Is Bounded Rationality Driven by Limited Ability?*

Amanda Friedenberg† Willemien Kets‡ Terri Kneeland§

First Draft: January 31, 2016
This Draft: July 26, 2018

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; Arad and Rubinstein, 2012; Crawford, Costa-Gomes, and Iriberri, 2013; Kneeland, 2015, amongst many others.) For example, a player may be rational and believe that

---

*We are indebted to Adam Brandenburger and Vince Crawford for many helpful conversations. We also thank Marina Agranov, Martin Dufwenberg, Jeff Erlich, 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, NYU Shanghai, Barcelona GSE Summer Forum, BRIC, the Game Theory Society World Congress, Alicante Level-k Workshop, the CRETA Conference on Identifying Rationality, the Society for Advancement of Economic Theory, and Logic & and the Foundations of Game and Decision Theory. Friedenberg thanks NSF grant SES-1358008 for financial support and Caltech, Kellogg School of Management and UCL for hospitality. Previous versions of the paper were circulated under the title "Bounded Reasoning: Rationality or Cognition."

†University of Arizona, amanda@amandafriedenberg.org, http://amandafriedenberg.org
‡University of Oxford, willemien.kets@economics.ox.ac.uk, http://wkets.org
§University College London, t.kneeland@ucl.ac.uk, http://tkneeland.com

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.

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 interactive reasoning. (See, e.g., Perner and Wimmer, 1985; Kinderman, Dunbar, and Bentall, 1998; Stiller and Dunbar, 2007, amongst many others.) Such limitations can, in turn, 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 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.

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? This paper provides a conceptual and practicable framework to address this important question. It shows that rationality bounds cannot entirely be explained by limited ability to engage in interactive reasoning. We, first, discuss the importance of addressing the question. Second, we turn to the approach, providing an overview of the framework. Third, we preview the results and their implications for out of sample predictions. Finally, we discuss connections to the literature.

Importance of the Question Addressing the question is important for at least two reasons: the external validity of laboratory experiments and for how behavioral game theory builds models of bounded reasoning.

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 (or willingness) 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.[1] 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.[2] 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. (See 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 best fitting model does exhibit such non-degenerate beliefs.

Our Approach The goal is to address the question: Is bounded reasoning about rationality determined by limitations in the ability to engage in interactive reasoning? The challenge is that the researcher does not observe the players' ability bounds. One might hope for an approach in which the researcher elicits the subject's ability to engage in such interactive reasoning. However, 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. In turn, this can suggest evidence of a gap between the ability and rationality bounds, even if none exists in practice (i.e., absent intervention of the researcher).

We develop a conceptual and practicable framework that allows us to address the question of interest. Importantly, the approach will allow us to sidestep the difficulty—that is, it will allow us to address the question without identifying the players' ability bounds. We begin by describing the conceptual framework.

Recall, limited ability to engage in interactive reasoning will limit the players' ability to reason about their opponent's rationality. More broadly, it would also limit their ability to reason about how their opponent plays the game. The approach will be to identify behavior that is (i) consistent with a player engaging in interactive reasoning about how her opponent plays the game, but (ii) inconsistent with interactive reasoning about rationality. Such behavior would indicate that

---

[1] This is true of both the level-k and cognitive hierarchy models. The single bound is typically referred to as a player's *depth of reasoning*; it captures both the rationality and ability bounds. (Though how the ability bounds are conceptualized differs across the two models.)

[2] This 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.

bounded reasoning about rationality is not (entirely) determined by the players' ability to engage in interactive reasoning.

With this in mind, it will be useful to distinguish between a rational player and a strategic player. Say that a player 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. Say that she is *strategic* if she has some theory (or method) for how to play the game. One example of such a theory is maximizing subjective expected utility; thus, a player who is rational is also strategic. However, a player may be strategic and irrational; that is, a player may have a decision criterion for playing the game which departs from subjective expected utility.[3]

We will distinguish between reasoning about rationality and strategic reasoning:

- *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.

- *Strategic Reasoning:* Say that Ann has a *strategic bound* of k if she is strategic, believes that Bob is strategic, believes that Bob believes she is strategic and so on, up to the statement that includes the word "strategic" 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 strategic bound k, i.e., $m \leq k$.

Now observe that strategic reasoning still requires an ability to engage in interactive reasoning. Thus, a subject's ability bound must be at least as high as her strategic bound. Thus, if a subject's strategic bound is strictly higher than her rationality bound, i.e., if $k > m$, then the subject's ability bound is also strictly higher than her rationality bound. In that case, bounded reasoning about rationality is not entirely determined by limited ability to engage in interactive reasoning.

In light of this, to identify a gap between the rationality bound and the ability bound, it suffices to identify a gap between the rationality bound and the strategic bound. However, doing so poses a second challenge: identifying the strategic bound. In principle, strategic reasoning is a broad concept; thus, it is not obvious what (if any) observable implications arise from strategic reasoning. In Sections 2-3, we point to observable implications in a particular class of games: permuted ring games, introduced by Kneeland (2015).

Let us preview the identification strategy. To identify both the rationality and strategic bounds, we assume that all subject's are rational (i.e., maximize their subjective expected utility). We identify the rationality bound based on iterated dominance. We identify the strategic bound based on two identification assumptions: the Principle of Strategic Reasoning and the Principle of Non-Strategic Reasoning. The logic behind these principles rests on an assumption that strategic behavior depends, in systematic ways, on payoffs of the game. The logic presumes that strategic

---

[3] 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.) Note, for us, a rational player conforms to the "textbook notion of rationality," while strategic a strategic player conforms to a "generalized notion of rationality."

behavior does not depend on certain fine presentation effects. (Sections 2-3.2.1-3.2.2 will clarify the presentation effects that are ruled out.) Permuted ring games are well-suited for applying these principles and, as a consequence, are well-suited for identifying the strategic bound (Remark 3.1).

Preview of Results Section 4 applies the identification strategy to Kneeland's (2015) experimental dataset. We find that 12% of our subjects have a strategic bound of 1, 22% have a strategic bound of 2, 28% have a strategic bound of 3, and 38% have a strategic bound of 4. Moreover, there is a nontrivial gap between subjects' rationality and strategic bounds. We find that 47% of the subjects identified as having a low rationality bound (i.e., either 1 or 2) have a higher strategic bound. The gap between the rationality and strategic bounds is most pronounced for subjects that have the highest strategic bound (i.e., 4). Further, subjects identified as having a gap between their rationality and strategic bounds outperform (in terms of average expected payoffs) subjects identified as having no gap. Thus, players have an incentive to choose a rationality bound lower than their bound on ability.

We go on to explore the nature of strategic reasoning: If there were no gap between the strategic and rationality bounds, strategic reasoning would be equivalent to reasoning about rationality. However, because there is a gap, players can be classified by the extent to which they reason about rationality. For instance, consider a subject who has a strategic bound of 2. If she also has a rationality bound of 2 then—not only does she believe that the other players are strategic—she also believes they are rational. However, if she has a rationality bound of 1, then she only assigns some probability $p \in [0, 1)$ to her opponents' rationality. That is, strategic reasoning may involve (possibility degenerate) beliefs about rationality, even if the player does not assign probability 1 to her opponents rationality.

Section 5 imposes discipline on the nature of these beliefs by studying a Population-Based (PB) model. The model imposes two assumptions. First, to the extent that players do reason about rationality, they do so in a way that treats all members of the population symmetrically. Second, to the extent that players do reason about rationality, they reason that other members of the population reason as they do. There are many different PB models: one special case of the PB model is a model in which the rationality bound and the strategic bounds coincide. However, we show that this is not the PB model that best fits the data. Instead, the PB model that best fits the data involves non-degenerate beliefs about rationality. That is, it *both* precludes probabilities 0 and 1, in reasoning about rationality. (Below we discuss the importance.)

Section 6 uses the best fitting PB model to explore 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 strategic bounds and, instead, certain observed behavior is an artifact of noise. We show that the best fitting PB model outperforms the best fitting noisy decision-making model.

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

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

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

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 higher-order reasoning. Thus, the subject must also have a strategic bound of $k = 2$.

By contrast, if there is a gap between a subject's rationality and strategic bounds, then there is scope for non-degenerate beliefs in reasoning about rationality. As discussed, Section 5 shows that the best fitting PB model must have such non-degenerate beliefs. So, if a subject is identified as having a gap between her rationality and strategic bounds, her behavior is best explained by non-degenerate beliefs about rationality. We already argued that these non-degenerate beliefs have important implications for building models of behavioral game theory. (See page 3.) But, as we now discuss, they also have important implications for interpreting extant experimental results.

Some papers have found that subjects' identified levels of reasoning change across games. (See Georganas et al., 2015, Cooper et al., 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—i.e., two rounds of iterated dominance—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—i.e., one round of iterated dominance—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) shows 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; Arad and Rubinstein, 2012). 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.) This idea is so engrained in the literature that papers typically use the phrase "depth of reasoning" to refer to *both* the ability and the rationality bounds. (This is so even in the current literature.) In fact, that one bound is identified based on the extent to which subjects iterate over best responses: A subject is identified as a level-k thinker if she performs exactly k rounds of iterated best responses.[4]

That said, Arad and Rubinstein (2012) do note that the rationality bound may be distinct from the ability bound. In fact, one might hope to use their results to provide evidence that the rationality bound is not driven by limits in the ability to engage in interactive reasoning: In the context of the 11-20 game, they show that most subjects are level-k thinkers for $k = 1, 2, 3$. Because the 11-20 game is strategically simple, one might posit that—within that specific game—subjects have unlimited ability to engage in interactive reasoning; thus, the fact that subjects are level-k thinkers (for low k) may point to a gap between the rationality and ability bounds. While this evidence is suggestive of a gap, the paper does not identify a gap. First, the paper does not provide a method for identifying an ability bound. (While the game is strategically simple, it does not imply that *a fortiori* subjects engage in an unlimited number of steps of "I think, you think, I think...") Second, all strategies in the 11-20 game are rationalizable. As a consequence, all behavior is consistent with unbounded reasoning about rationality.[5]

Analogously, one might hope to use the results in Agranov, Potamites, Schotter, and Tergiman (2012), Georganas, Healy, and Weber (2015), Alaoui and Penta (2016), and Gill and Prowse (2016) to show that reasoning about rationality is not driven by limits in ability. In the context of level-k models, they show that subjects' rationality bounds may vary based on whether they are playing against more versus less sophisticated players.[6] 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

---

[4] Sometimes 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.

[5] Arad and Rubinstein (2012) employ the level-k model, which allows for auxiliary assumptions about beliefs. They anchor level-0 players as players that choose 20. Typically, such an anchor is viewed as reflecting the behavior of irrational players. However, 20 may well be played by "fully rational players," who expect others to use the Arad and Rubinstein iterative reasoning and play 11. (In that case, 20 is the unique best response.)

