

Bounded Reasoning: Rationality or Cognition^{*}

Amanda Friedenberg[†]

Willemien Kets[‡]

Terri Kneeland[§]

First Draft: January 31, 2016

This Draft: November 18, 2017

Abstract

Bounded reasoning about rationality can have important implications for behavior. These bounds are typically viewed as an artifact of limits in the ability to engage in interactive reasoning, i.e., to reason through “I think, you think, I think, etc ...” However, in principle, these bounds need not be determined by limits in ability. This paper develops a novel identification strategy to show that bounded reasoning about rationality is not determined by limitations in ability. It goes on to show that non-degenerate beliefs about rationality can be an important determinant of behavior. This has important implications for out-of-sample predictions.

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

^{*}We are indebted to Adam Brandenburger and Vince Crawford for many helpful conversations. We also thank Marina Agranov, Martin Dufwenberg, Guillaume Fréchette, Ed Green, P.J. Healy, Toru Kitagawa, Brian Rogers, Pablo Schenone, and participants of seminars at Caltech, Texas, Penn State, UCL, Queen Mary, WZB, CERGE-EI, LSE, Utah, UCSD, York, Oxford, Barcelona GSE Summer Forum, BRIC, the Game Theory Society World Congress, Alicante Level- k Workshop, the CRETA Conference on Identifying Rationality, and the Society for Advancement of Economic Theory. Friedenberg thanks NSF grant SES-1358008 for financial support and Caltech, Kellogg School of Management and UCL for hospitality.

[†]W.P. Carey School of Business, Arizona State University, afrieden@asu.edu, <http://amandafriedenberg.org>

[‡]Department of Economics, University of Oxford, willemien.kets@economics.ox.ac.uk, <http://wkets.org>

[§]Department of Economics, University College London, t.kneeland@ucl.ac.uk, <http://tkneeland.com>

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.) Limitations on the ability to engage in interactive reasoning can limit players’ ability to engage in 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? Addressing this question is important both for the external validity of laboratory experiments and for how behavioral game theory builds models of bounded reasoning. We now explain.

First, in a laboratory setting, experimental game theory has shown that there is bounded reasoning about rationality. But, it leaves open whether those bounds are behaviorally relevant when it comes to important economic and social decisions (i.e., outside of the laboratory). In fact, 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](#).) As a consequence, limitations on ability may not be binding on important decisions. That said, if bounds on reasoning about rationality arise from other sources—and those other sources are independent of the stakes of the game—then 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 or with a population of like-minded 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 Bob were prepared to bet on the fact that Ann is rational, Ann may consider the possibility that Bob considers the possibility that she is irrational. (This might be the case if, in the past, Ann chose rationally but, given Bob’s past behavior, she concluded that Bob did not understand important parameters of her problem.) And so on. Thus, despite the fact that the executives can engage in interactive reasoning, bounded reasoning about rationality may well be important for understanding how the executives act.

Second, the standard approach in behavioral game theory is to model players as if they are

each characterized by a single bound. (This is true of the level- k and cognitive hierarchy models.) Implicit in this choice is an assumption that the rationality bound is determined by the player’s ability to engage in interactive reasoning. An implication is that, if a player reasons *only* 2 rounds about rationality, she cannot consider the possibility that the other player reasons 3 rounds.¹ However, if the rationality bound is not determined by the ability to engage in interactive reasoning, then a player who only reasons 2 rounds about rationality may well consider the possibility that the other player reasons 3 rounds. We expand on this point on page 5 below. The implication is that players may have a wide range of non-degenerate beliefs about rationality—beliefs that are not captured by existing models in behavioral game theory. In fact, our analysis will show that the behavior of certain ‘gap subjects’ is best explained by such non-degenerate beliefs.

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 criterion for playing the game which departs from subjective expected utility.

Consider a player who engages in full interactive reasoning, i.e., 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 playing a best response given her belief.

With this in mind, we focus on subjects who are rational, i.e., who play a best response given

¹This is true of the level- k model. The caveat is that the cognitive hierarchy model, nonetheless, allows a limited form of such beliefs. See Appendix A.

²Both the terms “cognitive” and “rational” are words that take on different meanings in different literatures. Our usage of the term *cognitive* is consistent with the usage of the term “social cognition” in psychology and neuroscience. (See Adolphs, 1999.) Our usage of the term *rational* is consistent with the usage in epistemic game theory. (See, e.g., Brandenburger, 2007 and Dekel and Siniscalchi, 2014.) In particular, here, rationality only refers to the idea that a player chooses a best response given her beliefs. Likewise, here, cognitive will take on a precise meaning clarified in Sections 3.2.1-3.2.2 below.

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 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 k if she is cognitive, believes that Bob is cognitive, believes that Bob believes she is cognitive and so on, up to the statement that includes the word “cognitive” k times, but no further.

Because rationality is one theory about how to play the game, a subject’s rationality bound m cannot be higher than her cognitive bound k , i.e., $m \leq k$. If a subject’s rationality bound is lower than her cognitive bound, i.e., if $m < k$, 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, 3.2.1 and 3.2.2 will clarify the precise presentation effects that are ruled out.)

To sum up, our approach is to use observed play of the game, to deduce that there is a gap between the rationality and cognitive bounds. This allows the researcher to deduce that the rationality bound is not entirely determined by limited ability to engage in interactive reasoning. One might hope for an alternate approach to address this question: one where the researcher elicits the subject’s hierarchies of beliefs and uses this to show that players engage in interactive reasoning beyond their reasoning about rationality. Such an approach would be more direct. However, the concern is that the very act of attempting to elicit the subjects’ hierarchies may cause subjects to engage in higher levels of interactive reasoning than they may otherwise do. With this in mind, we instead draw inferences from the play of the game.

Preview of Results We provide a novel strategy to identify a gap between reasoning about rationality versus reasoning about cognition. Section 2 provides a preview of the identification strategy; it is based on Kneeland’s (2015) ring game experiment. To identify the gap, we assume that each subject is rational, i.e., plays a best response given her belief. We identify behavior that is consistent with reasoning about cognition but inconsistent with reasoning about rationality.

Section 4 applies the 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 4. Moreover, there is a nontrivial gap between subjects’ rationality and cognitive bounds. We find that 47% of the subjects identified as

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

Figure 1.1. A Game Parameterized by $x \in (-\infty, 10)$

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 4.

Section 6 explores an important model selection question. 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.

Non-Degenerate Beliefs If a subject’s rationality bound is entirely determined by her ability to engage in interactive reasoning, then she necessarily has degenerate beliefs in reasoning about rationality. For instance, consider a subject who has a rationality bound of $m = 2$. The subject assigns probability one to “Bob is rational,” but not “Bob is rational and believes I am rational.” If her rationality bound is determined by her ability to engage in interactive reasoning, the failure to believe “Bob is rational and believes I am rational” results from her inability to specify a third-order belief. Thus, the subject must also have a cognitive bound of $k = 2$.

By contrast, if there is a gap between a subject’s rationality and cognitive bounds, then there is scope for non-degenerate beliefs in reasoning about rationality. In fact, Section 7 shows that such non-degenerate beliefs exist: If a subject is identified as having a gap between her rationality and cognitive bound, her behavior is best explained by non-degenerate beliefs about rationality.

We already argued that non-degenerate beliefs in reasoning about rationality has important implications for building models of behavioral game theory. (See page 2.) But, it is also important for interpreting extant experimental results. Some paper have found that subjects’ identified levels of reasoning change across games. (See [Georganas, Healy, and Weber, 2015](#), [Cooper, Fatas, Morales, and Qi, 2016](#).) This has been interpreted as an instability in the levels of reasoning—that a subject may reason m levels in one game and m' levels in another game. However, the same behavior can instead reflect non-degenerate beliefs about rationality. To see why, refer to Figure 1.1. If Ann plays a strategy 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 hold 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 for any $x \in (-\infty, 10)$. However, if 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. In fact, [Cooper et al. \(2016\)](#) show that varying a player’s payoff parameter in this way does impact the play of the game.

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 using level- k and cognitive hierarchy models (e.g., [Nagel, 1995](#); [Stahl and Wilson, 1995](#); [Costa-Gomes, Crawford, and Broseta, 2001](#); [Camerer, Ho, and Chong, 2004](#); [Costa-Gomes and Crawford, 2006](#)). There is a subtle relationship between rationalizability, the level- k model, and the cognitive hierarchy model. See [Appendix A](#).

The level- k and cognitive hierarchy models are often motivated by limitations on the players’ ability to engage in interactive reasoning. (See, e.g., [Nagel, 1995](#), pg. 1313 or [Camerer, Ho, and Chong, 2004](#), pg. 864.) That said, typically, the literature identifies subjects’ ability to engage in interactive reasoning based on the extent to which they iterate over best responses: A subject is identified as a level- k thinker if she performs exactly k rounds of iterated best responses.³ To the best of our knowledge, this is the first paper that can directly address the question whether the rationality bounds are determined by limited ability.⁴

That said, the literature has noted that a subject may have a low rationality bound even if she is capable of reasoning to higher orders. Specifically, in the context of level- k models, [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.⁵ At first glance, this variation might suggest that reasoning about rationality is not driven by limited ability: If a subject’s ability to engage in interactive reasoning does not depend on who her opponents are, then

³Sometimes this identification is augmented with auxiliary data that is suggestive of ability limitations. 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.

⁴A contemporaneous paper by [Jin \(2016\)](#) seeks to identify whether rationality bounds are determined by ability. The paper, in fact, studies whether reasoning about rationality in a sequential game corresponds to reasoning about rationality in a different simultaneous move game.

⁵There 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.

any variation in her rationality bounds must indicate that the bounds are not entirely determined by the difficulties of interactive reasoning. However, it is not clear that the premise holds. In particular, the premise would be false if the subject adapts her effort in interactive reasoning (i.e., how much effort she exerts on “I think, you think, I think...”), based on who her 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.

To identify the rationality and cognitive bounds, we build on the work of [Kneeland \(2015\)](#). A key conceptual difference between her work and ours is that [Kneeland](#) does not distinguish between the rationality bound and the cognitive bound. In fact, she implicitly assumes that the rationality bound is determined by limitations on a player’s ability to engage in interactive reasoning. We introduce the notion of a cognitive bound and use this to show that the rationality bound is not determined by such limitations in ability.⁶

In the course of our analysis, we show that non-degenerate beliefs about rationality are an important determinant of behavior. There is a long theoretical literature that explicitly models non-degenerate hierarchies of beliefs (e.g., [Monderer and Samet, 1989](#) and [Morris, 1999](#)) and those ideas can be used to provide an explicit model of non-degenerate beliefs about rationality (e.g., [Hu, 2007](#)).⁷ Moreover, as [Appendix A](#) describes, the level- k and cognitive hierarchy models can be used to implicitly model non-degenerate beliefs about rationality. However, to the best of our knowledge, this is the first paper to provide a general empirical analysis of beliefs about rationality—one that is suitable to identify (in the sense of model selection) non-degenerate beliefs.

As a consequence of this finding, seemingly irrelevant payoff parameters may have important implications for behavior. At the surface, this message may seem familiar from [Alaoui and Penta \(2016\)](#). However, there is an important difference. [Alaoui and Penta](#) argue that seemingly irrelevant payoff parameters can change the number of steps a player takes in interactive reasoning. We argue that—even if it does not change the number of steps a player takes in interactive reasoning—it can still influence behavior. This is because players have non-degenerate beliefs about rationality.

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 empirical result: a gap between the rationality and cognitive bounds. [Section 5](#) illustrates that, even if there is a gap between the rationality and cognitive bounds, a subject’s beliefs about play can be reinterpreted as beliefs about rationality. [Section 6](#) uses this

⁶Whether there can be a gap between the bounds has implications for the interpretation of the rationality bound. [Kneeland \(2015\)](#) assumes that there is no gap and uses that hypothesis to identify the *exact* level of reasoning about rationality consistent with the data. We show that there can be a gap between the rationality and the cognitive bound. This suggests that the rationality bounds identified in [Kneeland](#) are best interpreted as the *maximum* level of reasoning about rationality consistent with the data.

⁷[Hu \(2007\)](#) carries out his analysis in the context of a standard type structure, which does not capture limits in ability. However, in principle, the frameworks in [Kets \(2011\)](#) and [Heifetz and Kets \(2017\)](#) can be used to provide model non-degenerate beliefs about rationality when there are ability limitations.

characterization to show that the observed gap is not an artifact of noisy decision-making. Finally, Section 7 uses the empirical analysis in Section 6 to argue that the observed gap is best explained by non-degenerate beliefs about rationality.

2 An Illustrative Example

Figures 2.1a-2.1b describe two games, G and G_* . The payoff matrices on the left represent player 1's payoffs and the payoff matrices on the right represent player 2's payoffs. We will write (d, e_*) to denote that a player chooses action d in G and action e_* in G_* . We often refer to such an action profile as 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 labeled c_* in G_* , the row b in G is labeled a_* in G_* , and the row c in G is labeled 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.

		P1's Payoffs					P2's Payoffs				
		P2					P1				
			a	b	c				a	b	c
P1		a	12	16	14	P2		a	20	14	8
		b	8	12	10			b	16	2	18
		c	6	10	8			c	0	16	16

(a) Figure G

		P1's Payoffs					P2's Payoffs				
		P2					P1				
			a_*	b_*	c_*				a_*	b_*	c_*
P1		a_*	8	12	10	P2		a_*	20	14	8
		b_*	6	10	8			b_*	16	2	18
		c_*	12	16	14			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 chooses a best response given her subjective belief about how P2 plays the game. Then, she would play the strategy (a, 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 (c, 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 (a, c_*) . As P2 is rational, he chooses a best response to this belief; he thus plays (a, 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 $4/5$ to the rational strategy (a, c_*) and probability $1/5$ to the lucky-6 strategy (c, b_*) . In this case, his best response is to play (a, 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_* . And, similarly, for b (resp. c) and b_* (resp. c_*). 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., (a, a_*) , (b, b_*) , or (c, 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.

Notice, however, that a player’s cognitive bound may be strictly higher than her rationality bound. If it is, then it indicates that bounded reasoning about rationality is not entirely determined

by limits in ability.⁸ With this in mind, our question is: Does there exist a gap between the cognitive and rationality bounds? We seek a conservative estimate of the gap. As such, we seek to identify:

- (i) The *maximum* level of reasoning about rationality consistent with observed behavior.
- (ii) The *minimum* level of reasoning about cognition consistent with observed 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 a player’s choosing a best response given her belief. (Most of the observations in our dataset are consistent with rational behavior; we restrict attention to those observations.) As a consequence, we assume that all behavior is cognitive. That is, we do not attempt to distinguish rational behavior from cognitive behavior. Instead, our identification focuses on reasoning about rationality vs. reasoning about cognition. In light of this, we focus on the observed behavior of P2. We distinguish between three scenarios.

