})$	$(\mathrm{b}, \mathrm{c}_*)$,$(\mathrm{a}, \mathrm{a}_*)$,$(\mathrm{b}, \mathrm{a}_*)$,$(\mathrm{c}, \mathrm{a}_*)$
$[0, \frac{1}{11}]$	$(\mathrm{b}, \mathrm{c}_*)$, $(\mathrm{b}, \mathrm{a}_*)$,$(\mathrm{c}, \mathrm{a}_*)$

Table C.1: Assigning probability p to Rationality

This next table provides the lower- and upper-bounds of P3:
In the role of P4, we observe three strategies played: $(\mathrm{a}, \mathrm{b}_*)$, $(\mathrm{a}, \mathrm{c}_*)$, and $(\mathrm{c}, \mathrm{a}_*)$. The next table provides the lower- and upper-bounds of P4 for those observations.

	$(\mathrm{a}, \mathrm{a}_*)$	$(\mathrm{a}, \mathrm{b}_*)$	$(\mathrm{a}, \mathrm{c}_*)$	$(\mathrm{b}, \mathrm{a}_*)$	$(\mathrm{b}, \mathrm{b}_*)$	$(\mathrm{b}, \mathrm{c}_*)$	$(\mathrm{c}, \mathrm{a}_*)$	$(\mathrm{c}, \mathrm{b}_*)$	$(\mathrm{c}, \mathrm{c}_*)$
$q \in [\frac{7}{8}, 1]$	$[0, \frac{5}{8})$	NA	$(0, \frac{15}{32})$	$(0, 1]$	$[0, \frac{14}{31})$	$(0, \frac{7}{8})$	NA	NA	$[0, \frac{15}{62})$
$q \in [\frac{5}{7}, \frac{7}{8})$	$[0, \frac{5}{8})$	$(0, \frac{19}{48})$	$(0, \frac{5}{8})$	$(0, 1]$	$[0, \frac{5}{6})$	$(0, 1]$	NA	NA	$[0, \frac{15}{62})$
$q \in [\frac{62}{133}, \frac{5}{7})$	$[0, \frac{3}{4})$	$(0, \frac{5}{6})$	$(0, 1]$	$(0, 1]$	$[0, \frac{5}{6})$	$(0, 1]$	NA	$(0, \frac{7}{12})$	$[0, \frac{7}{8})$
$q \in [\frac{2}{5}, \frac{62}{133})$	$[0, \frac{231}{248})$	$(0, 1]$	$(0, 1]$	$(0, 1]$	$[0, 1]$	$(0, 1]$	NA	$(0, \frac{3}{4})$	$[0, \frac{7}{8})$
$q \in [\frac{29}{133}, \frac{2}{5})$	$[0, 1]$	$(0, 1]$	$(0, 1]$	$(0, 1]$	$[0, 1]$	$(0, 1]$	$(0, \frac{15}{31})$	$(0, \frac{7}{8})$	$[0, \frac{7}{8})$
$q \in [0, \frac{29}{133})$	$[0, 1]$	$(0, 1]$	$(0, 1]$	$(0, 1]$	$[0, 1]$	$(0, 1]$	$(0, \frac{215}{248})$	$(0, 1]$	$[0, 1]$

Table C.2: P3: Probability p to “Rationality and q-Belief of Rationality”

	$(\mathrm{a}, \mathrm{b}_*)$	$(\mathrm{a}, \mathrm{c}_*)$	$(\mathrm{c}, \mathrm{a}_*)$
$q \in [\frac{7}{8}, 1]$ $r \in [\frac{7}{8}, 1]$	$(0, \frac{1}{3})$	$(0, 1]$	NA
$q \in [\frac{15}{32}, \frac{7}{8})$ $r \in [\frac{7}{8}, 1]$	$(0, 1]$	$(0, 1]$	NA
$q \in [\frac{7}{8}, 1]$ $r \in [\frac{5}{7}, \frac{7}{8})$	$(0, 1]$	$(0, 1]$	NA
$q \in [\frac{7}{8}, 1]$ $r \in [\frac{62}{133}, \frac{5}{7})$	$(0, 1]$	$(0, 1]$	$(0, 1]$
$q \in (0, \frac{15}{32})$ $r \in [\frac{7}{8}, 1]$	$(0, 1]$	$(0, 1]$	$(0, 1]$
$q \in (0, \frac{5}{8})$ $r \in [\frac{5}{7}, \frac{7}{8})$	$(0, 1]$	$(0, 1]$	$(0, 1]$

Table C.3: P4: Probability p to “Rationality and q-Belief of ‘Rationality and r-Belief of Rationality’ ”

Appendix D The Anonymity Assumptions

D.1 Proof of Proposition 5.1

It will be convenient to introduce some notation: Write R_i for the event that player i is rational. Write $B_i^p(E_j)$ for the event that i assigs probability p to the event E_j. Write $\tilde{B}_i^p(E_j)$ for the event that i assigs probability at least probability p to event E_j.

It will be convenient to record four properties of these belief operators.

Property 1 $B_i^p(E_j) \implies \tilde{B}_i^p(E_j)$.