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

interactive reasoning does not depend on who her opponents are, then variation in her rationality bound must indicate that the bound is 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. (This is suggested by Alaoui and Penta (2016).) 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.

A recent paper by Alaoui and Penta (2017) also addresses whether rationality bounds are determined by limits in ability. Their identification strategy rests on a "tutorial method." It presumes teaching subjects about game theory (e.g., iterative best responses) can influence the ability bound, but not the rationality bound.

We identify a gap between the rationality and ability bounds, by identifying a gap between the rationality and strategic bounds. We build on the work of Kneeland (2015). Importantly, Kneeland implicitly assumes that the rationality bound is determined by limitations on a players' ability to engage in interactive reasoning. We introduce the notion of a strategic bound and use this to show that the rationality bound is not determined by such limitations in ability.[7]

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).[8]

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 strategic and rationality bounds. Section 5 illustrates that, even if there is a gap between the strategic and rationality bounds, a subject's beliefs about play can be reinterpreted as beliefs about rationality. It shows that the observed gap is best explained by non-degenerate beliefs about rationality. Section 6 uses this characterization to show that the observed gap is not an artifact of noisy decision-making.

---

[7] Whether 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 strategic 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.

[8] 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 model non-degenerate beliefs about rationality when there are ability limitations.

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 $(\mathrm{d}, \mathrm{e}_*)$ 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

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

P2's Payoffs

P1

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

(a) Figure G

P1's Payoffs

P2

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

P2's Payoffs

P1

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

(b) Figure G_*

Figure 2.1. A Two-Player Example

Rational vs. Strategic To illustrate the relationship between rational and strategic behavior, 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 $(\mathrm{a}, \mathrm{c}_*)$. Notice that, if she is rational, then she has a specific theory about how to play the game. This is the sense in which we will say that she is also *strategic*.

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

Reasoning about Rationality vs. Strategic Reasoning To illustrate the relationship between reasoning about rationality and strategic reasoning, we focus on P2. Throughout the discussion, we suppose that P2 is rational (and so strategic). We will distinguish between three scenarios.

First, suppose that P2 reasons about rationality. By this we mean, P2 *believes*—i.e., assigns probability 1 to the event—that P1 is rational. In this case, he must assign probability 1 to P1 playing $(\mathrm{a}, \mathrm{c}_*)$. As P2 is rational, he chooses a best response to this belief; he thus plays $(\mathrm{a}, \mathrm{b}_*)$. Notice, because P2 believes that P1 is rational, P2 also believes that P1 is strategic. Put differently, since P2 reasons that P1 is rational, he also reasons that P1 is strategic.

Second, suppose that P2 does not assign probability 1 to P1's rationality, but does assign probability 1 to P1 being strategic. For instance, he may assign probability $4/5$ to the rational strategy $(\mathrm{a}, \mathrm{c}_*)$ and probability $1/5$ to the lucky-6 strategy $(\mathrm{c}, \mathrm{b}_*)$. In this case, his best response is to play $(\mathrm{a}, \mathrm{c}_*)$.

Third, suppose that, unlike the two scenarios above, P2 reasons that P1 is *not strategic*. 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., $(\mathrm{a}, \mathrm{a}_*)$, $(\mathrm{b}, \mathrm{b}_*)$, or $(\mathrm{c}, \mathrm{c}_*)$.

Observe that both the first and third scenarios involve no gap between reasoning about rationality and strategic reasoning. In the first case, P2 reasons both that P1 is rational and that P1 is strategic. In that case, he rationally plays the only strategy that survives two rounds of iterated dominance. In the third case, P2 reasons that P1 is not strategic and, so, he also reasons that P1 is irrational. 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 rationality and strategic reasoning: P2 assigns probability 1 to P1 being strategic, but does not assign probability 1 to P1 being rational. He, then, rationally plays a non-constant strategy—one that does not survive iterated dominance.

Identification A player's strategic bound must be at least as high as her rationality bound: If she is not strategic, then she cannot be rational. So, if she reasons that the other player is not strategic, then she also reasons that the other player is irrational.

Notice, however, that a player's strategic 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.[9] With this in mind, our question is: Does there exist a gap between the strategic 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.

---

[9] Recall from the Introduction: The strategic bound need not correspond to an ability bound. Instead, we use the strategic bound as a vehicle to show that the rationality bound is not entirely determined by limitations in ability.

(ii) The *minimum* level of strategic reasoning 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 strategic. That is, we do not attempt to distinguish rational behavior from strategic behavior. Instead, our identification focuses on reasoning about rationality vs. strategic reasoning. 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 $(\mathrm{a}, \mathrm{b}_*)$. Notice that this behavior—i.e., the strategy $(\mathrm{a}, \mathrm{b}_*)$—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 $(\mathrm{a}, \mathrm{b}_*)$, if he assigns probabilities $\Pr(\mathrm{a}, \mathrm{c}_*) = 9/10$ and $\Pr(\mathrm{b}, \mathrm{a}_*) = \Pr(\mathrm{c}, \mathrm{a}_\mathrm{b}) = 1/20$.) 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 also identify P2 as having a strategic 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 strategic. Importantly, such behavior is inconsistent with a rational P2 believing that P1 is not strategic: Recall, if P2 is rational and believes that P1 is not strategic, then P2's behavior does not vary across G and G_*. Thus, if we observe P2 play the non-constant action profile $(\mathrm{a}, \mathrm{b}_*)$, we must conclude that P2 reasons that P1 is strategic. As such, the minimum level of strategic reasoning consistent with observed behavior is 2; we identify this behavior as having a *strategic bound of 2*.

Second, we identify P2 as having a *strategic bound of 1* if his behavior is consistent with being rational and believing (i.e., assigning probability 1 to the event) that P1 is not strategic. Thus, we identify P2 as having a strategic bound of 1 if and only if he plays a constant strategy. To understand why, recall that if a rational P2 believes that P1 is not strategic, then he plays a constant action profile. Moreover, each constant action profile is consistent with a rational P2 believing that P1 is not strategic. (For instance, it is a best response for P2 to play $(\mathrm{a}, \mathrm{a}_*)$, if he assigns probability 1 to P1 also playing $(\mathrm{a}, \mathrm{a}_*)$.) Notice that this behavior—i.e., the constant strategy $(\mathrm{a}, \mathrm{a}_*)$—is also consistent with a rational P2 believing that P1 is strategic. (For instance, it is a best response for P2 to play $(\mathrm{a}, \mathrm{a}_*)$ if he assigns probabilities $\Pr(\mathrm{a}, \mathrm{c}_*) = 2/5$ and $\Pr(\mathrm{b}, \mathrm{a}_*) = \Pr(\mathrm{c}, \mathrm{b}_*) = 3/10$.) But because we seek the minimum level of strategic reasoning consistent with observed behavior, we identify the strategic bound as 1.

If we identify P2 as having a strategic bound of 1, then we 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; 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 strategic bounds. We will identify a gap between P2's rationality and strategic bounds if he plays a non-constant strategy profile that differs from $(\mathrm{a}, \mathrm{b}_*)$. As an illustration, suppose we observe P2 play $(\mathrm{a}, \mathrm{c}_*)$. 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 is not strategic. Moreover, we have also seen that it is consistent with a rational P2 believing that P1 is strategic. Thus, we identify the subject as having a strategic 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's strategic reasoning as believing—i.e., assigning probability 1 to the event—that P1 is strategic. But, if P2 assigns probability $p = 4/5$ to the event that P1 is strategic, we do not think of this as lack of strategic reasoning. Such a belief would exhibit an ability to engage in interactive reasoning. Thus, we only identify P2's strategic bound as 1 if he believes that P1 is not strategic—i.e., if he assigns probability 0 to the event that P1 is strategic.

This conceptual point has a pragmatic implication for how we identify the strategic bound: We only identify P2's strategic bound if his behavior is consistent with rationality and assigning probability $p \in \{0, 1\}$ to the event that P1 is strategic. We identify P2 as having a strategic bound of 1 if we can take $p = 0$—i.e., if P2's behavior is consistent with rationality and belief that P1 is not strategic. We identify P2 as having a strategic 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 that P1 is strategic, but inconsistent with rationality and belief that P1 is not strategic.[10] By restricting p to be in $\{0, 1\}$, we limit our ability to rationalize the data.

To sum up, we have used this example to illustrate how we can separately identify the strategic 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.

---

[10] Notice, there is a difference between (a) not believing that a player is strategic and (b) believing that the player is not strategic. 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.

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 (Pi's) payoffs depend only on the behavior of Player $(i-1)$ (P$(i-1)$). (We adopt the convention that P0 $\equiv$ P4).

P1's Payoffs

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

P2's Payoffs

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

P3's Payoffs

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

P4's Payoffs

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

(a) Figure G

P1's Payoffs

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

P2's Payoffs

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

P3's Payoffs

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

P4's Payoffs

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

(b) Figure G_*

Figure 3.1. Kneeland's (2015) Ring Game

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 strategic 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 $x(i) \in \{\mathrm{a}, \mathrm{b}, \mathrm{c}\} \times \{\mathrm{a}_*, \mathrm{b}_*, \mathrm{c}_*\}$ indicates the subject's behavior in the role of Pi across both G and G_*. We assume that each subject is rational (and, so, strategic). Thus, we can use the subjects' behavior across both the games and the player roles to provide a lower bound on strategic reasoning and an upper bound on reasoning about rationality. This provides us with a conservative estimate (i.e., an underestimate) of the gap between the strategic and rationality bounds.

Table 3.1 previews of how we identify the rationality and strategic bounds. It focuses on the case where there is no gap between the two 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 Pi—but not in the role of P$(m+1)$. If the subject has a strategic bound of m then, for each $i > m$, the subject plays a constant strategy profile in the role of Pi—but a non-constant strategy profile in the role of Pm. The next subsections elaborate on the identification strategy.

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

Table 3.1. Identifying Bounds: No Gap

3.1 Identifying the Rationality Bounds

Return to the example in Section 2: We pointed out that a rational subject in the role of P1 will play $(\mathrm{a}, \mathrm{c}_*)$. This corresponded to the fact that $(\mathrm{a}, \mathrm{c}_*)$ is a dominant strategy for P1. We also pointed out that a rational subject who believes "P1 is rational" will play $(\mathrm{a}, \mathrm{b}_*)$ in the role of P2. This corresponded to the fact that $(\mathrm{a}, \mathrm{b}_*)$ is the only strategy that survives two rounds of iterated dominance for P2. 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 Pi, 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 Pi, 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 Pi, she plays a strategy that survives m rounds of iterated dominance.[11]

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

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

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 $(x(1), \ldots, x(k))$ is IU. Under iterated dominance, P1 would play $(\mathrm{a}, \mathrm{c}_*)$, P2 would play $(\mathrm{a}, \mathrm{b}_*)$, P3 would play $(\mathrm{b}, \mathrm{a}_*)$, and P4 would play $(\mathrm{a}, \mathrm{c}_*)$. Table 3.2 then summarizes how we identify the rationality bound.