First, we identify P2 as having a *rationality bound of 2* if his behavior is consistent with being rational and believing (i.e., assigning probability 1 to the event) that P1 is rational. Thus, we identify P2 as having a rationality bound of 2 if and only if his behavior is consistent with two rounds of iterated dominance—i.e., if and only if we observe P2 play (a, b_*) . Notice that this behavior—i.e., the strategy (a, b_*) —is also consistent with P2 being rational and assigning (only) probability $9/10$ to P1’s rationality. (For instance, it is a best response for P2 to play (a, b_*) , if he assigns probabilities $\Pr(a, c_*) = 9/10$ and $\Pr(b, a_*) = \Pr(c, a_b) = 1/20$.) But because we seek the maximum level of reasoning about rationality consistent with observed behavior, we identify the rationality bound as 2.

If we identify P2 as having a rationality bound of 2, then we will also identify P2 as having a cognitive bound of 2. To understand why, notice that if P2’s behavior is consistent with a rational P2 believing that P1 is rational, then it is also consistent with a rational P2 believing that P1 is cognitive. Importantly, such behavior is inconsistent with a rational P2 believing that P1 lacks cognition: Recall, if P2 is rational and believes that P1 lacks cognition, then P2’s behavior does not vary across G and G_* . Thus, if we observe P2 play the non-constant action profile (a, b_*) , we must conclude that P2 reasons that P1 is cognitive. As such, the minimum level of cognition consistent with observed behavior is 2 and, so, we identify this behavior as having a *cognitive bound of 2*.

Second, we identify P2 as having a *cognitive bound of 1* if his behavior is consistent with being rational and believing (i.e., assigning probability 1 to the event) that P1 lacks cognition. Thus, we identify P2 as having a cognitive bound of 1 if and only if he plays a constant strategy. To understand why, notice that if a rational P2 believes that P1 lacks cognition, then he plays a constant action profile. Moreover, each constant action profile is consistent with a rational P2 believing that P1 lacks cognition. For instance, it is a best response for P2 to play (a, a_*) , if he

⁸An important but subtle point: The cognitive bound need not map one-for-one with a bound on ability. Rather, we use the cognitive bound as a vehicle to show that the rationality bound is not entirely determined by limitations in ability. The gap between the cognitive and rationality bounds show just that.

assigns probability 1 to P1 also playing (a, a_*) ; such a constant belief is consistent with believing that P1 lacks cognition. Notice that this behavior—i.e., the constant strategy (a, a_*) —is also consistent with a rational P2 believing that P1 is cognitive. (For instance, it is a best response for P2 to play (a, a_*) if he assigns probabilities $\Pr(a, c_*) = 2/5$ and $\Pr(b, a_*) = \Pr(c, b_*) = 3/10$.) But because we seek the minimum level of reasoning about cognition consistent with observed behavior, we identify the cognitive bound as 1.

If we identify P2 as having a cognitive bound of 1, then we will also identify P2 as having a rationality bound of 1. To understand why, notice that if a rational P2 plays a constant strategy, then he cannot believe that P1 is rational. As such, the maximum level of reasoning about rationality consistent with observed behavior is 1 and, so, we identify this behavior as having a *rationality bound of 1*.

In these first two scenarios, we would not identify a gap between the rationality and cognitive bounds. We will identify a gap between P2’s rationality and cognitive bounds if he plays a non-constant strategy profile that differs from (a, b_*) . As an illustration, suppose we observe P2 play (a, c_*) . This behavior is consistent with one round—but not two rounds—of iterated dominance. Thus, we would identify P2 as having a rationality bound of 1. At the same time, because this observation is not a constant strategy, it is inconsistent with a rational P2 believing that P1 lacks cognition. Moreover, we have also seen that it is consistent with a rational P2 believing that P1 is cognitive. Thus, we identify the subject as having a cognitive bound of 2.

Identifying the Bounds: A Comment Suppose that P2 reasons about P1’s rationality. In the above discussion, we think of this scenario as one in which P2 believes—i.e., assigns probability 1 to the event—that P1 is rational. If P2 assigns probability $p = 4/5$ to P1’s rationality, we think of this as a departure from reasoning about rationality.

Likewise, 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 $p = 4/5$ 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.

This conceptual point has a pragmatic implication for how we identify cognition: We only identify P2’s level of cognition if his behavior is consistent with rationality and assigning probability $p \in \{0, 1\}$ to the event that P1 is cognitive. We identify P2 as having a cognitive bound of 1 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 having a cognitive bound of 2 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.⁹ By restricting p to be in $\{0, 1\}$, we limit our ability

⁹Notice, there is a difference between (a) not believing that a player is cognitive and (b) believing that the player lacks cognition. In the former case, there is uncertainty whether the player has a theory for playing the game. In the latter case, it is certain that the player does not have a theory for playing the game (and, so, it is certain that her behavior does not depend on specific parameters of the game). We restrict attention to the latter scenario.

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 paper studies a four-player game and experiment. This allows us to identify the gap up to four levels of reasoning. The next section explains the identification strategy.

3 Identification

Figures 3.1a-3.1b describe two games, G and G_* , from Kneeland (2015). Each of the games has a *ring structure*: Player i 's (P_i 's) payoffs depend only on the behavior of Player $(i - 1)$ ($P_{(i - 1)}$). (We adopt the convention that $P_0 \equiv P_4$).

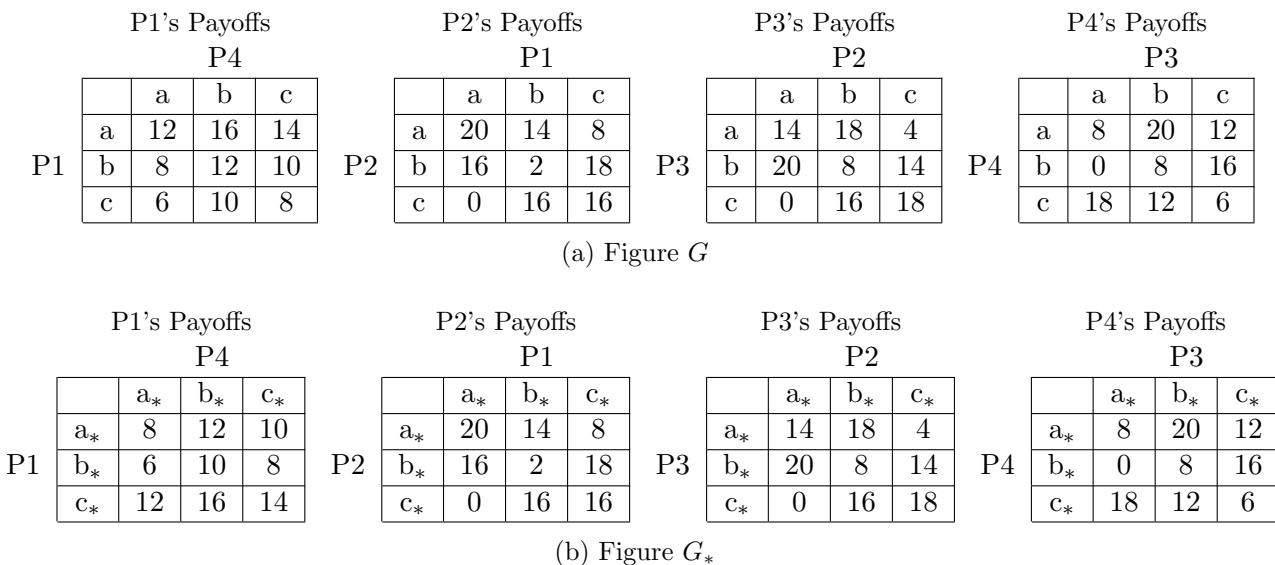


Figure 3.1. Kneeland's (2015) Ring Game

Let us point to two features of the games. First, G and G_* are both dominance solvable. This will be useful for identifying the rationality bound. Second, P1's and P2's payoff matrices are as in the example in Section 2. So, for P1, the payoff matrix in G_* is a relabeling of the payoff matrix in G . P2 has the same payoff matrix across the two games. The same is also true for P3 and P4. This will be useful for identifying the cognitive bound.

Each subject plays both games (G and G_*) in each of the player roles (P1, P2, P3, and P4). As such, an observation consist of a subject's behavior across eight games—that is, an observation is an $x = (x(1), x(2), x(3), x(4))$, where each $x(i) \in \{a, b, c\} \times \{a_*, b_*, c_*\}$ indicates the subject's behavior in the role of P_i across both G and G_* . We assume that each subject is rational (and, so, 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 the cognitive and rationality bounds.

Bounds	P1	P2	P3	P4
Rationality = Cognitive = 1	IU	Constant	Constant	Constant
Rationality = Cognitive = 2	IU	IU	Constant	Constant
Rationality = Cognitive = 3	IU	IU	IU	Constant
Rationality = Cognitive = 4	IU	IU	IU	IU

Table 3.1. Identifying Bounds: No Gap

Table 3.1 provides a preview of how we identify the rationality and cognitive bounds. It focuses on the case where there is no gap between the rationality and cognitive bounds. Observe, if a subject has a rationality bound of m then, for each $i \leq m$, the subject plays the *iteratively undominated (IU) strategy* in the role of P_i —but not in the role of $P(m+1)$. If the subject has a cognitive bound of m then, for each $i > m$, the subject plays a constant strategy profile in the role of P_i —but a non-constant strategy profile in the role of P_m . The next subsections elaborate on the identification strategy.

3.1 Identifying the Rationality Bounds

Return to the example in Section 2: We pointed out that a rational subject in the role of P_1 will play (a, c_*) . This corresponded to the fact that (a, c_*) is a dominant strategy for P_1 . We also pointed out that a rational subject who believes “ P_1 is rational” will play (a, b_*) in the role of P_2 . This corresponded to the fact that (a, b_*) is the only strategy that survives two rounds of iterated dominance for P_2 . Thus, we used iterated dominance to identify the level of reasoning about rationality. This will be the approach that we take more generally.

To better understand what is involved, it will be useful to introduce some terminology: Say a subject is *1-rational* if, in the role of each P_i , she plays a best response given a belief about $P(i-1)$ ’s play of the game. Say a subject is *m-rational* if, in the role of each P_i , she plays a best response given a belief that assigns probability one to the event that $P(i-1)$ is *(m-1)-rational*. There is a tight connection between *m-rationality* and iterated dominance: A subject is *m-rational* if and only if, in the role of each P_i , she plays a strategy that survives m rounds of iterated dominance.¹⁰

Identification (Rationality Bound). *Given an observation $x = (x(1), x(2), x(3), x(4))$, we assign a rationality bound of k if*

(i) $x = (x(1), x(2), x(3), x(4))$ survives k rounds of iterated dominance, and

¹⁰Our definition of *m-rationality* is consistent with formalizations in the epistemic literature. See, e.g., [Tan and da Costa Werlang \(1988\)](#). Using standard results, a subject’s behavior is consistent with *m-rationality* if and only if it survives m rounds of rationalizability ([Bernheim, 1984](#); [Pearce, 1984](#)). See, e.g., [Tan and da Costa Werlang \(1988\)](#), [Battigalli and Siniscalchi \(2002\)](#), amongst others. A strategy survives m rounds of rationalizability if and only if it survives m rounds of iterated strict dominance. See [Pearce \(1984\)](#).

(ii) either $x = (x(1), x(2), x(3), x(4))$ does not survive $(k + 1)$ rounds of iterated dominance or $k = 4$.

A subject’s behavior is identified as having a rationality bound of k if her behavior is consistent with k -rationality and her behavior is inconsistent with $(k + 1)$ -rationality when $k \neq 4$. Note, if the behavior is identified as having a rationality bound of 4, then it is consistent with “rationality and common belief of rationality” (i.e., k -rationality for all k).

In these games, an observation $x = (x(1), x(2), x(3), x(4))$ survives k rounds of iterated dominance if and only if $(x(1), \dots, x(k))$ is IU. Under iterated dominance, P1 would play (a, c_*) , P2 would play (a, b_*) , P3 would play (b, a_*) , and P4 would play (a, c_*) . Table 3.2 then summarizes how we identify the rationality bound.

Bound	P1	P2	P3	P4
Rationality = 1	IU = (a, c_*)	not IU		
Rationality = 2	IU = (a, c_*)	IU = (a, b_*)	not IU	
Rationality = 3	IU = (a, c_*)	IU = (a, b_*)	IU = (b, a_*)	not IU
Rationality = 4	IU = (a, c_*)	IU = (a, b_*)	IU = (b, a_*)	IU = (a, c_*)

Table 3.2. Identifying Rationality Bound

It is important to note that the identification strategy makes use of the subject’s behavior across all player roles. For instance, suppose we observe some $x = (x(1), x(2), x(3), x(4))$, where $x(2) = (b, a_*)$ and $x(4) = (a, c_*)$. 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 who assigns probability one to the event that “P1 is rational.” Thus, we identify the subject’s rationality bound as 1.

3.2 Identifying the Cognitive Bounds

Return to Section 2: We identified P2’s cognitive bound based on whether he played a constant versus a non-constant strategy. (Recall, a *constant strategy* is some strategy (d, d_*) .) In particular, we saw that, if a rational subject in the role of P2 believes “P1 lacks cognition,” the subject will play a constant strategy. On the other hand, if a rational subject in the role of P2 believes “P1 is cognitive,” the subject will play a non-constant strategy.

This is the approach that we will take more generally. Refer to Table 3.3. If we identify a subject as having a cognitive bound of k , then the subject plays a non-constant strategy in the role of P k and a constant strategy in the role of P j for all $j > k$. (Recall, we assume that each subject is rational. Thus, we restrict attention to subjects that play the dominant strategy in the role of P1. This accounts for the P1 column in Table 3.3.)

Bound	P1	P2	P3	P4
Cognitive = 1	Dominant (a, c _*)	Constant	Constant	Constant
Cognitive = 2	Dominant (a, c _*)	Non-Constant	Constant	Constant
Cognitive = 3	Dominant (a, c _*)		Non-Constant	Constant
Cognitive = 4	Dominant (a, c _*)			Non-Constant

Table 3.3. Identifying Cognitive Bound

To identify the cognitive bound, we make two interrelated assumptions. First, we assume that behavior is not an artifact of a subject’s indifference. That is, we assume that no subject—in any player role—is indifferent between any two actions. (Kneeland’s footnote 20 points out that the data from this experiment supports this assumption.¹¹) Because we also assume that each subject is rational, this implies that each subject chooses amongst pure strategies. Second, we assume that each subject believes that other subjects choose pure strategies. Importantly, this does *not* imply that a subject’s beliefs about the strategies of the other players is degenerate. (Below we provide examples of non-degenerate beliefs.)