Property 2 If E_j, F_j are events with $E_j \subseteq F_j$, then $\tilde{B}_i^p(E_j) \implies \tilde{B}_i^p(F_j)$.

Property 3 If $q \leq p$, then $\tilde{B}_i^p(E_j) \implies \tilde{B}_i^q(E_j)$.

Property 4 If $\tilde{B}_i^p(E_j) \wedge B_i^q(E_j)$, then $q \geq p$.

Property 1 says that if i assigns probability p to E_j then i also assigns *at least* probability p to E_j. Property 2 says that, if i assigns probability at least p to E_j, then i assigns probability at least p

to any larger event F_j.[20] Property 3 says that if, if i assigns probability at least p to E_j, then i also assigns probability at least $q \leq p$ to E_j. Property 4 says that if i assigns probability at least p to E_j and probability of exactly q to E_j, then $q \geq p$.

Lemma D.1. *Fix $i \in \{3, 4\}$. If the Anonymity Assumption holds and the subject is i-cognitive, then $p_2 \geq p_i$.*

Proof. Suppose that Pi is i-cognitive. By the P2 Requirement, $B_2^{p_2}(R_1)$ holds. So by the Anonymity Assumption $B_i^{p_2}(R_{i-1})$ holds. At the same time, by the Pi Requirement and Properties 1-2, $\tilde{B}_i^{p_i}(R_{i-1})$ holds. So applying Property 4, $p_2 \geq p_i$. ☐

Lemma D.2. *Suppose the Anonymity Assumption holds and the subject is 4-cognitive. If $q_2 \leq q_3$ then $p_3 \geq p_4$.*

Proof. By the P4 Requirement

$$B_4^{p_4}\left(R_3 \cap \tilde{B}_3^{q_3}(R_2 \cap \tilde{B}_2^{r_1}(R_1))\right),$$

i.e., P4 assigns probability p_4 to $R_3 \cap \tilde{B}_3^{q_3}(R_2 \cap \tilde{B}_2^{r_1}(R_1))$. Applying Property 1, it follows that

$$\tilde{B}_4^{p_4}\left(R_3 \cap \tilde{B}_3^{q_3}(R_2 \cap \tilde{B}_2^{r_1}(R_1))\right),$$

i.e., P4 assigns at least probability p_4 to $R_3 \cap \tilde{B}_3^{q_3}(R_2 \cap \tilde{B}_2^{r_1}(R_1))$. Observe that $R_2 \cap \tilde{B}_2^{r_1}(R_1) \subseteq R_2$ and so, applying Property 2,

$$\tilde{B}_4^{p_4}\left(R_3 \cap \tilde{B}_3^{q_3}(R_2)\right).$$

And implication of the Anonymity Assumption is that

$$\tilde{B}_3^{p_4}\left(R_2 \cap \tilde{B}_2^{q_3}(R_1)\right).$$

Now, suppose that $q_2 \leq q_3$. By Property 3, $\tilde{B}_2^{q_3}(R_1)$ implies $\tilde{B}_2^{q_2}(R_1)$. From this, $R_2 \cap \tilde{B}_2^{q_3}(R_1) \subseteq R_2 \cap \tilde{B}_2^{q_2}(R_1)$. So, again applying Property 2,

$$\tilde{B}_3^{p_4}\left(R_2 \cap \tilde{B}_2^{q_2}(R_1)\right)$$

But by the P3 requirement, $B_3^{p_3}(R_2 \cap B_2^{q_2}(R_1))$. So, applying Property 4, $p_3 \geq p_4$. ☐

D.2 Proposition 5.2

We now show Proposition 5.2 for the case where the subject's cognitive bound is identified as $k = 4$. Since there is a gap between the subject's cognitive and rationality bounds, the subject's rationality bound is either $m = 1, 2$, or 3. We will argue that, for this subject, either

---

[20] A premise of this property is that both E_j and F_j are events, i.e., measurable sets. In what follows, we apply this property to sets that are measurable. The proof that those sets are measurable is standard in the epistemic literature and, so, omitted.

(A) $p_4 > 0$ and $(p_4, q_3, r_2) \neq (1, 1, 1)$, or

(B) $p_3 > 0$ and $(p_3, q_2) \neq (1, 1)$.

The argument begins as in the case of $k = 3$: First, since the subject is identified as having a cognitive bound of $k = 4$, the subject plays a non-constant strategy in the role of P4. From this $p_4 > 0$. (Table C.3 verifies this holds in the observed data.) If the subject plays a strategy that does not survive 4 rounds of iterated dominance—i.e., a strategy $(\mathrm{d}, \mathrm{e}_*) \neq (\mathrm{a}, \mathrm{c}_*)$—then it is immediate that we cannot have $(p_4, q_3, r_2) \neq (1, 1, 1)$. Likewise, if the subject instead has a rationality bound of $m = 1$, then $p_2 < 1$. So by the AA, $1 > p_2 \geq p_4$. From this we again conclude that $(p_4, q_3, r_2) \neq (1, 1, 1)$.