---

[11] Our 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).

Bound	**P1**	**P2**	**P3**	**P4**
Rationality = 1	IU = $(\mathrm{a}, \mathrm{c}_*)$	not IU		
Rationality = 2	IU = $(\mathrm{a}, \mathrm{c}_*)$	IU = $(\mathrm{a}, \mathrm{b}_*)$	not IU	
Rationality = 3	IU = $(\mathrm{a}, \mathrm{c}_*)$	IU = $(\mathrm{a}, \mathrm{b}_*)$	IU = $(\mathrm{b}, \mathrm{a}_*)$	not IU
Rationality = 4	IU = $(\mathrm{a}, \mathrm{c}_*)$	IU = $(\mathrm{a}, \mathrm{b}_*)$	IU = $(\mathrm{b}, \mathrm{a}_*)$	IU = $(\mathrm{a}, \mathrm{c}_*)$

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 $x(2) = (\mathrm{b}, \mathrm{a}_*)$ and $x(4) = (\mathrm{a}, \mathrm{c}_*)$. 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 Strategic Bounds

Return to Section 2: We identified P2's strategic bound based on whether he played a constant versus a non-constant strategy. (Recall, a *constant strategy* is some strategy $(\mathrm{d}, \mathrm{d}_*)$.) In particular, we saw that, if a rational subject in the role of P2 believes "P1 is not strategic," the subject will play a constant strategy. On the other hand, if a rational subject in the role of P2 believes "P1 is strategic," the subject may play a non-constant strategy.

Bound	**P1**	**P2**	**P3**	**P4**
Strategic = 1	Dominant $(\mathrm{a}, \mathrm{c}_*)$	Constant	Constant	Constant
Strategic = 2	Dominant $(\mathrm{a}, \mathrm{c}_*)$	Non-Constant	Constant	Constant
Strategic = 3	Dominant $(\mathrm{a}, \mathrm{c}_*)$		Non-Constant	Constant
Strategic = 4	Dominant $(\mathrm{a}, \mathrm{c}_*)$			Non-Constant

Table 3.3. Identifying Strategic Bound

This is the approach that we will take more generally. Refer to Table 3.3. If we identify a subject as having a strategic bound of k, then the subject plays a non-constant strategy in the role of Pk and a constant strategy in the role of Pj 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.)

To identify the strategic 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 the assumption.[12]) 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.)

The remainder of Section 3.2 spells out the assumptions underlying the identification strategy. Some readers may prefer to skip this material, so as to get more quickly to the results.

3.2.1 Strategic Optimality

We pointed out that a player may be strategic 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_*.

In our analysis, a strategic (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 *strategically optimal* action.

Consider a strategic 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. 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 strategically optimal in G, then c_* is also strategically optimal in G_*. And, likewise, if b (resp. c) is strategically optimal in G, then a_* (resp. b_*) is also strategically optimal in G_*.

The same idea applies to $\mathrm{P}i = \mathrm{P2}, \mathrm{P3}, \mathrm{P4}$. Suppose $\mathrm{P}i$ has the same belief about $\mathrm{P}(i-1)$'s play or adopts a rule of thumb that does not depend on her belief about $\mathrm{P}(i-1)$'s play. Then, if her theory of how to play the game leads her to play $\mathrm{d} \in \{\mathrm{a}, \mathrm{b}, \mathrm{c}\}$ in G, it should lead her to play $\mathrm{d}_* \in \{\mathrm{a}_*, \mathrm{b}_*, \mathrm{c}_*\}$ in G_*—after all, the two games involve the very same payoff matrix for $\mathrm{P}i$.

We can abstract a general principle from these two scenarios: For each player $\mathrm{P}i$ there is a permutation Π_i of i's actions from G to G_* that preserves $\mathrm{P}i$'s payoff matrix. For $i = 1$, the

---

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

permutation maps $\mathrm{a} \mapsto \Pi_1(\mathrm{a}) = \mathrm{c}_*$, $\mathrm{b} \mapsto \Pi_1(\mathrm{b}) = \mathrm{a}_*$, $\mathrm{c} \mapsto \Pi_1(\mathrm{c}) = \mathrm{b}_*$; for $i = 2, 3, 4$, the permutation maps each $\mathrm{d} \mapsto \Pi_i(\mathrm{d}) = \mathrm{d}_*$. If a strategic 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 $\Pi_i(\mathrm{d})$ in G_*. The same conclusion holds for any strategic Pi, 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 $\Pi_{(i-1)}(\mathrm{d})$ in G_*). Put differently, in these cases, Pi's strategic optimality is invariant to the permutation of payoff-equivalent action labels.

3.2.2 Two Principles for Identification

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

We will say that a subject believes strategic behavior if in each player role Pi she satisfies the following principle:

> **Principle of Strategic Reasoning:** Suppose Pi believes that "P$(i-1)$ is strategic and P$(i-1)$ has the same beliefs about P$(i-2)$'s behavior across G and G_*." Then, $\Pr_i(\mathrm{d}, \mathrm{e}_*) > 0$ implies $\mathrm{e}_* = \Pi_{(i-1)}(\mathrm{d})$.

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

To better understand this principle, suppose that P2 believes that P1 is strategic and that P1 has the same beliefs about P4's behavior across G and G_*. P2 can still have non-degenerate beliefs about P1's play. For instance, P2 may assign probability $1/2$ to P1 playing a best response and probability $1/2$ to P1 adopting the lucky-6 rule of thumb. In that case, $\Pr_2(\mathrm{a}, \mathrm{c}_*) = \Pr_2(\mathrm{c}, \mathrm{b}_*) = 1/2$. The Principle of Strategic Reasoning implicitly requires that—across the games G and G_*—P2 has the same belief about the nature of P1's strategic optimality criterion. 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).[13]

---

[13] Another example may be useful: Suppose P3 believes that P2 maximizes her expected payoffs and that P2 assigns probability $5/7 : 2/7$ to $(\mathrm{a}, \mathrm{a}_*) : (\mathrm{c}, \mathrm{c}_*)$. Then P3 believes that P2 is indifferent between playing a and b in G. The Principle of Strategic Reasoning 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

Consider now the case where a subject believes that others are not strategic. To better understand the approach, return to the example (Section 2). We explained that, if P1 is not strategic, 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 is not strategic, 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 take more generally. If a subject believes that others are not strategic, then she reasons that their behavior 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 subjects do not observe the identity of their co-players.)

With this in mind, we will say that a subject believes other players are not strategic if she satisfies the following principle:

> **Principle of Non-Strategic Reasoning:** The subject has the same belief Pr in each player role, i.e, $\mathrm{Pr}_i = \mathrm{Pr}$ for each $i = 1, 2, 3, 4$. Moreover, this belief satisfies $\mathrm{Pr}(\mathrm{a}, \mathrm{a}_*) + \mathrm{Pr}(\mathrm{b}, \mathrm{b}_*) + \mathrm{Pr}(\mathrm{c}, \mathrm{c}_*) = 1$.

We will call a belief for Pi, Pr_i, a *constant belief (for Pi)* if $\mathrm{Pr}_i(\mathrm{a}, \mathrm{a}_*) + \mathrm{Pr}_i(\mathrm{b}, \mathrm{b}_*) + \mathrm{Pr}_i(\mathrm{c}, \mathrm{c}_*) = 1$. The Principle of Non-Strategic Reasoning says that the subject has the same constant belief in each player role.

To better understand this principle, suppose that P2 believes that P1 is not strategic. 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, $\mathrm{Pr}(\mathrm{a}, \mathrm{a}_*) = \mathrm{Pr}(\mathrm{b}, \mathrm{b}_*) = 1/2$. The Principle of Non-Strategic Reasoning implicitly requires, however, that a player has the same belief about the nature of others' non-strategic behavior across 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 $(\mathrm{a}, \mathrm{a}_*)$ in the role of P2, and probability 1 to $(\mathrm{b}, \mathrm{b}_*)$ 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-strategic (an analogue of m-rationality) and the strategic bounds. Call a subject *1-strategic* if, in each player role, she is strategic. Call a subject *2-strategic* if she is *1-strategic* and, in each player role, she satisfies the Principle of Strategic Reasoning. Inductively, say a subject is *k-strategic* if she is $(k-1)$-strategic and she believes (i.e., assigns probability one to the event) that the other player is $(k-1)$-strategic. Say that a subject has a *strategic bound of 1* if she is 1-strategic and satisfies the Principle of Non-Strategic Reasoning.

---

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 strategic optimality across the two games.

Inductively, a subject has a *strategic bound of k* if she is k-strategic and believes that other player has a strategic bound of $(k-1)$.

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

3.2.3 Identifying the Strategic Bounds

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

Strategic Bound of 1 Consider a rational subject who has a strategic bound of 1. By the Principle of Non-Strategic Reasoning, the subject has the same constant belief across player roles. As a consequence, in the roles of $\mathrm{P}i = \mathrm{P2}, \mathrm{P3}, \mathrm{P4}$, this subject must play a constant strategy: If d is a best response for $\mathrm{P}i$, then d_* is also a best response for $\mathrm{P}i$. Because we assume that the subject is not indifferent between any two actions, d_* must be her unique best response.

Identification (Strategic Bound 1)**.** *We identify* $(x(1), x(2), x(3), x(4))$ *as having a* strategic bound of 1 *if there exists a belief* $\Pr$ *on* $\{\mathrm{a},\mathrm{b},\mathrm{c}\} \times \{\mathrm{a}_*, \mathrm{b}_*, \mathrm{c}_*\}$ *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 strategic bound of 1, then $x(1)$ must be the dominant strategy $(\mathrm{a}, \mathrm{c}_*)$. (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, $x \in \{(\mathrm{a}, \mathrm{c}_*)\} \times \{(\mathrm{a}, \mathrm{a}_*), (\mathrm{b}, \mathrm{b}_*), (\mathrm{c}, \mathrm{c}_*)\}^3$. Importantly, there may be observations in $\{(\mathrm{a}, \mathrm{c}_*)\} \times \{(\mathrm{a}, \mathrm{a}_*), (\mathrm{b}, \mathrm{b}_*), (\mathrm{c}, \mathrm{c}_*)\}^3$ that would not be identified as having a strategic 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 strategic bound of 1.