3.2.1 Cognition

We pointed out that a player may be cognitive but irrational: That is, a player may have a purpose for choosing the action she does, even if she does not play a best response given her belief about play. For instance, in Section 2, we said that a player may adopt a rule of thumb in which she chooses an action that could lead to her lucky number 6 whenever such an action is available. The implication was that such a player would play c in G and the payoff-equivalent action b_* in G_* .

Our analysis focuses on a particular idea of cognition: A cognitive (but potentially irrational) player is a player whose decisions are determined by her payoff matrix and potentially her beliefs about play. For instance, she may adopt a rule of thumb whereby she always chooses the action that generates the highest arithmetic mean. Or, she may adopt a rule of thumb whereby she, first, chooses an action that could lead to a payoff of 6 if such an action exists and, second, if not, she plays a best response given her subjective belief about the play of the game. Or, alternatively, she may adopt the rule of thumb whereby she plays a best response given her subjective beliefs (i.e., she may be rational). Each of these rules of thumb correspond to a theory about how to play the game. In each of the examples, the decision depends on her payoff matrix. In the two latter examples, the decision also depends on her beliefs about play. Each of these rules of thumb selects actions to be played. For instance, the first rule of thumb can select any action that generates the highest mean. We will refer to any action that the rule of thumb can select as a *cognitively optimal* action.

Consider a cognitive P1. An important case is where she has the same belief about P4’s play (across G and G_*) or adopts a rule of thumb that does not depend on her belief about P4’s play.

¹¹Her footnote 20 discusses rationalizable actions, but the same argument applies more broadly in the data.

In that case, if her theory of how to play the game leads her to play a in G , it should lead her to play c_* in G_* . Put differently, if a is cognitively optimal in G , then c_* is also cognitively optimal in G_* . And, likewise, if b (resp. c) is cognitively optimal in G , then a_* (resp. b_*) is also cognitively optimal in G_* .

The same idea applies to $P_i = P_2, P_3, P_4$. Suppose P_i has the same belief about $P(i-1)$'s play or adopts a rule of thumb that does not depend on her belief about $P(i-1)$'s play. Then, if her theory of how to play the game leads her to play $d \in \{a, b, c\}$ in G , it should lead her to play $d_* \in \{a_*, b_*, c_*\}$ in G_* —after all, the two games involve the very same payoff matrix for P_i .

We can abstract a general principle from these two scenarios: For each player P_i there is a permutation Π_i of i 's actions from G to G_* that preserves P_i 's payoff matrix. For $i = 1$, the permutation maps $a \mapsto \Pi_1(a) = c_*$, $b \mapsto \Pi_1(b) = a_*$, $c \mapsto \Pi_1(c) = b_*$; for $i = 2, 3, 4$, the permutation maps each $d \mapsto \Pi_i(d) = d_*$. If a cognitive P_i adopts a rule of thumb that does not depend on her belief about $P(i-1)$'s behavior, then her rationale for choosing action d in G would also serve as a rationale for playing the permuted action $\Pi_i(d)$ in G_* . The same conclusion holds for any cognitive P_i , if she has the same beliefs about $P(i-1)$'s behavior across G and G_* (i.e., if the probability she assigns to d in G is the same as the probability she assigns to $\Pi_{(i-1)}(d)$ in G_*). Put differently, in these cases, P_i 's cognitive optimality is invariant to the permutation of payoff-equivalent action labels.

3.2.2 Reasoning about Cognition and Lack of Cognition

When we identify the subjects' reasoning, we will assume that they are rational—not simply cognitive. Instead, we use the ideas above to restrict the beliefs of a subject. We will think of P_i as having a belief about the strategies of $P(i-1)$ across G and G_* . We write Pr_i for P_i 's distribution on $\{a, b, c\} \times \{a_*, b_*, c_*\}$. So, $\text{Pr}_i(d, e_*)$ is the probability that P_i assigns to $P(i-1)$ playing the strategy (d, e_*) .

We will say that a subject believes cognition if in each player role P_i she satisfies the following principle:

Principle of Belief in Cognition: Suppose P_i believes that “ $P(i-1)$ is cognitive and $P(i-1)$ has the same beliefs about $P(i-2)$'s behavior across G and G_* .” Then, $\text{Pr}_i(d, e_*) > 0$ implies $e_* = \Pi_{(i-1)}(d)$.

Suppose that P_i believes that “ $P(i-1)$ is cognitive and $P(i-1)$ has the same beliefs about $P(i-2)$'s behavior across G and G_* .” Then P_i believes that $P(i-1)$'s cognitive optimality is invariant to the permutation of payoff-equivalent action labels (for $P(i-1)$). The Principle of Belief in Cognition, thus, requires that, if P_i assigns probability p to $P(i-1)$ playing the action d in G , then P_i also assigns probability p to $P(i-1)$ playing the strategy $(d, \Pi_{i-1}(d))$.

To better understand this principle, suppose that P_2 believes that P_1 is cognitive and that P_1 has the same beliefs about P_4 's behavior across G and G_* . P_2 can still have non-degenerate beliefs about P_1 's play. For instance, P_2 may assign probability $1/2$ to P_1 playing a best response and

probability $1/2$ to P1 adopting the lucky-6 rule of thumb. In that case, $\Pr_2(a, c_*) = \Pr_2(c, b_*) = 1/2$. The Principle of Belief in Cognition implicitly requires, however, that P2 has the same belief about the nature of P1's cognition across G and G_* . So, for instance, P2 cannot assign probability 1 to P1 playing a best response in G and probability 1 to P1 playing the lucky-6 strategy in G_* . If P2 had such a belief, he would believe that P1's theory of how to play the game changes across G and G_* , despite the fact that the two games are payoff equivalent (up to the permutation of action labels).¹²

Consider now the case where a subject believes that others lack cognition. To better understand the approach, return to the example (Section 2). We explained that, if P1 lacks cognition, then P1 does not have a theory about how to play the game. As a consequence, 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_* : If P2 assigns probability p to P1 playing d in G , then he also assigns probability p to P1 playing d_* in G_* .

This is the approach we will take more generally. If a subject believes that others lack cognition, then she reasons that the behavior of other subjects does not depend on the details of the game. This implies that, within a given player role, she reasons that the behavior of other subjects does not depend on whether G versus G_* is played. But, within a given a game, it also implies that she reasons that the behavior of other subjects does not depend on the player role. (This is a reasonable assumption in the context of the experiment, where the subject does not observe the identity of her co-players.)

With this in mind, we will say that a subject believes lack of cognition if she satisfies the following principle:

Principle of Belief in Lack of Cognition: The subject has the same belief \Pr in each player role, i.e, $\Pr_i = \Pr$ for each $i = 1, 2, 3, 4$. Moreover, this belief satisfies $\Pr(a, a_*) + \Pr(b, b_*) + \Pr(c, c_*) = 1$.

We will call a belief for P_i , \Pr_i , a *constant belief (for P_i)* if $\Pr_i(a, a_*) + \Pr_i(b, b_*) + \Pr_i(c, c_*) = 1$. The Principle of Belief in Lack of Cognition says that the subject has the same constant belief in each player role.

To better understand this principle, suppose that P2 believes that P1 lacks cognition. P2 can still have non-degenerate beliefs about P1's play. For example, P2 may assign probability $1/2$ to P1 choosing a in G and a_* in G_* , and probability $1/2$ to P1 choosing b in G and b_* in G_* . In that case, $\Pr(a, a_*) = \Pr(b, b_*) = 1/2$. The Principle of Belief in Lack of Cognition implicitly requires, however, that a player has the same belief about the nature of others' lack of cognition across

¹²Another example may be useful: Suppose P3 believes that P2 maximizes her expected payoffs and that P2 assigns probability $5/7 : 2/7$ to $(a, a_*) : (c, c_*)$. Then P3 believes that P2 is indifferent between playing a and b in G . The Principle of Belief in Cognition requires that P3 thinks that the method that P2 uses to resolves this indifference in G gets translated into G_* . For instance, P3 may reason that P2 resolves this indifference in G by choosing the action with the highest maximum payoff (i.e., a); but, if so, then P3 must also reason that P2 resolves this indifference in G_* by choosing the action with the highest maximum payoff (i.e., a_*). If not, P2 would effectively be using a different notion of cognition across the two games.

roles and across games. So, for instance, P2 cannot assign probability 1 to P1 playing a in G and probability 1 to b_* in G_* . Likewise, a subject cannot assign probability 1 to the strategy (a, a_*) in the role of P2, and probability 1 to (b, b_*) in the role of P3. If a subject has such a belief, he would believe that other subjects' behavior depend on the details of the game.

We will use these principles to inductively define k -cognitive (an analogue of m -rationality) and the cognitive bounds. Call a subject *1-cognitive* if, in each player role, she is cognitive. Call a subject *2-cognitive* if she is *1-cognitive* and, in each player role, she satisfies the Principle of Belief in Cognition. Inductively, say a subject is *k-cognitive* if she is $(k - 1)$ -cognitive and she believes (i.e., assigns probability one to the event) that the other player is $(k - 1)$ -cognitive. Say that a subject has a *cognitive bound of 1* if she is 1-cognitive and satisfies the Principle of Belief in Lack of Cognition. Inductively, a subject has a *cognitive bound of k* if she is k -cognitive and believes that other player has a cognitive bound of $(k - 1)$.

Next, we turn to how we identify cognitive bounds. Importantly, when we do so, we assume that each player is rational—not only cognitive. (Thus, we will not make use of 1-cognition as an axiom on behavior.) So, for instance, when we identify a subject as having a cognitive bound of 1, her behavior will be consistent with playing a best response given a belief that satisfies the Principle of Belief in Lack of Cognition. This, of course, implies that her behavior is consistent with being 1-cognitive and satisfying the Principle of Belief in Lack of Cognition.

3.2.3 Identifying the Cognitive Bounds

We now turn to identify the cognitive bounds. To do so, we assume that each subject is rational. We seek to identify the minimum cognitive bound consistent with observed behavior.

Cognitive Bound of 1 Consider a rational subject who has a cognitive bound of 1. By the Principle of Belief in Lack of Cognition, the subject has the same constant belief across player roles. As a consequence, in the roles of $P_i = P_2, P_3, P_4$, this subject must play a constant strategy: If d is a best response for P_i , then d_* is also a best response for P_i . Because we assume that the subject is not indifferent between any two actions, d_* must be her unique best response.

Identification (Cognitive Bound 1). *We identify $(x(1), x(2), x(3), x(4))$ as having a cognitive bound of 1 if there exists a belief Pr on $\{a, b, c\} \times \{a_*, b_*, c_*\}$ so that Pr is a constant belief (in each player role) and each $x(i)$ is a unique best response under Pr .*

If we identify an observation $x = (x(1), x(2), x(3), x(4))$ as having a cognitive bound of 1, then $x(1)$ must be the dominant strategy (a, c_*) . (This is the only strategy that can be a best response to any belief.) So, the observation $(x(1), x(2), x(3), x(4))$ involves behavior in the role of P1 that is not constant. However, in the roles of P2, P3, and P4, this observation involves behavior that is constant. Thus, $x \in \{(a, c_*)\} \times \{(a, a_*), (b, b_*), (c, c_*)\}^3$. Importantly, there may be observations in $\{(a, c_*)\} \times \{(a, a_*), (b, b_*), (c, c_*)\}^3$ that would not be identified as having a cognitive bound of

1. This is because those observations cannot be a best response given a single belief Pr . Table B.1 in Appendix B provides the observations that are identified as having a cognitive bound of 1.

Cognitive Bound of 2 Consider a subject who has a cognitive bound of 2. In the role of each P_i , this subject satisfies the Principle of Belief in Cognition and believes “ $P(i-1)$ is cognitive and believes $P(i-2)$ lacks cognition.” Recall, if $P(i-1)$ believes that “ $P(i-2)$ lacks cognition,” then $P(i-1)$ has the same belief about $P(i-2)$ ’s behavior across G and G_* . Thus, the Principle of Belief in Cognition says that P_i must believe that $P(i-1)$ ’s behavior is invariant to permuting equivalent action labels.

This pins down the subject’s belief about the strategies played. In the role of P_2 , the subject believes that, if P_1 plays a (resp. b , resp. c) in G then she plays $\Pi_1(a) = c_*$ (resp. $\Pi_1(b) = a_*$, $\Pi_1(c) = b_*$) in G_* . Thus, P_2 ’s belief, namely Pr_2 , must satisfy

$$\text{Pr}_2(a, c_*) + \text{Pr}_2(b, a_*) + \text{Pr}_2(c, b_*) = 1.$$

We refer to such a belief as a *2-cognitive belief*. In the role of $P_i \neq P_2$, the subject believes that, if $P(i-1)$ plays a (resp. b , resp. c) in G then she plays $\Pi_{(i-1)}(a) = a_*$ (resp. $\Pi_{(i-1)}(b) = b_*$, $\Pi_{(i-1)}(c) = c_*$) in G_* . Thus, for each $P_i = P_1, P_3, P_4$, Pr_i is a constant belief for P_i .

A rational subject who has a cognitive bound of 2 must play a best response given these beliefs. In the role of P_1 , this implies that she plays the dominant strategy (a, c_*) . In the role of P_2 , this implies that she plays a unique best response given a 2-cognitive belief. In the roles of $P_i = P_3, P_4$, she must play a best response to a constant belief. Observe that, if $d \in \{a, b, c\}$ is a best response (for $P_i = P_3, P_4$) in G , then $d_* \in \{a_*, b_*, c_*\}$ is also a best response (for $P_i = P_3, P_4$) in G_* . Because we assume that no subject is indifferent between any two actions, this implies that the subject must play a constant strategy in the roles of $P_i = P_3, P_4$. With this in mind:

Identification (Cognitive Bound 2). *We identify $(x(1), x(2), x(3), x(4))$ as having a cognitive bound of 2 if*

(i) $x(1) = (a, c_*)$,

(ii) $x(2)$ is a non-constant strategy that is a unique best response under a 2-cognitive belief, and

(iii) $x(3)$ and $x(4)$ are constant strategies.