However, we observe 10 observations that don't fit this mold: In those observations, subjects play the play the iteratively undominated strategies in the roles of both P2 and P4 (i.e., $m = 2$). Thus, in principle, we may have $(p_4, q_3, r_2) = (1, 1, 1)$. That said, we will now argue that, if $m = 2$ and such a subject is characterized by a $(p_2; p_3, q_2; p_4, q_3, r_2)$ with $(p_4, q_3) = (1, 1)$, then we must have both $p_3 > 0$ and $(p_3, q_2) \neq (1, 1)$.

Suppose that $(p_4, q_3) = (1, 1)$. In that case, $q_2 \leq q_3 = 1$ and so, by the AA, $p_3 \geq p_4 \geq 1 > 0$. At the same time, because the subject has a rationality bound of $m = 2$, $(p_3, q_2) \neq (1, 1)$. This establishes the claim.

References

Agranov, M., E. Potamites, A. Schotter, and C. Tergiman (2012). Beliefs and endogenous cognitive levels: An experimental study. *Games and Economic Behavior 75*, 449–463.

Alaoui, L. and A. Penta (2016). Endogenous depth of reasoning. *Review of Economic Studies*. Forthcoming.

Battigalli, P. and M. Siniscalchi (2002). Strong belief and forward induction reasoning. *Journal of Economic Theory 106*(2), 356–391.

Bernheim, B. D. (1984). Rationalizable strategic behavior. *Econometrica*, 1007–1028.

Bhatt, M. and C. F. Camerer (2005). Self-referential thinking and equilibrium as states of mind in games: fMRI evidence. *Games and Economic Behavior 52*, 424–459.

Blume, A. and U. Gneezy (2010). Cognitive forward induction and coordination without common knowledge: An experimental study. *Games and Economic Behavior 68*, 488–511.

Brandenburger, A. and A. Friedenberg (2014). How many levels do players reason? an observational challenge and solution.

Brandenburger, A. and X. Li (2015). Thinking about thinking and its cognitive limits. Working paper, New York University.

Burchardi, K. B. and S. P. Penczynski (2014). Out of your mind: Eliciting individual reasoning in one shot games. *Games and Economic Behavior 84*, 39–57.

Camerer, C. F., T.-H. Ho, and J.-K. Chong (2004). A cognitive hierarchy model of games. *Quarterly Journal of Economics 119*(3), 861–898.

Chen, C.-T., C.-Y. Huang, and J. T.-y. Wang (2009). A window of cognition: Eyetracking the reasoning process in spatial beauty contest games. Working paper, National Taiwan University.

Coricelli, G. and R. . P. N. A. S. U. Nagel (2009). Neural correlates of depth of strategic reasoning in medial prefrontalcortex. *Proceedings of the National Academy of Sciences 106*, 9163–9168.

Costa-Gomes, M., V. P. Crawford, and B. Broseta (2001). Cognition and behavior in normal-form games: An experimental study. *Econometrica 69*, 1193–1235.

Costa-Gomes, M. A. and V. P. Crawford (2006). Cognition and behavior in two-person guessing games: An experimental study. *American Economic Review 96*, 1737–1768.

Crawford, V. P., M. A. Costa-Gomes, and N. Iriberri (2013). Structural models of nonequilibrium strategic thinking: Theory, evidence, and applications. *Journal of Economic Literature 51*, 5–62.

Fragiadakis, D. E., D. T. Knoepfle, and M. Niederle (2013). Identifying predictable players: Relating behavioral types and subjects with deterministic rules. Working paper.

Georganas, S., P. J. Healy, and R. A. Weber (2015). On the persistence of strategic sophistication. *Journal of Economic Theory 159*, 369–400.

Gill, D. and V. Prowse (2016). Cognitive ability, character skills, and learning to play equilibrium: A level-k analysis. *Journal of Political Economy*. Forthcoming.

Kinderman, P., R. Dunbar, and R. Bentall (1998). Theory-of-Mind deficits and causal attributions. *British Journal of Psychology 89*, 191–204.

Kneeland, T. (2015). Identifying higher-order rationality. *Econometrica 83*, 2065–2079.

Nagel, R. (1995). Unraveling in guessing games: An experimental study. *American Economic Review 85*, 1313–1326.

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

Perner, J. and H. Wimmer (1985). "John thinks that Mary thinks that…": Attribution of second-order beliefs by 5- to 10-year-old children. *Journal of Experimental Child Psychology 39*, 437–471.

Rubinstein, A. (2007). Instinctive and cognitive reasoning: A study of response times. *The Economic Journal 117*(523), 1243–1259.

Stahl, D. O. and P. W. Wilson (1995). On players' models of other players: Theory and experimental evidence. *Games and Economic Behavior 10*, 218–254.

Stiller, J. and R. Dunbar (2007). Perspective-taking and memory capacity predict social network size. *Social Networks 29*, 93–104.

Tan, T. C.-C. and S. R. da Costa Werlang (1988). The bayesian foundations of solution concepts of games. *Journal of Economic Theory 45*(2), 370–391.

Wang, J. T.-y., M. Spezio, and C. F. Camerer (2009). Pinocchio's pupil: Using eyetracking and pupil dilation to understand truth telling and deception in senderreceiver games. *American Economic Review 100*, 984–1007.