Strategic Bound of 2 Consider a subject who has a strategic bound of 2. In the role of each $\mathrm{P}i$, this subject satisfies the Principle of Strategic Reasoning and believes "$\mathrm{P}(i-1)$ is strategic and believes $\mathrm{P}(i-2)$ is non-strategic." Recall, if $\mathrm{P}(i-1)$ believes that "$\mathrm{P}(i-2)$ is non-strategic," then $\mathrm{P}(i-1)$ has the same belief about $\mathrm{P}(i-2)$'s behavior across G and G_*. Thus, the Principle of Strategic Reasoning says that $\mathrm{P}i$ must believe that $\mathrm{P}(i-1)$'s behavior is invariant to permuting equivalent action labels.

This pins down the subject's belief about the strategies played. In the role of P2, the subject believes that, if P1 plays a (resp. b, resp. c) in G then she plays $\Pi_1(\mathrm{a}) = \mathrm{c}_*$ (resp. $\Pi_1(\mathrm{b}) = \mathrm{a}_*$,

Optimal Given a 2-Strategic Belief	
Constant	Non-Constant
$(\mathrm{a}, \mathrm{a}_*)$	$(\mathrm{a}, \mathrm{b}_*)$ $(\mathrm{a}, \mathrm{c}_*)$ $(\mathrm{b}, \mathrm{a}_*)$ $(\mathrm{b}, \mathrm{c}_*)$ $(\mathrm{c}, \mathrm{a}_*)$

Table 3.4. P2's 2-Strategic Behavior

$\Pi_1(\mathrm{c}) = \mathrm{b}_*)$ in G_*. Thus, P2's belief, namely Pr_2, must satisfy

$$\mathrm{Pr}_2(\mathrm{a}, \mathrm{c}_*) + \mathrm{Pr}_2(\mathrm{b}, \mathrm{a}_*) + \mathrm{Pr}_2(\mathrm{c}, \mathrm{b}_*) = 1.$$

We refer to such a belief as a *2-strategic belief.* In the role of P$i \neq$ P2 , the subject believes that, if P$(i-1)$ plays a (resp. b, resp. c) in G then she plays $\Pi_{(i-1)}(\mathrm{a}) = \mathrm{a}_*$ (resp. $\Pi_{(i-1)}(\mathrm{b}) = \mathrm{b}_*$, $\Pi_{(i-1)}(\mathrm{c}) = \mathrm{c}_*$)) in G_*. Thus, for each Pi = P1, P3, P4, Pr_i is a constant belief for Pi.

A rational subject who has a strategic bound of 2 must play a best response given these beliefs. In the role of P1, this implies that she plays the dominant strategy $(\mathrm{a}, \mathrm{c}_*)$. In the role of P2, this implies that she plays a unique best response given a 2-strategic belief. In the roles of Pi=P3,P4, she must play a best response to a constant belief. Observe that, if $\mathrm{d} \in \{\mathrm{a}, \mathrm{b}, \mathrm{c}\}$ is a best response (for Pi = P3, P4) in G, then $\mathrm{d}_* \in \{\mathrm{a}_*, \mathrm{b}_*, \mathrm{c}_*\}$ is also a best response (for Pi = P3, P4) 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 Pi = P3, P4. With this in mind:

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

(i) $x(1) = (\mathrm{a}, \mathrm{c}_*)$,

(ii) $x(2)$ *is a non-constant strategy that is a unique best response under a 2-strategic 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 strategic bound of 2. Then $x(3)$ and $x(4)$ must be constant strategies. Moreover, $x(2)$ must be optimal under a 2-strategic belief. Referring to Table 3.4, there are five non-constant strategies that are optimal under such a belief. (One non-constant strategy is precluded.) There is one constant strategy, namely $(\mathrm{a}, \mathrm{a}_*)$, that is optimal under such a belief. However, we require that $x(2)$ be non-constant: If $x(2)$ were constant, then the minimum strategic bound consistent with behavior may well be 1. (For instance, the observation $(x(1), x(2), x(3), x(4))$ with $x(1) = (\mathrm{a}, \mathrm{c}_*)$ and $x(2) = x(3) = x(4) = (\mathrm{a}, \mathrm{a}_*)$ would be identified as having a strategic bound of 1.) Even if it were not, observing a constant profile does not provide evidence in favor of the Principle of Strategic Reasoning. Thus, we tie our hands by not-classifying observations with $x(2) = (\mathrm{a}, \mathrm{a}_*)$, unless they can be identified as having a strategic bound of 1. This limits our ability to rationalize the data.

Strategic Bound of 3 Consider a rational subject who has a strategic bound of 3. As before, this subject must play the dominant $(\mathrm{a}, \mathrm{c}_*)$ 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 Pi, such a subject believes "P$(i-1)$ is 2-strategic and believes P$(i-2)$ has a strategic bound of 1."

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

In the role of P3, the subject can play any strategy. To understand why, note that the subject believes "P2 is 2-strategic and believes P1 has a strategic bound of 1." This implies that the subject believes "P2 is strategic and believes that P1 is strategic," i.e., that P2 plays a strategically optimal strategy, given a 2-strategic belief. Because we have taken a broad view of what preferences strategic optimality might represent, it imposes no restrictions on P2's behavior when she holds a 2-strategic belief. Thus, P3 can hold any belief about P2's behavior and, in turn, any strategy of P3 can be a unique best response.[14]

In the role of P4, the subject must play a constant strategy. To understand why, note that the subject believes "P3 is 2-strategic and believes P2 has a strategic bound of 1." This implies that the subject has a belief that "P3 is strategic and P3 believes that P2 plays a constant strategy." (This uses both principles; see Lemma B.1.) Thus, applying the Principle of Strategic Reasoning, 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 (Strategic Bound 3)**.** *We identify* $(x(1), x(2), x(3), x(4))$ *as having a* strategic bound of 3 *if*

(i) $x(1) = (\mathrm{a}, \mathrm{c}_*)$,

(ii) $x(2)$ *is a unique best response given a 2-strategic 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-strategic and believes P1

---

[14] We 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 strategic optimality. For instance, suppose P3 assigns some probability $p \in (0,1)$ to "P2 is rational and believes that P1 has a strategic 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.)

has a strategic bound of 1." However, because we focus on the minimum strategic 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 strategic bound. Second, unlike how we identify a strategic bound of 2, we do not require that $x(2)$ is non-constant. Now a non-constant $x(3)$ provides evidence in favor of the Principle of Strategic Reasoning. Thus, no requirement is needed of $x(2)$ beyond the requirement that it be a best response to a 2-strategic belief.

Strategic Bound of 4 Consider a rational subject who has a strategic 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 strategic 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-strategic and believes P2 has a strategic bound of 2." This implies that the subject has a belief that "P3 is strategic and P3 can hold any belief about P2's play." (To see this, repeat the argument for P3's behavior when there is a strategic bound of 3.) Thus, P4 can hold any belief about P3's behavior and, so, any strategy can be a best response.

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

(i) $x(1) = (\mathrm{a}, \mathrm{c}_*)$*,*

(ii) $x(2)$ *is a unique best response given a 2-strategic 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. A constant strategy could be a best response for a P4 that believes "P3 is 3-strategic and believes that P2 has a strategic bound of 2." However, because we focus on the minimum strategic bound consistent with observed behavior, we would assign such an observation—i.e., an observation with $x(4)$ constant—a lower strategic 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 strategic bound.

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 strategic 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 strategic bound. If all players had a constant permutation, then we would expect constant behavior across the ring games, independent of whether a subject reasons that others are strategic versus others are not strategic. Because P1's permutation is non-constant, we can separate a strategic bound of 1 from a strategic bound of $k = 2, 3, 4$. If, on the other hand, P2's permutation were also not constant, then a subject who has a strategic 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 strategic bound of 1 and subjects with a strategic 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 strategic bound). The rows in grey correspond to the case where there is no gap between the identified rationality and strategic bounds (c.f. Table 3.1.)

Rat	**Str**	**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 strategic 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 strategic 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 strategic 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 strategic 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 $x(2) = (\mathrm{b}, \mathrm{b}_*)$ and $x(3) = x(4) = (\mathrm{a}, \mathrm{a}_*)$.) Second, to identify the observation as having a strategic bound of $k = 2, 3, 4$, $x(2)$ must be a best response under a 2-strategic belief. (This, for instance, rules out an observation with $x(2) = (\mathrm{c}, \mathrm{b}_*)$.) 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 Bounds

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. (That is, in the experiment,

Strategic 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 Strategic Bound from Observed Behavior

the actions a and a_* receive the same label. We use different 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. (Refer to Kneeland's paper for a detailed description of the data.)

There are 80 subjects and thus 80 observations $x = (x(1), x(2), x(3), x(4))$. Table 4.1 shows the identified strategic 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 strategic bound of 4 and 12% are classified as having a strategic bound of 1. Recall, we identify the strategic bound as the minimum strategic bound consistent with the observed data. This has two implications. First, subjects identified as having a strategic bound of 4 may in fact have a higher level of strategic reasoning. (Given the nature of the 4-player ring game, we cannot distinguish a bound of four from higher strategic bounds.) Second, subjects identified as having a strategic bound of 1 may actually have a higher level of strategic reasoning. If so, their behavior would indicate a gap between reasoning about rationality and strategic reasoning. (After all, subjects identified with a strategic 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 rationality and strategic 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 strategic bound of at least 2 and, of those subjects, 23 are identified as having a gap between their strategic and rationality bounds.

We point to two specific features. First, 47% of the subjects identified as having a low rationality

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

Table 4.2. Gap Between Rationality and Strategic Bounds

bound—i.e., a rationality bound of 1 or 2—have a higher strategic bound. (The same is not true for those with a rationality bound of 3.) Second, the gap appears more pronounced for subjects identified as having a higher strategic bound. In particular, 54% of the subject identified as having a strategic bound of 4 have a rationality bound that is strictly less than 4.[15] Put differently, only 46% of the subjects identified as having a strategic bound of 4 also have a rationality bound of 4. But, 76% (resp. 81%) of the subjects who have a strategic 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 strategic 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 strategic bound of 2, 3 or 4 can have a gap between their strategic and rationality bounds. Table 4.3 shows that such gap subjects outperform their no gap counterparts.

	Strategic 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

The fact that gap subjects outperform their no gap counterparts is of interest. As discussed, 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 deliberate decision to *not* believe other players are rational (or engage in reasoning about rationality). The fact that the gap subjects outperform their no gap counterparts suggests that

---

[15] Eight 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 strategic 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.