Suppose that we identify $(x(1), x(2), x(3), x(4))$ as having a cognitive bound of 2. Then $x(3)$ and $x(4)$ must be a constant strategy profile. Moreover, $x(2)$ must be a non-constant strategy profile: If it were constant, then the minimum cognitive bound consistent with the behavior would be 1 and so we would identify the cognitive bound as 1. In addition, $x(2)$ must be uniquely optimal under a 2-cognitive belief. Referring to Table 3.4, there are five non-constant strategies that are optimal under such a belief. (There is one non-constant strategy that is precluded and one constant strategy that is also optimal under a 2-cognitive belief.)

Optimal Given a 2-Cognitive Belief	
Constant	Non-Constant
(a, a_*)	$(a, b_*) (a, c_*)$ $(b, a_*) (b, c_*) (c, a_*)$

Table 3.4. P2’s 2-Cognitive Behavior

Cognitive Bound of 3 Consider a rational subject who has a cognitive bound of 3. As before, this subject must play the dominant (a, c_*) in the role of P1. Thus, we focus on the behavior in the roles of P2, P3, and P4. To identify the behavior, we use the fact that, in each player role P_i , such a subject believes “ $P(i - 1)$ is 2-cognitive and believes $P(i - 2)$ has a cognitive bound of 1.”

In the role of P2, this subject must play a best response given a 2-cognitive belief. The key is that, if the subject has a cognitive bound of 3, then the subject must believe “P1 is cognitive and P1 believes that P4 plays a constant strategy.” (This uses both principles; see Lemma B.1.) Thus, applying the Principle of Belief in Cognition, P2 must believe that P1’s behavior is invariant to the permutation of equivalent action labels and, so, P2 must have a 2-cognitive belief.

In the role of P3, the subject can play any strategy. To understand why, note that the subject believes “P2 is 2-cognitive and believes P1 has a cognitive bound of 1.” This implies that the subject believes “P2 is cognitive and believes that P1 is cognitive,” i.e., that P2 plays a cognitively optimal strategy, given a 2-cognitive belief. Because we have taken a broad view of what preferences cognition might represent, cognitive optimality imposes no restrictions on P2’s behavior when she holds a 2-cognitive belief. Thus, P3 can hold any belief about P2’s behavior and, in turn, any strategy of P3 can be a unique best response.¹³

In the role of P4, the subject must play a constant strategy. To understand why, note that the subject believes “P3 is 2-cognitive and believes P2 has a cognitive bound of 1.” This implies that the subject has a belief that “P3 is cognitive and P3 believes that P2 plays a constant strategy.” (This uses both principles; see Lemma B.1.) Thus, applying the Principle of Belief in Cognition, P4 must believe that P3’s behavior is invariant to the permutation of equivalent action labels and, so, P4 must have a constant belief. Since the subject is not indifferent between any two actions, she plays a constant strategy in the role of P4.

Identification (Cognitive Bound 3). *We identify $(x(1), x(2), x(3), x(4))$ as having a cognitive bound of 3 if*

$$(i) \ x(1) = (a, c_*),$$

¹³We come to the conclusion that any strategy of P3 can be a best response given this belief. We would come to the *same* conclusion even if we imposed strong restrictions on P3’s beliefs about P2’s cognition. For instance, suppose P3 assigns some probability $p \in (0, 1)$ to “P2 is rational and believes that P1 has a cognitive bound of 1” and probability $(1 - p) \in (0, 1)$ to “P2’s preferences do not depend on her beliefs about P1.” In that case, P3’s beliefs can be seen as a convex combination of (i) a belief that assigns probability one to the actions in Table 3.4 and (ii) a constant belief. In that case, all strategies can be a best response for P3, depending on the value of p . (Calculations are available upon request.)

(ii) $x(2)$ is a unique best response given a 2-cognitive belief,

(iii) $x(3)$ is a non-constant strategy, and

(iv) $x(4)$ is a constant strategy.

Let us point to two features of the identification. First, we require that $x(3)$ be non-constant. A constant strategy would be a best response for a P3 that believes “P2 is 2-cognitive and believes P1 has a cognitive bound of 1.” However, because we focus on the minimum cognitive bound consistent with observed behavior, we would assign such an observation—i.e., an observation with both $x(3)$ and $x(4)$ constant—a lower cognitive bound. Second, unlike how we identify a cognitive bound of 2, we do not require that $x(2)$ is non-constant. In identifying the cognitive bound of 2, we only required that $x(2)$ is non-constant to ensure that we could not identify the observation as having a cognitive bound of 1.

Cognitive Bound of 4 Consider a rational subject who has a cognitive bound of 4. In the roles of P1, P2, and P3, this subject’s behavior is observationally equivalent to the behavior of a subject with a cognitive bound of 3. (Simply repeat the arguments above, replacing Lemma B.1 with Lemma B.2.) The difference comes in behavior in the role of P4. Now, in the role of P4, the subject believes “P3 is 3-cognitive and believes P2 has a cognitive bound of 2.” This implies that the subject has a belief that “P3 is cognitive and P3 can hold any belief about P2’s play.” (To see this, repeat the argument for P3’s behavior when there is a cognitive bound of 3.) Thus, P4 can hold any belief about P3’s behavior and, so, any strategy can be a best response.

Identification (Cognitive Bound 4). *We identify $(x(1), x(2), x(3), x(4))$ as having a cognitive bound of 4 if*

(i) $x(1) = (a, c_*)$,

(ii) $x(2)$ is a unique best response given a 2-cognitive belief,

(iii) $x(4)$ is a non-constant strategy.

Let us point to two features of the identification. First, we require $x(4)$ to be non-constant. As above, a constant strategy would be a best response for a P4 that believes “P3 is 3-cognitive and believes that P2 has a cognitive bound of 2.” However, because we focus on the minimum cognitive bound consistent with observed behavior, we would assign such an observation—i.e., an observation with $x(4)$ constant—a lower cognitive bound. Second, we do not require that $x(3)$ is non-constant. This is no longer required to ensure that that we cannot assign the observation a lower cognitive bound. (And, similarly, for $x(2)$.)

We conclude with an important comment about the experimental design.

Remark 3.1. Kneeland (2015) introduced the permuted ring games to identify the rationality bounds. In fact, the nature of the particular permutations also make the permuted ring games well

suited to identify the cognitive bounds. Notice that, for P1, the permutation is non-constant, i.e., it does not map *any* action d in G to the associated action d_* in G_* . However, for P2, P3, and P4, the permutation is constant, i.e., it maps *every* action d in G to the associated action d_* in G_* . These differences in the permutations are important for identifying the cognitive bound. If all players had a constant permutation, then we would expect constant behavior across the ring games, independent of how the subject reasons about cognition. Because P1’s permutation is non-constant, we can separate a cognitive bound of 1 from a cognitive bound of $k = 2, 3, 4$. If, on the other hand, P2’s permutation were also not constant, then a subject who has a cognitive bound of 1 would also play a non-constant strategy in the role of P3. This would, presumably, conflate the behavior of subjects with a cognitive bound of 1 and subjects with a cognitive bound of 2. And similarly for P3 and P4.

3.3 Identification: Wrap-Up

Table 3.5 summarizes Table 3.2 (identification of the rationality bound) and Table 3.3 (identification of the cognitive bound). The rows in grey correspond to the case where there is no gap between the identified rationality and cognitive bounds (c.f. Table 3.1.)

Rat	Cog	P1	P2	P3	P4
1	1	IU	constant	constant	constant
1	2	IU	not “constant or IU”	constant	constant
1	3	IU	not IU	not constant	constant
1	4	IU	not IU		not constant
2	2	IU	IU	constant	constant
2	3	IU	IU	not “constant or IU”	constant
2	4	IU	IU	not IU	not constant
3	3	IU	IU	IU	constant
3	4	IU	IU	IU	not “constant or IU”
4	4	IU	IU	IU	IU

Table 3.5. Identified Bounds

Consider an observation $(x(1), x(2), x(3), x(4))$ identified as having a cognitive bound of k . Then, $x(k)$ is non-constant and $x(i)$ is constant, for all $i > k$. When there is no gap between the identified rationality and cognitive bounds—that is, if the rationality bound is also identified as k — $x(i)$ is IU for all $i \leq k$. When there is a gap between the identified rationality and cognitive bounds, there is some $m < k$ so that $x(i)$ is IU for all $i \leq m$, and $x(m + 1)$ is not IU.

There are several subtleties obscured by Table 3.5 (resp. Table 3.3). First, to identify the observation as having a cognitive bound of 1, $x(2)$, $x(3)$, and $x(4)$ must be a unique best response under the *same* constant belief. (This, for instance, rules out an observation with $x(2) = (b, b_*)$ and $x(3) = x(4) = (a, a_*)$.) Second, to identify the observation as having a cognitive bound of $k = 2, 3, 4$, $x(2)$ must be a best response under a 2-cognitive belief. (This, for instance, rules out

Cognitive Bound	Potential Observations	Subjects
1	13	9
2	35	16
3	108	21
4	324	28
NC Rational	249	1
NC Irrational	5832	5
Total	6561	80

Table 4.1. Inferring the Minimum Cognition from Observed Behavior

an observation with $x(2) = (c, b_*)$.) As a consequence, there are strategy profiles that satisfy the conditions in Table 3.5 but are not classified under our identification strategy. See Table B.1.

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 eight 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 concerns.

There are 80 subjects and thus 80 observations $x = (x(1), x(2), x(3), x(4))$. Table 4.1 shows the identified cognitive bounds. There are 6071 potential observations that could lead to a non-classification: 5832 potential observations could lead to non-classification because they involve a dominated strategy in the role of P1, and 249 potential observations could lead to non-classification because they are inconsistent with our identifying assumptions. (Refer back to the discussion in Section 3.3.) In the data, five observations are not classified because they involve a dominated choice in the role of P1; these are labeled “NC Irrational” in Table 4.1. These five fall outside the purview of our analysis. Our analysis thus focuses on the behavior of the remaining 75 subjects. Of those subjects, 1 subject is not classified; this is labeled “NC Rational” in Table 4.1. (We include this subject in our analysis. In principle, the subject’s behavior is rational—the behavior is only ruled out by the assumptions we have made about the players’ beliefs.)

Refer to Table 4.1. More than 37% of the subjects 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 bound consistent with the observed data. This has two implications. First, subjects identified as having a cognitive bound of 4 may in fact have a higher level of cognition. (Given the nature of the 4-player ring game, we cannot distinguish a bound of four from higher cognitive bounds.) Second, subjects identified as having a cognitive bound of 1 may actually have a higher level of cognition. If so, their behavior would indicate a gap between reasoning about

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

Table 4.2. Gap Between Cognition and Rationality

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.)

Table 4.2 provides information about the gap between the cognitive and the rationality bounds. If there were no gap, 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.

We point to two specific features. First, 47% of the subjects identified as having a low rationality bound—i.e., a rationality bound of 1 or 2—have a higher cognitive bound. (The same is not true for those with a rationality bound of 3.) Second, the gap appears more pronounced at higher levels of cognition. In particular, 54% of the subject identified as having a cognitive bound of 4 have a rationality bound that is strictly less than 4.¹⁴ Put differently, only 46% of the subjects identified as having a cognitive bound of 4 also have a rationality bound of 4. But, 76% (resp. 81%) of the subjects who have a cognitive bound of 3 (resp. 2) also have a rationality bound of 3 (resp. 2).

Table 4.3 shows the average payoffs that a subject with a given cognitive bound could receive. These are computed by pairing each subject’s observed behavior with the observed behavior of every other subject. Thus, it is what the subject should expect to earn, if she has correct beliefs about the behavior of the population. Subjects identified as having a cognitive bound of 2, 3 or 4 can have a gap between their cognitive and rationality bounds. Table 4.3 shows that such gap subjects outperform their non-gap counterparts.

The fact that gap subjects outperform their non-gap counterparts is of interest. As previously suggested, bounds on reasoning about rationality are often interpreted as limits on the players’ sophistication—e.g., limits on their ability to engage in interactive reasoning. However, such bounds may instead reflect a strategic choice to *not* believe other players are rational (or engage in reasoning about rationality). The fact that the gap subjects outperform their non-gap counterparts suggests

¹⁴Eight of those subjects play IU in the role of P1-P2-P4 and (b, b_*) in the role of P3. That behavior can, for instance, be explained as a subject who, in each player role, assigns probability $\frac{3}{4} : \frac{1}{4}$ to “other subjects are rational” vs. “other subjects choose the action with the highest arithmetic mean.” Such first-order beliefs can be consistent with identifying the subject as having a cognitive bound of 4. Under such first-order beliefs, the subjects’ rationality bound would be lower than our identified rationality bound of 2; in that case, the identified gap is smaller than the subjects’ actual gap.

Cognitive Bound				
	1	2	3	4
No Gap	12.97	11.71	11.84	9.78
Gap	-	12.52	13.04	12.05

Table 4.3. Average Payoffs

that subjects are capable of reasoning higher levels about rationality, but simply choose not to do so.

5 Reasoning about Rationality

In Section 2, we pointed out that, if P2 assigns probability $4/5$ to P1 playing the rational strategy (a, c_*) and probability $1/5$ to the cognitive but irrational strategy (b, c_*) , then P2’s best response is to play the non-IU strategy (a, c_*) . Note two features of this example. First, P2 was identified as having a cognitive bound that was higher than his rationality bound. Second, P2 had non-degenerate beliefs about P1’s rationality.

In this section, we focus on subjects who have a gap between their rationality and cognitive bounds. We illustrate how such a subject’s beliefs (about strategies) can be reinterpreted in terms of beliefs about rationality—or, more loosely, in terms of reasoning about rationality—even when such a subject is *not* playing an IU strategy. In Section 6, we use this to show that subjects’ behavior can be explained using a limited set of “types,” each of whom have non-degenerate beliefs about rationality.

Recall, to identify the rationality bound, we used the behavior in the role of P_i to distinguish a subject who has a rationality bound of $(i - 1)$ from a subject who has a rationality bound of $m \geq i$. With this in mind, we begin by focusing on behavior in a single player role P_i : We focus on the case where the subject has a rationality bound of $m = i - 1$ and a cognitive bound of $k \geq i$. We first seek to understand what the behavior in the role of P_i tells us about reasoning about rationality. We then use the behavior across player roles to provide a more complete analysis on reasoning about rationality.

Player Role P2 Suppose that we identify a subject as having a rationality bound of 1 and a cognitive bound of $k \geq 2$. In that case, the subject plays some $x(2) = (d, e_*)$ in the role of P2 that is optimal under a 2-cognitive belief but it is not the IU strategy. For instance, in Section 2, we pointed out that the strategy (a, c_*) is rational for P2 if she assigns probability $4/5$ to P1 playing the rational strategy (a, c_*) and probability $1/5$ to the cognitive but irrational strategy (b, c_*) . If, in fact, P2 holds this belief, then she assigns probability $p_2 = 4/5$ to the event that “P1 is rational.”

More generally, the probability that P2 assigns to P1 playing the dominant strategy (a, c_*) is the probability that P2 assigns to the event that “P1 is rational.” Thus, for any observed strategy $x(2)$ that is optimal under a 2-cognitive belief, we can find the set of probabilities p_2 so that $x(2)$ is

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

Table 5.1. Bounds on P2’s Belief p_2 of Rationality Given Cognition Level ≥ 2

a unique best response under some 2-cognitive belief that assigns probability p_2 to the event that “P1 is rational.” These sets are given in Table 5.1.

Player Role P3 Next, consider a subject who is identified as having a rationality bound of 2 and a cognitive bound of $k \geq 3$. In that case, the subject plays some $x(3) = (d, e_*)$ in the role of P3 that is not the IU strategy, but is optimal under some belief Pr_3 about P2’s behavior. Since we have taken a broad view of cognition, Pr_3 can be any belief.

Suppose that we observe P3 play $x(3) = (b, c_*)$. This behavior is a best response under a belief Pr_3 with $\text{Pr}_3(a, c_*) = 1$. Referring to Table 5.1, this belief can be reconceptualized as a belief that assigns probability $p_3 = 1$ to “P2 is rational and assigns probability at least q_2 to P1’s rationality,” for any $q_2 < 7/8$. But, it is also a best response to a belief Pr'_3 with $\text{Pr}'_3(a, b_*) = \text{Pr}'_3(c, c_*) = 1/2$. This belief can be reconceptualized as a belief that assigns probability $p'_3 = 1/2$ to “P2 is rational and assigns probability at least 1 to P1’s rationality.” (Of course, it can also be reconceptualized as a belief that assigns probability $p'_3 = 1$ to “P2 is rational and assigns probability at least 0 to P1’s rationality.”)

This idea applies more generally. Any observed strategy $x(3) = (d, e_*)$ is a unique best response under a belief Pr_3 . Any such belief can be reconceptualized in terms of non-degenerate beliefs about reasoning about rationality. To do so, say that an event is q -believed if the event is assigned probability at least q . Then, the observed strategy can be viewed as a best response given a belief that assigns probability p_3 to

“P2 is rational and q_2 -believes that P1 is rational.”

As the examples highlight, the choice of a pair (p_3, q_2) will, quite generally, not be unique. Table C.2 in Appendix C provides the pairs of (p_3, q_2) that rationalize each observation.

Player Role P4 Finally, consider a subject who is identified as having a rationality bound of 3 and a cognitive bound of 4. In that case, the subject plays some $x(4) = (d, e_*)$ in the role of P4 that is not the IU strategy, but is optimal under some belief Pr_4 about P3’s behavior. Much as in above, such a belief can be formulated in terms of non-degenerate beliefs about reasoning about rationality. Specifically, it can now be reconceptualized as a belief that assigns probability p_4 to

“P3 is rational and q_3 -believes that ‘P2 is rational and r_2 -believes rationality’ .”

(Again, the choice of a triple (p_4, q_3, r_2) is not unique.) Table C.3 in the appendix presents the triples (p_4, q_3, r_2) that rationalize each strategy observed in the data.

Cross-Player Role Restriction: The Anonymity Assumption We can use the subjects’ behavior across player roles to obtain additional information on reasoning about rationality. To exemplify the approach, suppose we observe a subject play a non-IU strategy $x(2) = (b, a_*)$ in the role of P2 and an IU strategy $x(4) = (a, c_*)$ in the role of P4. If we only observed the subject’s behavior in the role of P4, then we would conclude that the subject’s behavior is consistent with her playing a best response given a belief that assigns $p_4 = 1$ to

“P3 is rational and 1-believes that ‘P2 is rational and 1-believes rationality.’”

So, certainly $x(4)$ is consistent with the subject playing a best response to a belief that assigns probability $p_4 = 1$ to “P3 is rational.” Arguably, if the subject assigns probability 1 to the event “P3 is rational,” then she should be prepared to assign probability 1 to the event that “P1 is rational.” However, referring to Table 5.1, $x(2)$ is only a best response to a belief that assigns probability $p_2 \leq 2/5$ to P1’s rationality. Thus, the fact that p_2 is at most $2/5$ implies that p_4 is also at most $2/5$.

More generally, a subject’s behavior across player roles can be used to obtain additional information about how she reason about rationality. We assume:

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

- (i) *A subject assigns probability p to the event “ P_i is rational” if and only if she assigns probability p to the event “ P_j is rational.”*
- (ii) *A subject assigns probability p to the event “ P_i is rational and q -believes P_k is rational” if and only if she assigns probability p to the event “ P_j is rational and q -believes P_ℓ is rational.”*

For the remainder of the paper, we assume the Anonymity Assumption (AA). The AA is natural since, in the experiment, the subjects are anonymous and there is a certain symmetry to the game. As suggested above, under the AA, the behavior of the subject in the role of P_j provides a bound on her reasoning about rationality in the role of P_i , for $i > j$. To see this, notice that an observation identified as having a cognitive bound of 3 (resp. 4) can be associated with numbers $(p_2; p_3, q_2)$ (resp. $(p_2; p_3, q_2; p_4, q_3, r_2)$) as above. The AA implies:

Proposition 5.1. *Fix a subject whose cognitive bound is $k \in \{3, 4\}$. Then, $p_2 \geq \max\{p_3, p_4\}$. Moreover, if $k = 4$ and $q_3 \geq q_2$, then $p_3 \geq p_4$.*

6 Deliberate Choice or Errors?

In Section 4, we argued that there can be 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 do so, we estimate three models, corresponding to the two view points on the off-diagonal entries. We refer to the first model as the Deliberate Choice model; it corresponds to the analysis in this paper. We refer to the other models as the Random Choice and the Logistic Choice models; they correspond to the alternate view of the off-diagonal entries. We argue that the data is not well explained by the Random or Logistic Choice models.¹⁵ In Section 7, we use the best fitting Deliberate Choice model to argue that subjects identified as having a gap between their rationality and cognitive bounds have non-degenerate beliefs about rationality.

Write $n = 1, \dots, N$ for a subject in the subject pool. In our dataset, $N = 75$. For each of the models, we will match subjects to “types of reasoners” by maximizing the log likelihood of observing the subject’s behavior (across both games and player roles).

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, in each player role $P_i = P_2, \dots, P_k$, the subject’s behavior can be reconceptualized as a best response to reasoning about rationality. Specifically, if the subject is identified as having a cognitive bound of 4, the subject’s beliefs can be reconceptualized as beliefs that satisfy the following P2-P3-P4 requirements:

P2 Requirement The subject assigns probability p_2 to the event that “P1 is rational.”

P3 Requirement The subject assigns probability p_3 to the event that “P2 is rational and q_2 -believes that ‘P1 is rational.’”

P4 Requirement The subject assigns probability p_4 to the event that “P3 is rational and q_3 -believes that ‘P2 is rational and r_2 -believes rationality.’”

And, analogously, if the subject is identified as having a cognitive bound of 3 (resp. 2), the subject’s beliefs can be reconceptualized as beliefs that satisfy the P2-P3 (resp. P2) Requirements (resp. Requirement).

In the Deliberate Choice model, we will think of a subject as characterized by parameters satisfying the P_i Requirements. To simplify the analysis, we assume that the subjects’ beliefs correspond to what they believe about the population’s beliefs.

Assumption 6.1 (Population Assumption).

¹⁵The Random Choice and Logistic Choice models are in the spirit of [Kneeland’s \(2015\)](#) approach. In particular, [Kneeland](#) assumes that the rationality bound is determined by the cognitive bound and any departures from IU are due to mistakes. She does not model the mistakes explicitly but instead classifies a subject’s rationality/cognitive bound based on its closeness to the IU profile. The Random Choice and Logistic Choice models below will classify subjects based on an explicit model of mistakes.

(i) If a subject has a cognitive bound of 3, then the subject satisfies the P2 and P3 Requirements for some $p_2 = q_2$.

(ii) If a subject has a cognitive bound of 4, then the subject satisfies the P2, P3, and P4 Requirements for some $p_2 = q_2 = r_2$ and $p_3 = q_3$.

We will show that, under the Population Assumption (PA), the Deliberate Choice model outperforms the other models. As a consequence, the same will be true absent the PA.

To understand what the 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 the event that “P1 is rational.” The PA assumption says that, in the role of P3, she acts as if she assigns probability p_3 to the event that “P2 is rational and q_2 -believes ‘P1 is rational.’” Thus, the subject’s beliefs about the population’s beliefs about “P1’s rationality,” i.e., q_2 , are determined by her own beliefs about “P1’s rationality,” i.e., p_2 . And so on.

From now on, we assume the PA (in addition to the AA). This implies that a subject identified as having a cognitive bound of 4 (resp. 3) satisfies the P2-P3-P4 Requirements for some triple (p_2, p_3, p_4) (resp. pair (p_2, p_3)).

The Deliberate Choice model views a subject’s *type* as reflecting both beliefs about rationality and a cognitive bound. As such, a type can be characterized by a triple in $([0, 1] \cup \{\text{cb}\})^3$. A type with a cognitive bound of 4 corresponds to a triple (p_2, p_3, p_4) that satisfies the P2-P3-P4 Requirements. A type with cognitive bound of 3 corresponds to a triple (p_2, p_3, cb) , where (p_2, p_3) satisfy the P2-P3 Requirements and cb indicates that the cognitive bound binds in the role of P4. A type with cognitive bound of 2 corresponds to a triple $(p_2, \text{cb}, \text{cb})$, where p_2 satisfies the P2 Requirement. A type with cognitive bound of 1 corresponds to the triple $(\text{cb}, \text{cb}, \text{cb})$. (By the AA, $p_2 \geq \max\{p_3, p_4\}$.)

An observation is an action profile for subject n , i.e., some $x_n = (x_n(1), \dots, x_n(4))$, where $x_n(i) \in \{a, b, c\} \times \{a_*, b_*, c_*\}$ denotes the behavior in the role of P_i . For a given type indexed t^ℓ , write $A^\ell \subseteq (\{a, b, c\} \times \{a_*, b_*, c_*\})^4$ for the set of strategy profiles that are a best response for type t^ℓ . Informally, A^ℓ is the set of strategies that can be played by type t^ℓ . For example, if $t^\ell = (1, \text{cb}, \text{cb})$, then $|A^\ell| = 9$ and, if $t^\ell = (0.8, \text{cb}, \text{cb})$, then $|A^\ell| = 18$. (See Table 5.1.)

Write \mathbb{T}^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 $(9/10, 9/10, 9/10)$. So, the types in \mathbb{T}^D can be partitioned into a finite number of subsets, so that types $t^\ell, t^{\ell'}$ are in the same partition member if and only if $A^\ell = A^{\ell'}$. With this in mind, it suffices to restrict attention to finite subsets of \mathbb{T}^D and we use the phrase ‘type’ as a label for an equivalence class. We look for the minimal set of types that rationalize the data.

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

In the DC model, the probability of observing x_n depends on the distribution of types (i.e., π) and the likelihood that each type t^ℓ plays the action profile x_n . Thus, the likelihood depends on the set of strategies that can be played by type t^ℓ , A^ℓ . We assume that each strategy in A^ℓ is equally likely to be played by type t^ℓ . The econometric model \mathcal{M}^D allows for type-specific noise. In particular, type t^ℓ follows the predicted strategies A^ℓ with probability $1 - \varepsilon^\ell$ and deviates from the predicted strategies with probability ε^ℓ . When type t^ℓ deviates from the predicted strategies, type t^ℓ plays each strategy in $A \setminus A^\ell$ with equal probability.

With this, the probability that observation x_n was generated by type t^ℓ given ε^ℓ is

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

The likelihood of observing x_n in model \mathcal{M}^D is then

$$\mathcal{L}_n(x_n; \mathcal{M}^D) = \sum_{t^\ell \in T^D} \pi(t^\ell) p(x_n, \varepsilon^\ell | t^\ell).$$

The aggregate log-likelihood of observing the experimental data, $\mathbf{x} = (x_n)_{n=1}^N$, is

$$\ln \mathcal{L}(\mathbf{x}; \mathcal{M}^D) = \sum_{n=1}^N \ln \mathcal{L}_n(x_n, \varepsilon, \mathcal{M}^D).$$

We can always maximize $\ln \mathcal{L}(\mathbf{x}; \mathcal{M}^D)$ by setting $T^D = \mathbb{T}^D$. To limit overfitting of the data, we choose amongst models $\mathcal{M}^D = (T^D, \pi, \varepsilon)$ by penalizing a model for having a larger set of types. We do so by using the Bayesian Information Criterion (BIC). For a given model \mathcal{M}^D , the BIC is given by

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

where f is the number of free parameters in \mathcal{M}^D . Note $f = 2 \cdot |T^D| - 1$ since each additional type adds 2 parameters, $(\pi(t^\ell), \varepsilon^\ell)$. We choose $\hat{\mathcal{M}}^D$ to maximize $\text{BIC}(\mathcal{M}^D)$.

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 the Random and Logistic Choice 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. So, in these models, the subject who actually has a cognitive bound of 3 can play the non-IU strategy (a, c_{*}) in the role of P2—but only if she trembles. If these models are correct, the DC model would misidentify this behavior by P2 as reflecting a gap between the cognitive and rationality bounds.

The Random and Logistic Choice models differ in how trembles are modeled. In the Random Choice model, if a subject trembles in a given game, she plays the remaining actions with equal probability. In the Logistic Choice model, the probability of playing an action depends on the logistic best response function. Thus, when a subject trembles in a given game, she is more likely to play the action that gives a higher expected payoff.

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

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

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

In each of these models, the researcher cannot infer that the subject's cognitive bound must be k or higher. (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 $x_n = (x_n(1), \dots, x_n(4))$. In particular, 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 $x^{\text{rat}}(i)$ for the IU strategy in the role of P_i . (So, $x^{\text{rat}}(1)$ is (a, c_*) , $x^{\text{rat}}(2)$ is (a, b_*) , etc.) Let $\eta[x_n, i]$ be the number of coordinates for which $x_n(i)$ and $x^{\text{rat}}(i)$ agree. So if $x_n(i) = x^{\text{rat}}(i)$, then $\eta[x_n, i] = 2$; if $x_n(i)$ and $x^{\text{rat}}(i)$ specify the same behavior in G but not G_* , then $\eta[x_n, i] = 1$. And so on. Write $x_g^{\text{rat}}(i)$ for action specified by $x^{\text{rat}}(i)$ 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.¹⁶ Thus, the probability of observing $x_n(4)$ is $(\frac{1}{3})^2$. In the role of P3, the subject plays the IU strategy (b, a_*) up to trembles. The RC 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 IU action 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_n in the model \mathcal{M}^R given a subject of type c^k is

$$p(x_n, \varepsilon^k | c^k) = \left(\frac{1}{3}\right)^{2(4-k)} \left(1 - \varepsilon^k\right)^{\sum_{i=1}^k \eta[x_n, i]} \left(\frac{\varepsilon^k}{2}\right)^{\sum_{i=1}^k (2 - \eta[x_n, i])}.$$

¹⁶Appendix F provides an alternate version of the RC model where, modulo errors, the subject instead randomizes equally amongst the constant strategies. (It explains why this is a desirable specification.) The results for that specification of the RC model are analogous to those presented in the main text.