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

5 The Population-Based Model

In Section 2, we pointed out that, if P2 assigns probability 4/5 to P1 playing the rational strategy $(\mathrm{a}, \mathrm{c}_*)$ and probability 1/5 to the strategic but irrational strategy $(\mathrm{b}, \mathrm{c}_*)$, then P2's best response is to play the non-IU strategy $(\mathrm{a}, \mathrm{c}_*)$. Note two features of this example. First, P2 was identified as having a strategic 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 strategic bounds. Section 5.1 illustrates how such a subject's beliefs about the strategies played can be reinterpreted in terms of beliefs about rationality—even when such a subject is *not* playing an IU strategy. Section 5.2 imposes discipline on those beliefs by studying, what we will call, the *Population-Based Model*. Section 5.3 uses Selten's measure of predictive success (Selten and Krischker, 1982; Selten, 1991) to discuss how well the Population-Based Model explains the data. In the course of doing so, we show that the best fitting model involves non-degenerate beliefs in reasoning about rationality. (Theoretically inclined readers may prefer to skip Section 5.1.)

5.1 Reasoning About Rationality

Recall, to identify the rationality bound, we used the behavior in the role of Pi 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 Pi: We focus on the case where the subject has a rationality bound of $m = i - 1$ and a strategic bound of $k \geq i$. We argue that the behavior in the role of Pi is tied to how the subject reasons about rationality. This will be an input for the Population-Based Model to follow.

Player Role P2 Suppose that we identify a subject as having a rationality bound of 1 and a strategic bound of $k \geq 2$. In that case, the subject plays some $x(2) = (\mathrm{d}, \mathrm{e}_*)$ in the role of P2, where $x(2)$ is optimal under a 2-strategic belief but is not the IU strategy. For instance, in Section 2, we pointed out that the strategy $(\mathrm{a}, \mathrm{c}_*)$ is rational for P2 if she assigns probability 4/5 to P1 playing the rational strategy $(\mathrm{a}, \mathrm{c}_*)$ and probability 1/5 to the strategic but irrational strategy $(\mathrm{b}, \mathrm{c}_*)$. 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 $(\mathrm{a}, \mathrm{c}_*)$ 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-strategic belief, we can find the set of probabilities p_2 so that $x(2)$ is a unique best response under some 2-strategic belief that assigns probability p_2 to the event that "P1 is rational." These sets are given in Table 5.1. Notice, if $x(2)$ is not the IU strategy then p_2 cannot be 1.

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

Table 5.1. Bounds on P2's Belief p_2 of Rationality Given Strategic Bound ≥ 2

Player Role P3 Next, consider a subject who is identified as having a rationality bound of 2 and a strategic bound of $k \geq 3$. In that case, the subject plays some $x(3) = (\mathrm{d}, \mathrm{e}_*)$ 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 strategic optimality, Pr_3 can be any belief.

Suppose that we observe P3 play $x(3) = (\mathrm{b}, \mathrm{c}_*)$. This behavior is a best response under a belief Pr_3 with $\mathrm{Pr}_3(\mathrm{a}, \mathrm{c}_*) = 1$. 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 $\mathrm{Pr}'_3(\mathrm{a}, \mathrm{b}_*) = \mathrm{Pr}'_3(\mathrm{c}, \mathrm{c}_*) = 1/2$. 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 $x(3) = (\mathrm{d}, \mathrm{e}_*)$ is a unique best response under a belief Pr_3. Any such belief can be reconceptualized in terms of beliefs about rationality (or, more precisely, 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. Again, if the observed behavior is not IU, then it cannot be the case that $p_3 = q_2 = 1$, i.e., there must be a departure from full blown reasoning about rationality.

Player Role P4 Finally, consider a subject who is identified as having a rationality bound of 3 and a strategic bound of 4. In that case, the subject plays some $x(4) = (\mathrm{d}, \mathrm{e}_*)$ in the role of P4 that is not the IU strategy, but is optimal under some belief Pr_4 about P3's behavior. As above, such a belief can be reconceptualized in terms of beliefs about rationality (or 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'."

The choice of a triple (p_4, q_3, r_2) is not unique. Table C.3 in Appendix C presents the triples (p_4, q_3, r_2) that rationalize each strategy observed in the data. Again, if the observed behavior is not IU, then it cannot be the case that $p_4 = q_3 = r_2 = 1$, i.e., there must be a departure from full blown reasoning about rationality.

5.2 The Population-Based Model

Above we argued that a subject's beliefs about the strategies played can be reconceptualized in terms of beliefs about rationality (or beliefs about "reasoning about rationality"). For instance, if a subject is identified as having a strategic bound of $k = 4$, 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.' "

Thus, the subject's behavior can be explained by a vector of six parameters $(p_4, q_3, r_2; p_3, q_2; p_2)$. Analogously, if the subject is identified as having a strategic bound of 3 (resp. 2), the subject's behavior can be explained by a vector $(p_3, q_2; p_2)$ (resp. (p_2)) that satisfies the P2-P3 (resp. P2) requirements. Moreover, if there is a gap between the subject's rationality and strategic bounds, then the associated vector is not the **1** vector.

The **Population-Based Model** (**PB Model**) imposes discipline on these vectors of beliefs, by imposing two assumptions: Anonymity and Introspective Beliefs. Each of these assumptions limit beliefs, by exploiting the fact that we observe behavior across player roles. The limitations come in two forms. First, we can explain the behavior of subjects with a strategic bound 4 (resp. 3) by three (resp. two) parameters instead of six (resp. three) parameters. Second, those three (resp. two) parameters must satisfy certain restrictions.

Assumption 5.1 (Anonymity Assumption)**.**

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

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

Anonymity says that, the subject assigns probability p to the event that "P1 is rational" if and only if she assigns probability p to the event that "Pi is rational" for each $i = 2, 3, 4$. And, analogously, for higher levels of reasoning. The idea is this: In the experiment, a subject does not observe who she is playing against and, moreover, the pool of subjects is the same across player roles. As such, the nature of reasoning about the rationality of opponents should not change across player roles.

To exemplify the power of the assumption, suppose we observe a subject play a non-IU strategy $x(2) = (\mathrm{b}, \mathrm{a}_*)$ in the role of P2 and an IU strategy $x(4) = (\mathrm{a}, \mathrm{c}_*)$ in the role of P4. Because $x(4)$ is the IU strategy, one might hope to explain the subjects behavior by parameters $(p_4, q_3, r_2; p_3, q_2; p_2) = (1, 1, 1; p_3, q_2; p_2)$: After all, playing IU in the role of P4 is consistent with a subject who plays a best response given a belief that assigns probability $p_4 = 1$ to

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

But, if that were the case, the subject also assigns probability $p_4 = 1$ to "P3 is rational." Anonymity then requires that the subject also assigns probability $p_2 = 1$ 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, Anonymity implies that p_4 is also at most ${}^2\!/_5$.

Assumption 5.2 (Introspective Beliefs Assumption)**.**

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

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

To understand what Introspective Beliefs delivers, consider a subject who has a strategic bound of at least 3. Then, in the role of P2, she assigns probability p_2 to the event that "P1 is rational." Introspective Beliefs requires that, in the role of P3, she acts as if she assigns probability p_3 to the event that "P2 is rational and p_2-believes 'P1 is rational.'" (As an example, suppose that, in the role of P2, the subject assigns probability 1 to "P1 is rational," and, so, plays the IU strategy. Then, Introspective Beliefs requires that, in the role of P3, she acts as if she assigns p_3 to the event that "P2 is rational and 1-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). Informally, her beliefs about the populations beliefs are determined by introspection.

The Introspective Beliefs assumption directly reduces the dimensionality of the beliefs: The behavior of a strategic bound 4 subject can be described by some $(\alpha^*, \beta^*, \gamma^*)$, where $\alpha^* = p_2 = q_2 = r_2$, $\beta^* = p_3 = q_3$, and $\gamma^* = p_4$. Likewise, the behavior of a strategic bound 3 (resp. 2) subject can be described by some (α^*, β^*) (resp. α^*), where $\alpha^* = p_2 = q_2$, and $\beta^* = q_3$ (resp. $\alpha^* = p_2$). The Anonymity assumption imposes discipline on these parameters.

Proposition 5.1. *If a subject is identified as having a strategic bound* $k = 4$ *(resp.* $k = 3$*), then the subject's behavior can be characterized by parameters* $(\alpha^*, \beta^*, \gamma^*)$ *(resp.* (α^*, β^*)*), where* $\alpha^* = p_2 = q_2 = r_2$*,* $\beta^* = p_3 = q_3$*, and* $\gamma^* = p_4$ *(resp.* $\alpha^* = p_2 = q_2$ *and* $\beta^* = p_3$*) satisfy the P2-P3-P4 (resp. P2-P3) Requirements. Moreover,* $\alpha^* \geq \max\{\beta^*, \gamma^*\}$*.*

5.3 The Best Fitting Population-Based Model

We now bring the PB Model to bear on the experimental data. The model will capture both (i) subjects who have no gap between their strategic and rationality bounds, and (ii) subjects who have a gap between their strategic and rationality bounds. We will describe the behavior of both *no gap subjects* (i.e., i) and *gap subjects* (i.e., ii) using the language of types. Toward that end, a Population-Based Model will consist of a set of types $T = T(\mathrm{n}) \cup T(\mathrm{g})$, where $T(\mathrm{n})$ is a set of *no gap types* and $T(\mathrm{g})$ is a set of *gap types*.

There are four no gap types in the model, i.e., $T(\mathrm{n}) = \{t(1), t(2), t(3), t(4)\}$. Type $t(i)$ represents the case where the strategic and rationality bounds are both i. For each $j \leq i$, type $t(i)$ is expected to play the IU strategy in the role of Pj. For each $j > i$, type $t(i)$ is expected to play a constant strategy. (The set $T(\mathrm{n})$ is the same set in every model.)

Each gap type will be described by a triple of parameters, which implicitly describes both the strategic bound and how it reasons about rationality. A gap type that captures a strategic bound of $k = 4$ (resp. $k = 3$) is characterized by a triple $(\alpha, \beta, \gamma) \in [0, 1]^3$ (resp. $(\alpha, \beta, \mathrm{cb}) \in [0, 1]^2 \times \{\mathrm{cb}\}$) with $(\alpha, \beta, \gamma) \neq (1, 1, 1)$ and $\alpha \geq \max\{\beta, \gamma\}$ (resp. $(\alpha, \beta, \mathrm{cb}) \neq (1, 1, \mathrm{cb})$ and $\alpha \geq \beta$). A gap type that captures a strategic bound of $k = 2$ is characterized by an $(\alpha, \mathrm{cb}, \mathrm{cb}) \in [0, 1) \times \{\mathrm{cb}\}^2$. Types are expected to play in accordance with Proposition 5.1. (See Tables 5.1-C.2-C.3.)