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

$$\mathcal{L}_n(x_n, \mathcal{M}^R) = \sum_{c^k \in T^R} \pi(c^k) p(x_n, \varepsilon^k | c^k).$$

And, the aggregate log-likelihood of observing the experimental dataset $\mathbf{x} = (x_n)_{n=1}^N$ is

$$\ln \mathcal{L}(\mathbf{x}, \mathcal{M}^R) = \sum_{n=1}^N \ln \mathcal{L}_n(x_n, \mathcal{M}^R).$$

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

Logistic Choice Model Consider again a subject whose cognitive bound is $k = 3$. As in the RC model, the probability of observing $x_n(4)$ 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 $d \in \{a, b, c\}$, write $EP_G(d|i)$ for the expected payoff from playing d in the role Pi , when Pi expects $P(i-1)$ to play the IU strategy. The logistic best response function for G is given by $\sigma_G : \mathbb{R}_+ \times \{1, 2, 3, 4\} \rightarrow \Delta(\{a, b, c\})$, where

$$\sigma_G(\lambda, i)(d) = \frac{\text{Exp}(\lambda \cdot EP_G(d|i))}{\sum_{e \in \{a, b, c\}} \text{Exp}(\lambda \cdot EP_G(e|i))},$$

where λ specifies a precision parameter. Define $EP_{G_*}(d_*|i)$ and the logistic best response $\sigma_{G_*} : \mathbb{R}_+ \times \{1, 2, 3, 4\} \rightarrow \Delta(\{a_*, b_*, c_*\})$ analogously.

For each player role Pi , 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 IU strategy. 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 . If λ is high then trembles are less likely than when λ is close to 0. In particular, observe that $\lim_{\lambda \rightarrow \infty} \sigma_G(\lambda, i)(d)$ is 1 if $d = x_G^{\text{rat}}(i)$ and is 0 otherwise.

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

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

where $x_n^G(i)$ and $x_n^{G_*}(i)$ are the actions that subject n chooses in the role of Pi in G and G_* ,

respectively. Then, the likelihood of observing behavior x_n in the model \mathcal{M}^L is

$$\mathcal{L}_n(x_n, \mathcal{M}^L) = \sum_{c^k \in T^L} \pi(c^k) p(x_n, \lambda^k | c^k).$$

And, the aggregate log-likelihood of observing the experimental dataset $\mathbf{x} = (x_n)_{n=1}^N$ is

$$\ln \mathcal{L}(\mathbf{x}, \mathcal{M}^L) = \sum_{n=1}^N \ln \mathcal{L}_n(x_n, \mathcal{M}^L).$$

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

6.3 Model Selection

Table 6.1 compares the DC and RC models. The first three columns contain the estimates of the DC model selected by the BIC, the middle three columns correspond to the RC model, and the last three columns correspond to the LC model.

Deliberate Choice			Random Choice			Logistic Choice		
$(\mathbf{p}_2, \mathbf{p}_3, \mathbf{p}_4)$	π	ε	\mathbf{c}^k	π	ε	\mathbf{c}^k	π	λ
(1, 1, 1)	.13	.006	c^4	.41	.105	c^4	.16	1.54
(1, 1, $\frac{4}{9}$)	.05	.027	c^3	.18	0	c^3	.24	∞
($\frac{7}{8}, \frac{5}{6}, \frac{7}{8}$)	.16	.015	c^2	.20	0	c^2	.34	1.90
($\frac{2}{7}, \frac{2}{7}, \frac{2}{7}$)	.03	.018	c^1	.21	.005	c^1	.26	∞
(1, 1, cb)	.21	.007						
($\frac{2}{5}, \frac{2}{5}, \text{cb}$)	.07	.012						
(1, cb, cb)	.13	.012						
($\frac{7}{8}, \text{cb}, \text{cb}$)	.08	.052						
(cb, cb, cb)	.12	.003						
Neg. Log-Likelihood	294.78			356.07			364.18	
BIC	662.73			742.36			758.58	
AIC	623.56			726.14			742.36	

Table 6.1. Deliberate Choice model versus Random Choice models

According to BIC, the DC model that best fits the data has nine types. We label the types, so that they correspond to the supremum over all behaviorally equivalent probabilities. For instance,

the type $(\frac{7}{8}, \frac{5}{6}, \frac{7}{8})$ corresponds to a continuum of types (p_2, p_3, p_4) with $p_2 \in [5/7, 7/8]$, $p_3 \in [15/62, 5/6]$, $p_4 \in [0, 7/8]$, and $p_2 \geq p_3$. This can be thought of as a situation in which P4 assigns probability p_4 to “P3 is rational and p_3 -believes that ‘P2 is rational and p_4 -believes rationality,’” where the upper bounds on p_2 and p_4 is $7/8$ and the upper bound on p_3 is $5/6$. Thus, our choice of labels provides some indication on the size of the gap between the cognitive and rationality bounds.¹⁷

The DC model explains the data better than both the RC model and the LC model according to model fit: The (negative) log-likelihood for the DC model is 294.78, while it is 356.07 for the RC model and 364.18 for the LC model. To account for the fact that the DC model has more parameters than both the RC model and the LC model, we can use the BIC and the Akaike Information Criterion (AIC). Under both these criteria, the DC model outperforms the other two models.

To test whether one model provides a significantly better fit of the data than another, we use the Vuong test.¹⁸ When comparing the DC model to the RC model, the DC model provides a significantly better fit of the data than the RC model at all standard significant levels (Vuong test statistic = 2.62; p-value=.009). The same holds true for the comparison between the DC model and the LC model (Vuong test statistic = 3.42; p-value=.001). The RC model also provides a significantly better fit of the data than the LC 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, they analyzed the distributions of types that best fit behavior in the data. This subsection estimates an individual-level econometric model, in which (in principle) each observation is described by a distribution of types and so each observation can be associated with a given type in the model. There are two reasons to estimate the individual-level models. First, doing so serves as a robustness check. It verifies that the distribution associated with the frequency of estimated types corresponds to the estimated aggregate analysis in Section 6.3. Second, and more importantly, it allows us to address an important model-selection question: Can a model driven by random choice generate the distribution of DC types observed in the data? We address this question in Section 6.5.

The individual-level analysis will be based on the econometric framework in Sections 6.1-6.2. The key difference is that now, the distributions π are subject specific (now π_n) and the noise/errors are also subject specific (now ε_n). We focus on the Deliberate Choice and Random Choice estimates. (Because the RC model outperforms the LC model in the aggregate analysis, we omit discussion of the latter.)

Deliberate Choice Estimation We choose each $(\hat{\pi}_n, \hat{\varepsilon}_n)$ to maximize $\ln \mathcal{L}_n(x_n, \mathcal{M}^D)$. As a consequence, we will estimate the subject-specific noise to be zero provided we can rationalize the

¹⁷This is why we focus on the supremum of the continuum of types that rationalize the data.

¹⁸The Vuong test is used to test whether one of two non-nested models provides a significantly better fit of the data. Our three models are all non-nested.

subject’s behavior by the DC model.¹⁹ That is, if there exists a type t^ℓ with $x_n \in A^\ell$, then $\hat{\varepsilon}_n$ must be 0. Moreover, since we have subject specific distributions, each $\hat{\pi}_n$ is degenerate on a particular type. Thus, we can view $\hat{\pi}_n$ itself as a type categorizing subject n . As such, we equate subject x_n with the type t^ℓ so that $\hat{\pi}_n(t^\ell) = 1$.

We estimate the DC model that corresponds to the model in Section 6.3. Specifically, we focus on models associated with the same nine 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 $(\frac{2}{7}, \frac{2}{7}, \frac{2}{7})$ with noise.²⁰ In light of this, we think of this model as one in which there are effectively nine *nice types*—corresponding to the analysis in Section 6.3—and one *error type*.

Random Choice Estimation We choose each $(\hat{\pi}_n, \hat{\varepsilon}_n)$ to maximize $\ln \mathcal{L}_n(x_n, \mathcal{M}^R)$. Again, for each subject n , $\hat{\pi}_n$ is degenerate on some type; we use that type to classify the subject. Notice that, if x_n is consistent with 4-rationality, subject n will be assigned to the type c^4 and the error will be estimated as 0 (i.e. $\hat{\pi}_n(c^4) = 1$ and $\hat{\varepsilon}_n = 0$). Similarly, for a subject whose action profile is consistent with m -rationality. If a subject is assigned to a type with an error (i.e., $\hat{\varepsilon}_n > 0$), then it is more likely that x_n was generated by a type with a higher cognitive bound who made a mistake vs. a type with a lower cognitive bound who did not make a mistake.

Estimates Appendix E provides the estimates of the individual analysis. For the DC model, the individual analysis tells a similar story as the aggregate analysis. In particular, the estimated proportion of each type is close to the estimated likelihood in the aggregate analysis. The aggregate analysis suggests that, when the cognitive bound is $k = 2, 4$, there is a slightly larger fraction of subjects that have a gap between the cognitive and rationality bounds. But the differences are not large. For the RC model, the differences are more pronounced. In particular, under the individual analysis, the estimated proportion of types places less weight on c^4 types and more weight on c^3 types. This is suggestive of (but only suggestive of) the fact that the RC model may not be a robust model.

6.5 Simulating the Random Choice Model

To further rule out the possibility that the data was generated by the RC model—and, hence, rule out the possibility that the gap between cognition and rationality was generated by noise—we simulate the (estimated) RC model. We then estimate the (individual) DC model for each simulated dataset. We report the mean DC distribution from 1,000 simulations of the RC model. This is shown in Figure 6.1. (Error bars represent standard errors.)

The estimates from the simulations of the RC model suggest a gap between reasoning about cognition and reasoning about rationality, with 15% of subjects being assigned to the ‘gap cognition 4 types’ $(1, 1, \frac{4}{9})$, $(\frac{7}{8}, \frac{5}{6}, \frac{7}{8})$, and $(\frac{2}{7}, \frac{2}{7}, \frac{2}{7})$. This is less than the 21% of subjects assigned to the ‘gap

¹⁹We include the noise to capture the behavior of subjects who cannot be classified according to our approach.

²⁰The subject’s behavior is not consistent with the Deliberate Choice model.




Figure 6.1. Simulated Data with the Random Choice model

cognition 4 types’ in the actual data. But, the difference is not statistically significant. Based on this fact, one might conclude that the gap could have been generated by the RC model. However, if the gap between the cognitive and rationality bounds were actually generated by the RC model, there would be other observable implications. Specifically, if the RC model did generate this distribution on types, it would be able to approximate the (full) distribution of types. But, referring to Figure 6.1, it does not. First, the proportion of error types in the simulated data is quite different from the proportion of error types in the actual data. The difference is striking. The simulations predict that we should observe a very large proportion of error types—over 40%. But, in the data, error types are rare, making up only 1% of observed behavior. Second, types with low cognitive bounds—i.e., $k = 1, 2$ —disappear in the simulated data, but make up more than 30% of the actual 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.

7 Non-Degenerate Beliefs about Rationality

In the previous section, we argued that the data is well-described by the DC model. In this section, we conclude by taking a closer look at the gap types—i.e., the types for whom the cognitive bound is strictly greater than the rationality bound—in the best fitting DC model. We argue that these gap types all exhibit non-degenerate beliefs about rationality.

Refer to Table 7.1. Each gap type is associated with some vector in $([0, 1] \cup \{cb\})^2$. Recall, this label reflects a set of equivalent probabilities. The table indicates the sets of associated equivalent probabilities. For instance, $(p_2, p_3, p_4) = (1, 1, 4/9)$ is associated with the type that can have any $\tilde{p}_2 \in [7/8, 1]$, $\tilde{p}_3 \in [7/8, 1]$ and $\tilde{p}_4 \in [1/3, 4/9]$, provided these probabilities meet the requirements in Proposition 5.1.

Note an important property of these sets of beliefs: Each gap type must have either $p_2 \in (0, 1)$ or $p_4 \in (0, 1)$. This conclusion is not inevitable. In particular, while a gap type cannot be characterized by $(p_2, p_3, p_4) = (1, 1, 1)$ (resp. $(p_2, p_3, cb) = (1, 1, cb)$ or $(p_2, cb, cb) = (1, cb, cb)$), a gap type can,

Type	p_2	p_3	p_4
(1, 1, 4/9)	[7/8, 1]	[7/8, 1]	[1/3, 4/9]
(7/8, 5/6, 7/8)	[5/6, 7/8]	[15/62, 5/6]	[0, 7/8]
(2/7, 2/7, 2/7)	[1/11, 2/7]	[0, 2/7]	[0, 2/7]
(2/5, 2/5, cb)	[2/7, 2/5]	[0, 2/5]	-
(7/8, cb, cb)	[2/5, 7/8]	-	-

Table 7.1. Reasoning About Rationality: Gap Types

in principle, be characterized by $(p_2, p_3, p_4) = (0, 0, 0)$ or $(p_2, p_3, p_4) = (1, 1, 0)$, etc. Table 7.1 indicates that this is not the case for the gap types in the (best fitting) DC model.

In the Introduction, we highlighted the fact that such non-degenerate beliefs have important implications for out-of-sample predictions. Refer to Figure 1.1: Rational behavior depends on the parameter x if and only if $p_2 \in (0, 1)$. With the exception of the type labeled (1, 1, 4/9), all types must have $p_2 \in (0, 1)$. The type labeled (1, 1, 4/9) *can* be associated with $p_2 \in (0, 1)$. But, even if this type represented a subject with $p_2 = p_3 = 1$, the type’s behavior could depend on parameters of the game (i.e., in other games). This is because $p_4 \in (0, 1)$.

		Bob		
		L	R	R'
Ann	U	10,0	0,5	0,-5
	D	$x,0$	10,5	0,-5
	D'	-5,5	-5,0	5,-5

Figure 7.1. A Game Parameterized by $x \in (-\infty, 10)$

Refer to the game in Figure 7.1, which adds the actions D' and R' to Figure 1.1. If a rational Ann has $(p_2, p_3, p_4) = (1, 1, 1)$, then she plays D . Likewise, if a rational Ann has $(p_2, p_3, p_4) = (1, 1, 0)$, then she plays U . As before, these conclusions are independent of the value of x . But, if a rational Ann has $(p_2, p_3, p_4) = (1, 1, p_4)$ for $p_4 \in (0, 1)$, then the conclusion does depend on x .

This idea applies more broadly: There is a class of games that are equivalent for types with degenerate beliefs—in particular, for types whose beliefs coincide with iterated dominance—but which are not equivalent for types with non-degenerate beliefs—i.e., for types with some $p_\ell \in (0, 1)$. Note, a necessary condition for non-degenerate beliefs is a gap between the rationality and cognitive bounds. The fact that there is a gap raises the possibility that such non-degenerate beliefs are an important determinant of behavior.

References

Adolphs, R. (1999). Social cognition and the human brain. *Trends in Cognitive Sciences* 3, 469–479.

- 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. (2007). The power of paradox: some recent developments in interactive epistemology. *International Journal of Game Theory* 35(4), 465–492.
- 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.
- Cooper, D., E. Fatas, A. J. Morales, and S. Qi (2016). Consistent depth of reasoning in level-k models. Technical report.
- 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. Iriberry (2013). Structural models of nonequilibrium strategic thinking: Theory, evidence, and applications. *Journal of Economic Literature* 51, 5–62.
- Dekel, E. and M. Siniscalchi (2014). Epistemic game theory. *Handbook of Game Theory* 4.

- 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.
- Heifetz, A. and W. Kets (2017). Robust multiplicity with a grain of naiveté. *Theoretical Economics*. Forthcoming.
- Hu, T.-W. (2007). On p-rationalizability and approximate common certainty of rationality. *Journal of Economic Theory* 136(1), 379–391.
- Jin, Y. (2016). Does level- k behavior imply level- k thinking? Working paper, U.C. Berkeley.
- Kets, W. (2011). Bounded reasoning and higher-order uncertainty. Working paper, Northwestern University.
- 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.
- Monderer, D. and D. Samet (1989). Approximating common knowledge with common beliefs. *Games and Economic Behavior* 1(2), 170–190.
- Morris, S. (1999). Approximate common knowledge revisited. *International Journal of Game Theory* 28(3), 385–408.
- 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.

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 \(1988\)](#) and [Battigalli and Siniscalchi \(2002\)](#).

Say that a strategy is an R1-strategy if it is rational, i.e., if it is a best response given *some* belief about play. For $m > 1$,

A strategy is an **R m -strategy** if it is a best response given some belief that assigns probability one to the R $(m - 1)$ -strategies.

A subject is an **R m reasoner** if she chooses a strategy that is in R m but not R $(m + 1)$. This is consistent with the definitions adopted in the main text.

Let us make two observations about R m -strategies. First, if a strategy is R m then it is also an R n -strategy for $n \leq m$. Second, a strategy is R m 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 model This is the model introduced by [Nagel \(1995\)](#). The level- k literature begins by specifying the behavior of so-called L0 reasoners. L0 reasoners are non-strategic. Thus, their behavior is characterized by an L0-distribution. An L0-distribution for Ann (resp. Bob) is denoted by p_a^0 (resp. p_b^0). For $k > 0$,

A strategy is an **L k -strategy** if it is a best response to an L $(k - 1)$ -distribution.

A subject is an **L k reasoner** if she chooses an L k -strategy. A distribution is an L k -distribution if it has support in the L k -strategies.

Because we seek to understand the concepts at an abstract level—with applicability to any game—we have described the concept in its full 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 L k -distribution is degenerate. That is, there is a unique L k -distribution and that distribution assigns probability one to a particular strategy. (This property would necessarily hold in a “generic” game, provided the L0-distributions 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 L k strategy, r_a is an L n strategy, and $k \neq n$, then $s_a \neq r_a$. That is, strategies played by L k reasoners are distinct from strategies played by all lower-order reasoners.

Cognitive Hierarchy Model Like level- k models, cognitive hierarchy models assume that players choose a best response to a belief that their opponent has a lower level. Unlike in level- k models, players think that any lower level is possible. We follow the specification in [Camerer, Ho, and Chong \(2004\)](#). As before, the starting point is distributions p_a^0 and p_b^0 that describe the behavior of CH0 reasoners, who are non-strategic. Define the CH0-distribution for Ann, denoted by q_a^0 , to be equal to p_a^0 . Likewise for Bob. Refer to the strategies in the support of q_a^0 and q_b^0 as CH0-strategies.

To derive the distributions for higher level reasoners, fix a parameter $\tau > 0$, and for $\ell = 0, 1, \dots$, let $f(\ell; \tau)$ be the Poisson density at ℓ (i.e., $f(\ell; \tau) = \tau^\ell e^{-\tau} / \ell!$). The idea is that $f(k; \tau)$ is the “true” fraction of players who reason up to level k . However, a player who reasons up to level k can conceive only of players who reason up to a lower level. Thus, if Ann reasons up to level k , then her belief over Bob’s reasoning levels is given by the truncated Poisson distribution which assigns probability $f(\ell; \tau) / \sum_{m=0}^{k-1} f(m; \tau)$ to Bob reasoning up to level $\ell \leq k - 1$ (and probability zero to levels greater than $k - 1$). For $k > 0$,

The **CH k -strategy** is the best response to the CH($k - 1$)-distribution. (If there are multiple best responses, the CH k -strategy assigns equal probability to each of them.)

A subject is a **CH k reasoner** if she plays the CH k -strategy. The CH k -strategy for Ann thus defines a distribution p_a^k over strategies. The CH k -distribution for Ann, denoted q_a^k , is the distribution over strategies if for $\ell \leq k - 1$, the fraction of CH ℓ reasoners is given by the truncated Poisson distribution and CH ℓ reasoners play according to p_a^ℓ . That is,

$$q_a^k = \sum_{\ell=0}^{k-1} \left(\frac{f(\ell; \tau)}{\sum_{m=0}^{k-1} f(m; \tau)} \right) p_a^\ell.$$

As in the case of level- k models, cognitive hierarchy models are often applied to games where CH k reasoners have a unique best response.

Connections Let us draw connections between bounded reasoning about rationality, level- k 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 L k -strategy is also an R k -strategy.²¹ However, the converse does not hold. There may be strategies that are R k but not L k -strategies. This is because level- k models fix an (exogenous) L0-distribution, and there can be strategies that are not a best response to the exogenous L0 belief even if they are a best response to some other belief.

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

²¹In fact, if $j \geq k \geq 1$, then an L j -strategy is an R k -strategy.

	P1	P2	P3	P4
L1	(a, c _*)	(a, a _*)	(b, b _*)	(a, a _*)
L2	(a, c _*)	(a, b _*)	(b, b _*)	(a, a _*)
L3	(a, c _*)	(a, b _*)	(b, a _*)	(a, a _*)
L4	(a, c _*)	(a, b _*)	(b, a _*)	(a, c _*)

(a) Level- k Reasoning

	P1	P2	P3	P4
CH1	(a, c _*)	(a, a _*)	(b, b _*)	(a, a _*)
CH2	(a, c _*)	(a, a _*)	(a, a _*)	(c, c _*)
CH3	(a, c _*)	(a, a _*)	(a, a _*)	(c, c _*)
CH4	(a, c _*)	(a, a _*)	(a, a _*)	(c, c _*)

(b) Cognitive Hierarchy Reasoning

Table A.1. Application to the Ring Game

to strategies in the support of an L0 distribution 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 R k -strategies are disjoint from the CH k -strategies, for all $k \geq 2$.

Third, typically, both L k reasoning and CH k 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,”

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 (a, c_{*}) in the role of P1, (a, a_{*}) in the role of P2, (b, b_{*}) in the role of P3, and (a, a_{*}) in the role of P4. Tables A.1a-A.1b give the behavior of the L k and CH k 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 L $k \neq$ L0 reasoner, we would conclude that the subject is an R k 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 L k reasoners, for any k . Second, if we were to observe the behavior of any CH $k \neq$ CH0 reasoner, we would conclude that the subjects’ cognitive bound is 1.

Appendix B Identifying Cognitive Bounds

Lemma B.1. *Let $P_i = P_1, P_3, P_4$. If P_i believes that $P(i-1)$ has a cognitive bound of 1, then P_i has a constant belief.*

Proof. If P_i believes that $P(i-1)$ has a cognitive bound of 1, then P_i believes that “ $P(i-1)$ is cognitive and satisfies the Principle of Belief in Lack of Cognition.” It follows from the Principle of Belief in Cognition, P_i ’s belief satisfies $\Pr_i(d, e_*) > 0$ only if $\pi_{(i-1)}(d) = e_*$. Since $i-1 \neq 1$, it follows that P_i has a constant belief. \clubsuit

Lemma B.2. *If P_1 believes that P_4 has a cognitive bound of 2, then P_1 has a constant belief.*

Proof. Suppose that P_1 believes that P_4 has a cognitive bound of 2. That is, P_1 believes “ P_4 is 2-cognitive and believes that P_3 has a cognitive bound of 1.” When P_4 believes that “ P_3 has a cognitive bound of 1,” P_4 has a constant belief (Lemma B.1). Thus, P_1 believes “ P_4 is 2-cognitive and has a constant belief.” Thus, applying the Principle of Belief in Cognition, P_1 believes that P_4 plays a constant strategy. So, again, P_1 has a constant belief. \clubsuit

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 _*), (b, b _*), (c, c _*)	(a, a _*), (b, b _*), (c, c _*)	45	16
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					6071	1
Total					6561	75

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

Appendix C Intervals for Section 5

Table C.1 re-expresses Table 5.1, in a way that easily permits computing the lower and upper bounds of p_2 . Table C.2 uses Table C.1 to provide lower and upper bounds of p_3 for every value of

q_2 . Finally, in the role of P4, we observe three strategies played: (a, b_*) , (a, c_*) , and (c, a_*) . Table C.3 uses Table C.2 to provide lower and upper bounds of p_4 for those observations, given various values of q_3 and r_2 .

p_2	Actions	p_2	Actions
$[\frac{7}{8}, 1]$	(a, b_*)	$[\frac{2}{5}, \frac{2}{5}]$	$(a, b_*), (b, c_*), (a, a_*)$
$[\frac{5}{7}, \frac{7}{8})$	$(a, b_*), (a, c_*)$	$(\frac{2}{7}, \frac{2}{5})$	$(a, b_*), (b, c_*), (a, a_*), (b, a_*)$
$[\frac{62}{133}, \frac{5}{7})$	$(a, b_*), (a, c_*), (b, c_*)$	$[\frac{29}{133}, \frac{2}{7}]$	$(b, c_*), (a, a_*), (b, a_*)$
$(\frac{2}{5}, \frac{62}{133})$	$(a, b_*), (a, c_*), (b, c_*), (a, a_*)$	$(\frac{1}{11}, \frac{29}{133})$	$(b, c_*), (a, a_*), (b, a_*), (c, a_*)$
		$[0, \frac{1}{11}]$	$(b, c_*), (b, a_*), (c, a_*)$

Table C.1. Assigning probability p_2 to Rationality

	(a, a_*)	(a, b_*)	(a, c_*)	(b, a_*)	(b, b_*)	(b, c_*)	(c, a_*)	(c, b_*)	(c, c_*)
$q_2 \in [\frac{7}{8}, 1]$	$[0, \frac{5}{8})$	$[0, \frac{14}{31})$	$[0, \frac{5}{8})$	$[0, 1]$	$[0, \frac{14}{31})$	$[0, \frac{7}{8})$	$[0, \frac{15}{62})$	$[0, \frac{15}{62})$	$[0, \frac{15}{62})$
$q_2 \in [\frac{5}{7}, \frac{7}{8})$	$[0, \frac{5}{8})$	$[0, \frac{5}{8})$	$[0, \frac{5}{8})$	$[0, 1]$	$[0, \frac{5}{6})$	$[0, 1]$	$[0, \frac{15}{62})$	$[0, \frac{15}{62})$	$[0, \frac{15}{62})$
$q_2 \in [\frac{62}{133}, \frac{5}{7})$	$[0, \frac{6}{7})$	$[0, \frac{5}{6})$	$[0, 1]$	$[0, 1]$	$[0, \frac{5}{6})$	$[0, 1]$	$[0, \frac{817}{2914})$	$[0, \frac{35}{48})$	$[0, \frac{7}{8})$
$q_2 \in [\frac{2}{5}, \frac{62}{133})$	$[0, \frac{231}{248})$	$[0, 1]$	$[0, 1]$	$[0, 1]$	$[0, 1]$	$[0, 1]$	$[0, \frac{817}{2914})$	$[0, \frac{3}{4})$	$[0, \frac{7}{8})$
$q_2 \in [\frac{29}{133}, \frac{2}{5})$	$[0, 1]$	$[0, 1]$	$[0, 1]$	$[0, 1]$	$[0, 1]$	$[0, 1]$	$[0, \frac{35}{66})$	$[0, \frac{7}{8})$	$[0, \frac{7}{8})$
$q_2 \in [0, \frac{29}{133})$	$[0, 1]$	$[0, 1]$	$[0, 1]$	$[0, 1]$	$[0, 1]$	$[0, 1]$	$[0, \frac{215}{248})$	$[0, 1]$	$[0, 1]$

Table C.2. P3: Probability p_3 to “Rationality and q_2 -Belief of Rationality”

	(a, b_*)	(a, c_*)	(c, a_*)
$r_2 \in [\frac{7}{8}, 1] \quad q_3 \in [\frac{7}{8}, 1]$	$[0, \frac{1}{3})$	$[0, 1]$	$[0, \frac{4}{9})$
$r_2 \in [\frac{7}{8}, 1] \quad q_3 \in [\frac{5}{8}, \frac{7}{8})$	$[0, 1]$	$[0, 1]$	$[0, \frac{5}{9})$
$r_2 \in [\frac{7}{8}, 1] \quad q_3 \in [0, \frac{5}{8})$	$[0, 1]$	$[0, 1]$	$[0, 1]$
$r_2 \in [\frac{5}{7}, \frac{7}{8}) \quad q_3 \in [\frac{5}{6}, 1]$	$[0, 1]$	$[0, 1]$	$[0, \frac{5}{9})$
$r_2 \in [\frac{5}{7}, \frac{7}{8}) \quad q_3 \in [0, \frac{5}{6})$	$[0, 1]$	$[0, 1]$	$[0, 1]$
$r_2 \in [0, \frac{5}{7}) \quad q_3 \in [0, 1]$	$[0, 1]$	$[0, 1]$	$[0, 1]$

Table C.3. P4: Probability p_4 to “Rationality and q_3 -Belief of ‘Rationality and r_2 -Belief of Rationality’”

Appendix D The Anonymity Assumption

It will be convenient to introduce some notation: Write R_i for the event that a subject is rational in the role of Pi . Write $B_i^p(E_j)$ for the event that i assigns probability p to the event E_j . Write $\tilde{B}_i^p(E_j)$ for the event that i assigns probability at least probability p to event E_j .