Write $\mathbb{T}(\mathrm{g})$ for the set of all possible gap types. Observe that this set is uncountable. However, many types are essentially equivalent. For instance, the set of predicted strategies for type $(19/20, 19/20, 19/20)$ is the same as the set of predicted strategies for type $(18/20, 18/20, 18/20)$. (Both only contain the IU strategies.) With this, the types in $\mathbb{T}(\mathrm{g})$ can be partitioned into a finite number of subsets, so that two types are in the same partition member if and only if they have the same predicted set of strategies. With this in mind, we focus on models $T = T(\mathrm{n}) \cup T(\mathrm{g})$, where $T(\mathrm{g})$ is a subset of the partition members in $\mathbb{T}(\mathrm{g})$. (Thus, $T(\mathrm{g})$ is finite. Note, we informally describe a gap type as a triple—e.g., (α, β, γ)—but formally represent a gap type as an equivalence class of such triples. No confusion should result.)

We seek to find the "best fitting" PB Model. To talk about the fit of the model, we use Selten's measure of predictive success (Selten and Krischker, 1982; Selten, 1991). Predictive success is an area-based method of trading off the accuracy of the model relative to the precision of the model. The accuracy of the model is measured by the hit rate. Specifically, the *Hit Rate* of T is the probability of observing a datapoint consistent with T. The precision of the model is measured by the relative area. Specifically, the *Relative Area* of T is the proportion of outcomes consistent with the model T. The **Predictive Success** of **T**, viz.

$$\mathcal{PS}(T) = \text{Hit Rate of } T - \text{ Area of } T,$$

is the difference between the accuracy and the precision of the model. Appendix C.3 discusses how this is computed in the data.

We seek a minimal set of types that maximizes the predictive success. That is, we choose $T^* = T^*(\mathrm{n}) \cup T^*(\mathrm{g})$ so that, for each model T: (i) $\mathcal{PS}(T^*) \geq \mathcal{PS}(T)$, and (ii) if $\mathcal{PS}(T) = \mathcal{PS}(T^*)$ then it is not the case that $T \subsetneq T^*$. (Appendix C.3 discusses the minimality requirement.)

There is exactly one minimal model T^* that maximizes predictive success. Table 5.2 describes the associated gap types.[16] There are four such types: one of strategic bound of 2, one of strategic bound 3, and two of strategic bound 4. The predictive success of that model is .843. This says that the model does 84.3% better than uniformly drawing from the set of possible outcomes. To put

---

[16] Recall, a type represents a set of behaviorally equivalent probability parameters. Thus, types are represented by subsets of intervals that satisfy $\alpha \geq \max\{\beta, \gamma\}$. All elements of these intervals are behaviorally equivalent.

	SB 2 Type	SB 3 Type	SB 4 Type	SB 4 Type
α	$[\frac{5}{7}, \frac{7}{8})$	$[\frac{7}{8}, 1]$	$[\frac{7}{8}, 1]$	$[\frac{5}{7}, \frac{7}{8})$
β	cb	$[\frac{15}{62}, \frac{14}{31})$	$[\frac{7}{8}, 1]$	$[\frac{5}{8}, \frac{5}{6})$
γ	cb	cb	$[\frac{1}{3}, \frac{4}{9}]$	$[\frac{5}{9}, \frac{7}{8}]$

Table 5.2. Model That Maximizes Predictive Success

this number in perspective, note that the predictive success of the model with only no gap types is .644. Thus, gap types appear central to understanding the data.

5.4 The Nature of Beliefs

Refer to Table 5.2. Each of the gap types involve non-degenerate beliefs about rationality. For instance, the gap type of strategic bound 2 is described as assigning probability $\alpha \in [\frac{5}{7}, \frac{7}{8})$ to the event that "Pi is rational." The gap type of of strategic bound 3 can be described as assigning probability $\beta \in [\frac{15}{62}, \frac{14}{31})$ to "Pi is rational and believes P$(i-1)$ is rational." Likewise, one of the strategic bound 4 gap types has analogous non-degenerate beliefs—i.e., it can be described as assigning probability $\gamma \in [\frac{1}{3}, \frac{4}{9}]$ to

> "Pi is rational and believes 'P$(i-1)$ is rational and believes P$(i-2)$ is rational'."

The second gap type has non-degenerate beliefs at each level of reasoning—i.e., for the second gap type, $\alpha, \beta, \gamma \in (0, 1)$.

With this framework in place, we are well suited to ask questions of how restrictions on the model impacts the predictive success. We consider two increasingly more restrictive assumptions on beliefs and show that, in each case, not much is lost (in terms of predictive success) by imposing those restrictions. This suggests that the restrictions on beliefs are well-suited to explain the data.

First, the level-k and cognitive hierarchy models implicitly assume that all subjects of a given bound can be described by a single type. This raises the question: Can we explain all gap subjects of a given strategic bound k by a single gap type t_k? The best fitting PB model (Table 5.2) already describes gap subjects of strategic bound $k \in \{2, 3\}$ by a single gap type t_k. How much is lost if we require the same at all strategic levels? We address this question by looking for the the model that maximizes predictive success amongst all *singleton models*, i.e., models where $T(\mathrm{g})$ contains at most one type for each strategic bound $k \in \{2, 3, 4\}$. The model that best fits the data involves $T^*_{\mathrm{s}}(\mathrm{g}) = \{t_{2,\mathrm{s}}, t_{3,\mathrm{s}}, t_{4,\mathrm{s}}\}$, where $t_{2,\mathrm{s}}$ and $t_{3,\mathrm{s}}$ correspond to Table 5.2 and $t_{4,\mathrm{s}}$ is similar to the fourth column in Table 5.2. (It involves $\alpha \in [\frac{5}{7}, \frac{7}{8})$, $\beta \in [\frac{5}{8}, \frac{5}{6})$, and $\gamma \in [\frac{1}{4}, \frac{5}{9}]$.) The model also involves non-degenerate beliefs at each level of reasoning and the predictive success of that model is .828.

Second, suppose that, in addition, we require that *all* gap types reason the same way about rationality, i.e., irrespective of their strategic bound. So, for instance, if subjects of strategic bound 4 are best described by $(\alpha^*, \beta^*, \gamma^*)$, then subjects of strategic bound 3 (resp. 2) are best described by the same $(\alpha^*, \beta^*, \mathrm{cb})$ (resp. $(\alpha^*, \mathrm{cb}, \mathrm{cb})$). (This restricts the parameters of the model, i.e., going from a six parameter model—involving $(\alpha_4, \beta_4, \gamma_4)$, $(\alpha_3, \beta_3, \text{cb})$, and $(\alpha_2, \text{cb}, \text{cb})$—to a three parameter model $(\alpha^*, \beta^*, \gamma^*)$.) In this case, the best fitting model involves the same strategic 4 gap type $t_{4,\text{s}}$. The predictive success of the model is .81.

Overall, these findings suggests that not much is lost—in terms of predictive success—by imposing these auxiliary assumptions on beliefs. Moreover, it suggests that our analysis and findings are robust—that they do not depend on the flexibility of the PB model.

6 Deliberate Choice or Errors?

In Section 4, we argued that there is a gap between the strategic 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. In this section, we argue that this alternate hypothesis is incorrect.

To do so, we will estimate a model of noisy choice, which we refer to as the Random Choice model. In that model, the subjects' rationality bounds are determined by their strategic bounds, but subjects are prone to making mistakes.[17] Thus, the off-diagonal entries in Table 4.2 only reflect those mistakes. We then compare the Random Choice model to the best fitting PB model. To do so, we take two complementary approaches. Both address the extent to which draws from the Random Choice model can replicate the predictions of the best fitting PB model. The first approach does so by looking at an adjusted measure of predictive success. The second approach does so by looking at the distribution of draws predicted by the Random Choice model. These approaches suggest that the data is better explained by a model of Deliberate Choice versus a model of Random Choice.

6.1 Random Choice Model

The Random Choice model takes as given that a subject's rationality bound necessarily coincides with her strategic bound. It interprets what appears to be a gap (between strategic reasoning and reasoning about rationality) as an artifact of errors. For instance, consider a subject who *actually* has a strategic bound of 3. Under the Random Choice model, 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 and play a constant profile in the role of P4. In this model, the subject who actually has a strategic bound of 3 can play the non-IU strategy (a, c_*) in the role of P2—but only if she trembles. If Random Choice is correct, then our Deliberate Choice setting (Sections 3-4) would misidentify this behavior by P2 as reflecting a gap between the strategic and rationality bounds.

A **Random Choice Model** (**RC Model**) consists of a set of types, a prior on those types, and error rates for each type. There are four *types*, viz. r_1, r_2, r_3, r_4; the *probability of type* r_m is

---

[17] The Random Choice model is in the spirit of Kneeland's (2015) approach. Kneeland assumes that the rationality bound coincides with the strategic bound and departures from IU are due to mistakes. She does not model the mistakes explicitly but, instead, classifies a subject's rationality bound based on its closeness to the IU profile.

π_m. Type r_m corresponds to a type who has both a rationality and strategic bound of m. Thus, absent errors, that type is expected to play (i) the IU strategy in the role of P1,...,Pm, and (ii) a constant strategy in any role Pi with $i > m$. More precisely, each type r_m has an *error rate* ε_m. In each player role, r_m plays a predicted strategy with probability $(1 - \varepsilon_m)$ and plays the remaining action profiles with equal probability. For instance, consider type r_3: In each of the roles P1-P2-P3, r_3 plays the IU strategy with probability $(1 - \varepsilon_3)$ and plays any of the eight remaining strategy profiles with probability $\frac{\varepsilon_3}{8}$. In the role of P4, she plays each of the constant strategy profiles with $\frac{1-\varepsilon_3}{3}$ and the remaining six action profiles with probability $\frac{\varepsilon_3}{6}$. (See Appendix D.1 for a formal description of this likelihood function.[18])

	r_1	r_2	r_3	r_4
$\hat{\pi}$	.3645	.122	.096	.417
$\hat{\varepsilon}$	.2777	0	0	.16

Table 6.1. Random Choice Model

We use maximum likelihood, to estimate a probability distribution, $\hat{\pi} = (\hat{\pi}_1, \hat{\pi}_2, \hat{\pi}_3, \hat{\pi}_4)$, and a vector of error rates, $\hat{\varepsilon} = (\hat{\varepsilon}_1, \hat{\varepsilon}_2, \hat{\varepsilon}_3, \hat{\varepsilon}_4)$. This gives the probability distribution and error rates that maximize the likelihood of observing the experimental dataset. The distribution and errors are given in Table 6.1.

An important comment about the estimation: Because the Deliberate Choice model allows for mistakes in the role of P1, subjects may well choose a dominated strategy in the role of P1. This cannot occur in the Deliberate Choice setting studied earlier in this paper. For that reason, earlier in the paper, we restricted attention to the data of subjects who do not choose a dominated strategies. This led to a dataset with $N = 75$ observations. Because we want to give the Random Error model the best chance of explaining the data (relative to the PB model studied earlier), here and through the remainder of the paper, we look at the full dataset—including subjects who do choose a dominated strategy. This gives a dataset of $N = 80$.

6.2 Adjusted Predictive Success

Is the best fitting PB model is better versus worse at predicting the data relative to a theory of RC? To address this question, we first follow an approach proposed in Beatty and Crawford (2011) and use an adjusted measure of predictive success.

To understand the approach, it will be useful to provide a reinterpretation of the original measure. The area of the model can be viewed as the likelihood of seeing some model-prediction, when we draw uniformly from all possible outcomes. When $\mathcal{PS}(T) \approx 0$, the model T performs about as well as uniform draws from the set of all possible outcomes. When $\mathcal{PS}(T) > 0$, the model

---

[18] A previous version of the paper studied a variant of the Random Choice model, in which the likelihood of observing a tremble depends on the payoffs associated with that choice. The qualitative results of that analysis were quite similar to the ones presented here.

T predictions outperform the hypothesis of uniform draws from the set of all possible outcomes.

We are interested in the performance of the PB model relative to draws from the RC Model. So, instead of drawing uniformly from all possible outcomes, we will draw outcomes from the RC model in Table 6.1: The RC model specifies the likelihood of observing any outcome x. (This likelihood depends on both the estimated error rate of each type r_m, viz. $\hat{\varepsilon}_m$, and the estimated probability of each type r_m, viz. $\hat{\pi}_m$.) The *adjusted area* is the likelihood of observing some outcome consistent with the PB model T. With this, we can compute the **Adjusted Predictive Success** of **T**, i.e., the difference between the hit rate of T and the adjusted area. (See Appendix D.2.)

We look at the Adjusted Predictive Success of the best fitting PB model (Table 5.2) relative to the estimated RC model (Table 6.1).[19] This Adjusted Predictive Success is .331. It suggests that our model of deliberate choice—as represented by the PB model of Table 5.2—does more than 33% better than the Random Choice model.

6.3 Simulating the RC Model

Next, we look at the distribution of behavior predicted by the RC model associated with Table 6.1. We argue that it does not match the distribution observed in the data. To see this, we simulate 1000 datasets (of size $N = 75$), by taking draws from that RC model. Within each dataset $d = 1, \dots, 1000$, we match each simulated observation to types of the PB model T^* (Table 5.2): We assign simulated observations to either no-gap types or gap types. If a simulated observation cannot be assigned to any type in T^*, we refer to that as a non-classified (NC) observation.

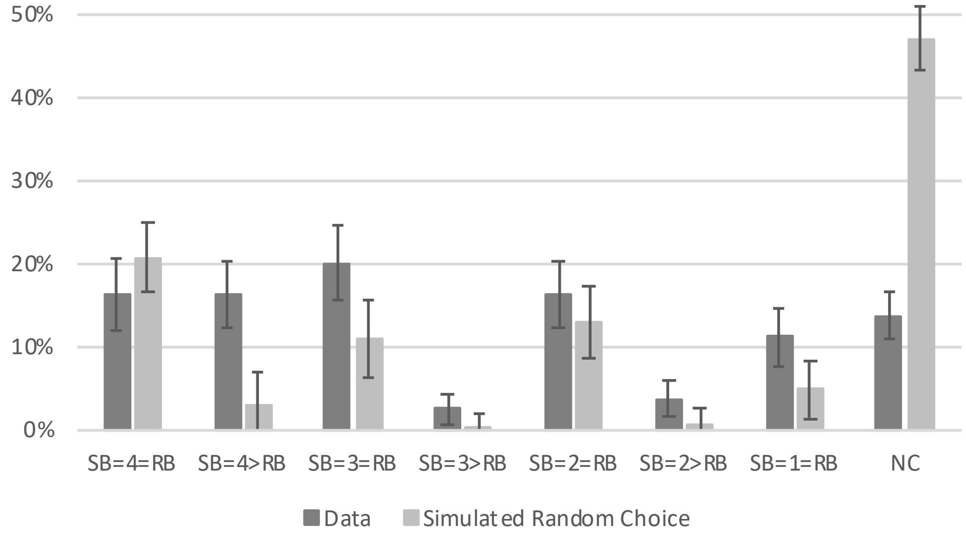

Figure 6.1. Simulated Data with the Alternate Random Choice model

---

[19] The PB model in Table 5.2 also maximizes predictive success when we include all $N = 80$ subjects—not just the 75 rational subjects, as in Section 5. When $N = 80$, the benchmark no-gap types model has a predictive success of .637, whereas the PB model of Table 5.2 has a predictive success of .785. (Likewise, the singleton and three type models that maximize predictive success remain unchanged, when we go from 75 to 80 observations.)

Each simulated dataset d generates a distribution of no gap types, gap types, and NC observations. Figure 6.1 depicts the mean distribution, i.e., taking the average across all 1000 simulations. Importantly, the RC model fails to generate the gap between the strategic and rationality bounds that are observed in the data. Moreover, the simulations suggest that we should observe far more non-classified subjects, if the data was generated by the RC model. These striking differences between the observed distribution and the simulated distributions give us confidence that the identified gap cannot be driven by noise.

References

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

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

Alaoui, L. and A. Penta (2017). Reasoning about others' reasoning. Unpublished Manuscript.

Arad, A. and A. Rubinstein (2012). The 11-20 money request game: A level-k reasoning study. *American Economic Review 102*, 3561–3573.

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

Beatty, T. K. and I. A. Crawford (2011). How demanding is the revealed preference approach to demand? *American Economic Review 101*(6), 2782–95.

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. Iriberri (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.

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.

Selten, R. (1991). Properties of a measure of predictive success. *Mathematical Social Sciences 21*, 153–167.

Selten, R. and W. Krischker (1982). Comparison of two theories for characteristic function experiments.

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

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

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

Level-k 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). The idea is that an L1 reasoner plays a best response given a belief that the other player is not-strategic. For each $k \geq 1$:

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

A subject is an **Lk reasoner** if she chooses an Lk-strategy. A distribution is an Lk-distribution if it has support in the Lk-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 Lk-distribution is degenerate. That is, there is a unique Lk-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 Lk strategy, r_a is an

Ln strategy, and $k \neq n$, then $s_a \neq r_a$. That is, strategies played by Lk reasoners are distinct from strategies played by all lower-order 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 $\ell = 0, 1, \ldots$, 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 $k \geq 1$,

> A **CHk-strategy** is a best response to a CH$(k-1)$-distribution. If there are multiple pure strategy best responses, a CHk-strategy is a mixture that assigns equal probability to each best response.

A subject is a **CHk reasoner** if she plays a CHk-strategy. The CHk-strategy for Ann thus defines a distribution p_a^k over strategies. The CHk-distribution for Ann, denoted q_a^k, is the distribution over strategies if for $\ell \leq k-1$, the fraction of CHℓ reasoners is given by the truncated Poisson distribution and CHℓ 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 CHk reasoners have a unique best response.

Connections: Rk Reasoning versus Lk Reasoning An L1-strategy is rational and, so, an R1-strategy. As a consequence, an L2-strategy is also an R2-strategy: It is a best response under a 1-distribution and a 1-distribution assigns probability one to L1—and, so, R1—strategies. More generally, for any $k \geq 1$, an Lk-strategy is also an Rk-strategy. (In fact, if $j \geq k \geq 1$, then an Lj-strategy is an Rk-strategy.) However, the converse does not hold. There may be strategies that are Rk but not Lk-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.

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

(a) Level-k Reasoning

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

(b) Cognitive Hierarchy Reasoning

Table A.1. Application to the Ring Game

Because an Lk-strategy is an Rk-strategy, Lk reasoning can only lead to behavior that is consistent with degenerate beliefs about rationality. More precisely, an Lk strategy is consistent with rationality and $(k-1)^{\text{th}}$-order belief of rationality. At first this degeneracy may not be obvious: 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 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 $p \in (0,1)$. However, because L2 reasoning is also R2 reasoning, that behavior is also consistent with Ann being rational and assigning probability one to

> "Bob is rational and assigns probability 1 to my rationality."

Connections: Rk Reasoning versus CHk Reasoning A CH1 is rational and, so, an R1-strategy. However, a CH2-strategy need not be an R2-strategy. This is because a CH2-strategy is optimal under a distribution that assigns positive probability to strategies in the support of 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 Rk-strategies are disjoint from the CHk-strategies, for all $k \geq 2$. As a consequence, CHk behavior may only be consistent with non-degenerate beliefs about rationality.

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

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

Appendix B Identifying Strategic Bounds

Lemma B.1. *Let* $Pi = P1, P3, P4$*. If* Pi *believes that* $P(i-1)$ *has a strategic bound of* 1*, then* Pi *has a constant belief.*

Proof. If Pi believes that P$(i-1)$ has a strategic bound of 1, then Pi believes that "P$(i-1)$ is strategic and satisfies the Principle of Non-Strategic Reasoning." It follows from the Principle of Strategic Reasoning that Pi's belief satisfies $\Pr_i(\mathrm{d}, \mathrm{e}_*) > 0$ only if $\pi_{(i-1)}(\mathrm{d}) = \mathrm{e}_*$. Since $i - 1 \neq 1$, it follows that Pi has a constant belief. ∎

Lemma B.2. *If P1 believes that P4 has a strategic bound of* 2*, then P1 has a constant belief.*

Proof. Suppose that P1 believes that P4 has a strategic bound of 2. That is, P1 believes "P4 is 2-strategic and believes that P3 has a strategic bound of 1." When P4 believes that "P3 has a strategic bound of 1," P4 has a constant belief (Lemma B.1). Thus, P1 believes "P4 is 2-strategic and has a constant belief." Thus, applying the Principle of Strategic Reasoning, P1 believes that P4 plays a constant strategy. So, again, P1 has a constant belief. ∎