It will be convenient to record four properties of the standard belief operator.

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

Property 2 If E_j, F_j are events with $E_j \implies F_j$, 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 and E_j implies F_j , then i assigns probability at least p to F_j .²² 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 $q \geq p$.

Consider a subject who has a cognitive bound of $k = 3, 4$. If $k = 3$, this subject is characterized by a $(p_2; p_3, q_2)$ and, if $k = 4$, this subject is characterized by a $(p_2; p_3, q_2; p_4, q_3, r_2)$. We fix these characterizations in what follows below.

Lemma D.1. *Suppose the AA holds. If the subject has a cognitive bound of $k = 3, 4$ then, for each Pi with $P3 \leq Pi \leq Pk$, $p_2 \geq p_i$.*

Proof. Consider a subject who has a cognitive bound of $k = 3, 4$. Notice that $B_2^{p_2}(R_1)$ and so, by the Anonymity Assumption, $B_i^{p_2}(R_{i-1})$ holds in each player role $P3 \leq Pi \leq Pk$. At the same time, by Property 1, $\tilde{B}_i^{p_2}(R_{i-1})$ holds. So applying Property 4, $p_2 \geq p_i$. ♣

Lemma D.2. *Suppose the AA holds. If the subject has a cognitive bound of 4 and $q_2 \leq q_3$, then $p_3 \geq p_4$.*

Proof. Notice that

$$B_4^{p_4} \left(R_3 \wedge \tilde{B}_3^{q_3} (R_2 \wedge \tilde{B}_2^{r_2} (R_1)) \right),$$

i.e., $P4$ assigns probability p_4 to $R_3 \wedge \tilde{B}_3^{q_3} (R_2 \wedge \tilde{B}_2^{r_2} (R_1))$. Applying Property 1, it follows that

$$\tilde{B}_4^{p_4} \left(R_3 \wedge \tilde{B}_3^{q_3} (R_2 \wedge \tilde{B}_2^{r_2} (R_1)) \right),$$

²²A 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.

i.e., P4 assigns at least probability p_4 to $R_3 \wedge \tilde{B}_3^{q_3}(R_2 \wedge \tilde{B}_2^{r_2}(R_1))$. Observe that $R_2 \wedge \tilde{B}_2^{r_2}(R_1) \implies R_2$ and so, applying Property 2,

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

And implication of the Anonymity Assumption is that

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

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, $R_2 \wedge \tilde{B}_2^{q_3}(R_1) \implies R_2 \wedge \tilde{B}_2^{q_2}(R_1)$. So, again applying Property 2,

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

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

Appendix E Individual Estimates

Table E.1 compares the individual analysis for the DC and RC model. The first four columns contain the estimates of the DC model and the last four columns correspond to the RC model. In both cases, the column “Subject” indicates the number of subjects assigned to the type, according to the estimates $\hat{\pi}_n$. Likewise, in both cases, the column “Distribution” indicates the induced distribution of types.

Deliberate Choice				Random Choice			
(P_2, P_3, P_4)	ε	Subjects	Distribution	c^k	ε	Subjects	Distribution
$(1, 1, 1)$	0	13 (3.40)	.17 (.05)	c^4	0	13 (3.28)	.17 (.04)
$(1, 1, \frac{4}{9})$	0	2 (1.39)	.03 (.02)	c^4	$\frac{1}{8}$	11 (3.15)	.15 (.04)
$(\frac{7}{8}, \frac{5}{6}, \frac{7}{8})$	0	11 (3.05)	.15 (.04)	c^4	$\frac{2}{8}$	1 (1.03)	.01 (.01)
$(\frac{2}{7}, \frac{2}{7}, \frac{2}{7})$	0	2 (1.37)	.03 (.02)	c^3	0	19 (3.85)	.25 (.05)
$(1, 1, \text{cb})$	0	16 (3.54)	.22 (.05)	c^3	$\frac{1}{6}$	1 (0.97)	.01 (.01)
$(\frac{2}{5}, \frac{2}{5}, \text{cb})$	0	5 (2.18)	.07 (.03)	c^2	0	15 (3.50)	.2 (.05)
$(1, \text{cb}, \text{cb})$	0	13 (3.28)	.17 (.04)	c^1	0	15 (3.63)	.2 (.05)
$(\frac{7}{8}, \text{cb}, \text{cb})$	0	3 (1.69)	.04 (.02)				
$(\text{cb}, \text{cb}, \text{cb})$	0	9 (2.80)	.12 (.04)				
$(\frac{2}{7}, \frac{2}{7}, \frac{2}{7})$	1	1 (0.95)	.01 (.01)				
Neg. Log-Likelihood		146.94		258.80			

Table E.1. Deliberate Choice model versus Random Choice model - Individual Analysis

Appendix F Alternate Random Choice Model

This Appendix considers an Alternate Random Choice model, one that shares similarities with our Deliberate Choice model. Similar to the Random Choice and Logistic Choice models in Section 6.2, this model takes as given that a subject's rationality bound coincides with her cognitive bound; it interprets what appears to be a gap between cognition and rationality as an artifact of errors. This model differs from the Random and Logistic Choice models above in the predicted behavior when the cognitive bound binds. For example, consider a subject who has a cognitive bound of 3 and, so, a rationality bound of 3. If there are no errors, the subject should play the IU strategy in the role of P1-P2-P3 and a constant strategy in the role of P4. This draws on the analysis in the paper and, so, is in the spirit of the Deliberate Choice model. But it differs from the Random Choice model of Section 6.2: The Random Choice model assumes that, in the role of P4, the subject plays all strategies with equal probability. This Alternate Random Choice model instead assumes that, in the role of P4, the subject plays each constant strategy with equal probability and, when there is an error, each non-constant strategy with equal probability.

Write $\mathcal{M}^{\text{AR}} = (T^{\text{AR}}, \pi, \varepsilon)$ for a **Alternate Random Choice model**. Much as the Random Choice model, it has three components: Types, $T^{\text{AR}} = \{c^1, c^2, c^3, c^4\}$, a probability distribution on

types T^{AR} , namely, π , and type-specific trembles are given by $\varepsilon = (\varepsilon^k)_{k=1}^4 \in (0, 1)^4$. It differs from the Random Choice model in the likelihood of observing a given strategy profile.

To describe how the cognitive bound and trembles influence the likelihood of observing a given strategy profile, it will be convenient to introduce notation. Let $\eta_R[x_n, i]$ be an indicator for whether $x_n(i)$ and $x^{\text{rat}}(i)$ agree. So if $x_n(i) = x^{\text{rat}}(i)$, then $\eta_R[x_n, i] = 1$ and if $x_n(i) \neq x^{\text{rat}}(i)$, then $\eta_R[x_n, i] = 0$. Let $\eta_C[x_n, i]$ be an indicator for whether $x_n(i)$ is a constant strategy. So, $\eta_C[x_n, i] = 1$ if $x_n(i)$ is a constant strategy and $\eta_C[x_n, i] = 0$ otherwise.

Consider a subject whose cognitive bound is $k = 3$. In the role of P4, up to trembles, the subject randomizes equally amongst the three constant strategies. In the role of P3, up to trembles, the subject plays the IU strategy. The Alternate Random Choice model assumes that trembles are independent of the payoffs and the player role. So, in the role of P4, the subject plays one of the constant strategies with probability $\frac{1-\varepsilon}{3}$ and plays any of the other six strategies with probability $\frac{\varepsilon}{6}$. In the role of P3, the subject plays the IU strategy with probability $1 - \varepsilon$ and plays any of the other eight strategies with probability $\frac{\varepsilon}{8}$. And similarly for her play in the role of P2 and P1.

With this, the probability of observing x_n in the model \mathcal{M}^{AR} given a subject of type c^k is

$$p(x_n, \varepsilon^k | c^k) = \left(\frac{1 - \varepsilon^k}{3} \right)^{\sum_{i=k+1}^4 \eta_C[x_n, i]} \left(\frac{\varepsilon^k}{6} \right)^{\sum_{i=k+1}^4 (1 - \eta_C[x_n, i])} \left(1 - \varepsilon^k \right)^{\sum_{i=1}^k \eta_R[x_n, i]} \left(\frac{\varepsilon^k}{8} \right)^{\sum_{i=1}^k (1 - \eta_R[x_n, i])}$$

Then, the likelihood of observing behavior x_n in the model \mathcal{M}^{AR} is

$$\mathcal{L}_n(x_n, \mathcal{M}^{\text{AR}}) = \sum_{c^k \in T^{\text{AR}}} \pi(c^k) p(x_n, \varepsilon^k | c^k).$$

And, the aggregate log-likelihood of observing the experimental dataset $\mathbf{x} = (x_n)_{n=1}^N$ is

$$\ln \mathcal{L}(\mathbf{x}, \mathcal{M}^{\text{AR}}) = \sum_{n=1}^N \ln \mathcal{L}_n(x_n, \mathcal{M}^{\text{AR}}).$$

We choose $\hat{\mathcal{M}}^{\text{AR}}$ to maximize $\ln \mathcal{L}(\mathbf{x}, \mathcal{M}^{\text{AR}})$.

The aggregate results are given in Table F.1. The Deliberate Choice model outperforms the Alternate Random Choice model according to Log-Likelihood, BIC, and the AIC. The individual results are presented in Table F.2. Like the results for the Random Choice model, the differences between the aggregate and individual results are suggestive of the fact that the Alternate Random Choice model may not be a robust model.

We ran 1000 simulations of the estimated individual Alternate Random Choice Model and fit the simulated datasets to the individual Deliberate Choice Model. The mean Deliberate Choice type distribution is presented in Figure F.1. Importantly, the Alternate Random Choice model fails to generate the gap between the cognitive and rationality bounds that are observed in the data. Moreover, the simulations suggest that we should observe far more error types if the data was generated by our Alternate Random Choice model.

Deliberate Choice			Alternate Random Choice		
$(\mathbf{p}_2, \mathbf{p}_3, \mathbf{p}_4)$	π	ε	\mathbf{c}^k	π	ε
(1, 1, 1)	.13	.006	c^4	.46	.164
(1, 1, $\frac{4}{9}$)	.05	.027	c^3	.10	0
($\frac{7}{8}, \frac{5}{6}, \frac{7}{8}$)	.16	.015	c^2	.17	.034
($\frac{2}{7}, \frac{2}{7}, \frac{2}{7}$)	.03	.018	c^1	.27	.17
(1, 1, cb)	.21	.007			
($\frac{2}{5}, \frac{2}{5}, \text{cb}$)	.07	.012			
(1, cb, cb)	.13	.012			
($\frac{7}{8}, \text{cb}, \text{cb}$)	.08	.052			
(cb, cb, cb)	.12	.003			
Neg. Log-Likelihood		294.78			334.74
BIC		662.73			699.70
AIC		623.56			683.48

Table F.1. Deliberate Choice model versus Alternate Random Choice model

Deliberate Choice				Alternate Random Choice			
(p_2, p_3, p_4)	ε	Subjects	Distribution	c^k	ε	Subjects	Distribution
(1, 1, 1)	0	13 (3.40)	.17 (.05)	c^4	0	13 (3.07)	.17 (.04)
(1, 1, $\frac{4}{9}$)	0	2 (1.39)	.03 (.02)	c^4	$\frac{1}{4}$	14 (3.26)	.19 (.04)
($\frac{7}{8}, \frac{5}{6}, \frac{7}{8}$)	0	11 (3.05)	.15 (.04)	c^3	0	16 (3.53)	.21 (.05)
($\frac{2}{7}, \frac{2}{7}, \frac{2}{7}$)	0	2 (1.37)	.03 (.02)	c^3	$\frac{1}{4}$	4 (1.99)	.05 (.03)
(1, 1, cb)	0	16 (3.54)	.22 (.05)	c^2	0	13 (3.22)	.17 (.04)
($\frac{2}{5}, \frac{2}{5}, cb$)	0	5 (2.18)	.07 (.03)	c^1	0	9 (2.82)	.12 (.04)
(1, cb, cb)	0	13 (3.28)	.17 (.04)	c^1	$\frac{1}{4}$	6 (2.35)	.08 (.03)
($\frac{7}{8}, cb, cb$)	0	3 (1.69)	.04 (.02)				
(cb, cb, cb)	0	9 (2.80)	.12 (.04)				
($\frac{2}{7}, \frac{2}{7}, \frac{2}{7}$)	1	1 (0.95)	.01 (.01)				
Neg. Log-Likelihood		146.94		195.44			

Table F.2. Deliberate Choice model versus Alternate Random Choice model - Individual Analysis

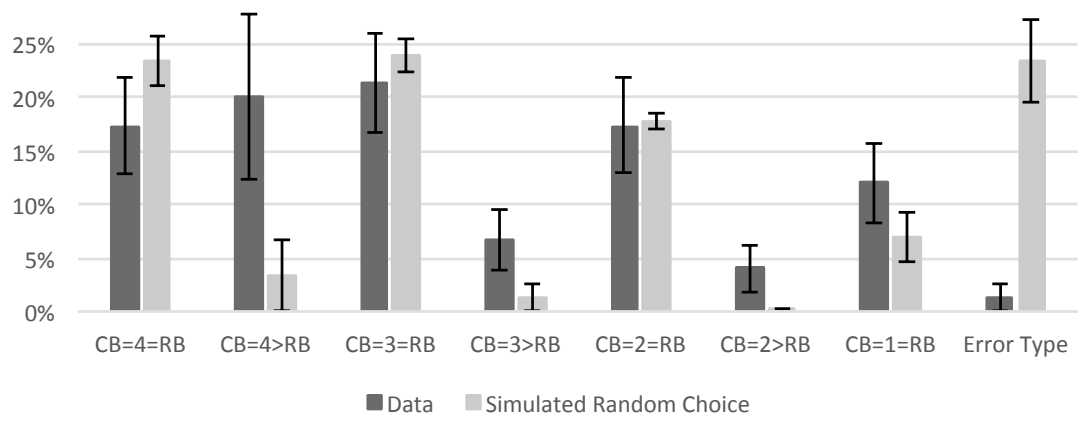


Figure F.1. Simulated Data with the Alternate Random Choice model