Strategic Bound	P1	P2	P3	P4	Strategy	Subject
1	$(\mathrm{a}, \mathrm{c}_*)$	$(\mathrm{a}, \mathrm{a}_*)$	$(\mathrm{a}, \mathrm{a}_*)$, $(\mathrm{b}, \mathrm{b}_*)$, $(\mathrm{c}, \mathrm{c}_*)$	$(\mathrm{a}, \mathrm{a}_*)$	3	**7**
1	$(\mathrm{a}, \mathrm{c}_*)$	$(\mathrm{a}, \mathrm{a}_*)$	$(\mathrm{a}, \mathrm{a}_*)$, $(\mathrm{b}, \mathrm{b}_*)$	$(\mathrm{c}, \mathrm{c}_*)$	2	**1**
1	$(\mathrm{a}, \mathrm{c}_*)$	$(\mathrm{b}, \mathrm{b}_*)$	$(\mathrm{b}, \mathrm{b}_*)$	$(\mathrm{a}, \mathrm{a}_*)$, $(\mathrm{b}, \mathrm{b}_*)$, $(\mathrm{c}, \mathrm{c}_*)$	3	**1**
1	$(\mathrm{a}, \mathrm{c}_*)$	$(\mathrm{b}, \mathrm{b}_*)$	$(\mathrm{c}, \mathrm{c}_*)$	$(\mathrm{a}, \mathrm{a}_*)$, $(\mathrm{b}, \mathrm{b}_*)$	2	**0**
1	$(\mathrm{a}, \mathrm{c}_*)$	$(\mathrm{c}, \mathrm{c}_*)$	$(\mathrm{a}, \mathrm{a}_*)$	$(\mathrm{a}, \mathrm{a}_*)$	1	**0**
1	$(\mathrm{a}, \mathrm{c}_*)$	$(\mathrm{c}, \mathrm{c}_*)$	$(\mathrm{c}, \mathrm{c}_*)$	$(\mathrm{a}, \mathrm{a}_*)$, $(\mathrm{b}, \mathrm{b}_*)$	2	**0**
2	$(\mathrm{a}, \mathrm{c}_*)$	$(\mathrm{a}, \mathrm{b}_*)$, $(\mathrm{a}, \mathrm{c}_*)$, $(\mathrm{b}, \mathrm{a}_*)$, $(\mathrm{b}, \mathrm{c}_*)$, $(\mathrm{c}, \mathrm{a}_*)$	$(\mathrm{a}, \mathrm{a}_*)$, $(\mathrm{b}, \mathrm{b}_*)$, $(\mathrm{c}, \mathrm{c}_*)$	$(\mathrm{a}, \mathrm{a}_*)$, $(\mathrm{b}, \mathrm{b}_*)$, $(\mathrm{c}, \mathrm{c}_*)$	45	**16**
3	$(\mathrm{a}, \mathrm{c}_*)$	$(\mathrm{a}, \mathrm{a}_*)$, $(\mathrm{a}, \mathrm{b}_*)$, $(\mathrm{a}, \mathrm{c}_*)$, $(\mathrm{b}, \mathrm{a}_*)$, $(\mathrm{b}, \mathrm{c}_*)$, $(\mathrm{c}, \mathrm{a}_*)$	$(\mathrm{a}, \mathrm{b}_*)$, $(\mathrm{a}, \mathrm{c}_*)$, $(\mathrm{b}, \mathrm{a}_*)$, $(\mathrm{b}, \mathrm{c}_*)$, $(\mathrm{c}, \mathrm{a}_*)$, $(\mathrm{c}, \mathrm{b}_*)$	$(\mathrm{a}, \mathrm{a}_*)$, $(\mathrm{b}, \mathrm{b}_*)$, $(\mathrm{c}, \mathrm{c}_*)$	108	**21**
4	$(\mathrm{a}, \mathrm{c}_*)$	$(\mathrm{a}, \mathrm{a}_*)$, $(\mathrm{a}, \mathrm{b}_*)$, $(\mathrm{a}, \mathrm{c}_*)$, $(\mathrm{b}, \mathrm{a}_*)$, $(\mathrm{b}, \mathrm{c}_*)$, $(\mathrm{c}, \mathrm{a}_*)$	$(\mathrm{a}, \mathrm{a}_*)$, $(\mathrm{a}, \mathrm{b}_*)$, $(\mathrm{a}, \mathrm{c}_*)$, $(\mathrm{b}, \mathrm{a}_*)$, $(\mathrm{b}, \mathrm{b}_*)$, $(\mathrm{b}, \mathrm{c}_*)$, $(\mathrm{c}, \mathrm{a}_*)$, $(\mathrm{c}, \mathrm{b}_*)$, $(\mathrm{c}, \mathrm{c}_*)$,	$(\mathrm{a}, \mathrm{b}_*)$, $(\mathrm{a}, \mathrm{c}_*)$, $(\mathrm{b}, \mathrm{a}_*)$, $(\mathrm{b}, \mathrm{c}_*)$, $(\mathrm{c}, \mathrm{a}_*)$, $(\mathrm{c}, \mathrm{b}_*)$	324	**28**
NC					6071	**1**
Total					6561	**75**

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

Appendix C Section 5

C.1 Interval Calculations

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: $(\mathrm{a}, \mathrm{b}_*)$, $(\mathrm{a}, \mathrm{c}_*)$, and $(\mathrm{c}, \mathrm{a}_*)$. 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. (An expanded version of Table C.3 was used in calculations of the Predictive Success; it is available by request.)

$\mathbf{p_2}$	**Actions**
$[\frac{7}{8}, 1]$	$(\mathrm{a}, \mathrm{b}_*)$
$[\frac{5}{7}, \frac{7}{8})$	$(\mathrm{a}, \mathrm{b}_*)$, $(\mathrm{a}, \mathrm{c}_*)$
$[\frac{62}{133}, \frac{5}{7})$	$(\mathrm{a}, \mathrm{b}_*)$, $(\mathrm{a}, \mathrm{c}_*)$, $(\mathrm{b}, \mathrm{c}_*)$
$(\frac{2}{5}, \frac{62}{133})$	$(\mathrm{a}, \mathrm{b}_*)$, $(\mathrm{a}, \mathrm{c}_*)$, $(\mathrm{b}, \mathrm{c}_*)$,$(\mathrm{a}, \mathrm{a}_*)$

$\mathbf{p_2}$	**Actions**
$[\frac{2}{5}, \frac{2}{5}]$	$(\mathrm{a}, \mathrm{b}_*)$, $(\mathrm{b}, \mathrm{c}_*)$,$(\mathrm{a}, \mathrm{a}_*)$
$(\frac{2}{7}, \frac{2}{5})$	$(\mathrm{a}, \mathrm{b}_*)$, $(\mathrm{b}, \mathrm{c}_*)$,$(\mathrm{a}, \mathrm{a}_*)$,$(\mathrm{b}, \mathrm{a}_*)$
$[\frac{29}{133}, \frac{2}{7}]$	$(\mathrm{b}, \mathrm{c}_*)$,$(\mathrm{a}, \mathrm{a}_*)$,$(\mathrm{b}, \mathrm{a}_*)$
$(\frac{1}{11}, \frac{29}{133})$	$(\mathrm{b}, \mathrm{c}_*)$,$(\mathrm{a}, \mathrm{a}_*)$,$(\mathrm{b}, \mathrm{a}_*)$,$(\mathrm{c}, \mathrm{a}_*)$
$[0, \frac{1}{11}]$	$(\mathrm{b}, \mathrm{c}_*)$, $(\mathrm{b}, \mathrm{a}_*)$,$(\mathrm{c}, \mathrm{a}_*)$

Table C.1. Assigning probability p_2 to Rationality

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

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

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

Table C.3. P4: Probability p_4 to "Rationality and q_3-Belief of 'Rationality and r_2-Belief of Rationality'"

C.2 Proposition 5.1

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 at least probability p to event E_j.

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

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

Property 2 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 probability of exactly q to E_j, then $q \geq p$.

Consider a subject who has a strategic 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 the lemma below. Proposition 5.1 follows immediately from that lemma.

Lemma C.1. *Suppose Anonymity holds. If the subject has a strategic 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 strategic bound of $k = 3, 4$. Notice that $B_2^{p_2}(R_1)$ and so, by Anonymity, $B_i^{p_2}(R_{i-1})$ holds in each player role P3 $\leq$ Pi $\leq$ Pk. At the same time, by Property 1, $\tilde{B}_i^{p_i}(R_{i-1})$ holds. So applying Property 2, $p_2 \geq p_i$.

C.3 Predictive Success

Recall, a Population-Based Model consists of a set of types $T = T(\mathrm{n}) \cup T(\mathrm{g})$, where $T(\mathrm{n})$ is a set of *no gap types* and $T(\mathrm{g})$ is a set of *gap types*. A no gap type $t(i)$ is a type whose rationality and strategic bound are both i. A gap type is characterized by a triple of parameters in $([0, 1] \cup \{\mathrm{cb}\})^3$. (See the main text.)

For any given type $t \in T$, write $S(t)$ for the set of strategy profiles that can be played by type t. These are the *predictions for type t*. Note, $S(t) \subseteq (\{\mathrm{a}, \mathrm{b}, \mathrm{c}\} \times \{\mathrm{a}_*, \mathrm{b}_*, \mathrm{c}_*\})^4$. Write $O(t)$ for the set of observed data points, viz. $\mathbf{x} = (x(1), x(2), x(3), x(4))$, so that $\mathbf{x} \in S(t)$. These are the *observations consistent with type* t.

The measure of Predictive Success of the model is the difference between the hit rate (i.e., the probability of observing a datapoint consistent with the model) and the area (i.e., the proportion of outcomes consistent with the model). Define the **Hit Rate of T** as

$$H(T) = \frac{|\bigcup_{t \in T} O(t)|}{N},$$

where N is the number of datapoints. Define the **Area of T** as

$$A(T) = \frac{|\bigcup_{t \in T} S(t)|}{Z},$$

where Z is the number of possible outcomes. The **Predictive Success of T** is given by

$$\mathcal{PS}(T) = H(T) - A(T).$$

When we compute this measure, we take $N = 75$ and $Z = 729$. That is, we focus on the rational observations (i.e., we take $N = 75$). In keeping with that, we only look at rational strategies when we look the possible outcomes (i.e., we take $Z = 729$).

We choose a minimal model which maximizes predictive success. That is, we choose T^* so that

(i) for each model T, $\mathcal{PS}(T^*) \geq \mathcal{PS}(T)$, and

(ii) if $\mathcal{PS}(T) = \mathcal{PS}(T^*)$ then $\neg(T \subsetneq T^*)$.

In the data, there is a single minimal model that maximizes predictive success.

To understand the import of the minimality criterion, fix a model $T^* = T^*(\mathrm{n}) \cup T^*(\mathrm{g})$ that maximizes predictive success. Suppose there is a type $t_k \in T^*(\mathrm{g})$ in the model and a type $u_k \notin T^*(\mathrm{g})$ not in the model, so that $S(u_k) \subseteq S(t_k)$ and $H(t_k) \subseteq H(u_k)$. (Here we write $H(t)$ for the hit rate of a model that only has the type t, i.e., $|O(t)|/N$.) Then adding u_k to $T^*(\mathrm{g})$ does not change the model's predictive success. We impose minimality so that we look for the smallest set of parameters that explain the data.

This criterion points to a modeling assumption: We assume that the model includes both gap and no gap types. An alternate would be to allow all types to be characterized by parameters (α, β, γ) (resp. $(\alpha, \beta, \mathrm{cb})$ or $(\alpha, \mathrm{cb}, \mathrm{cb})$), with "no gap types" being the special case of $(\alpha, \beta, \gamma) = (1, 1, 1)$ (resp. $(\alpha, \beta, \mathrm{cb}) = (1, 1, \mathrm{cb})$ or $(\alpha, \mathrm{cb}, \mathrm{cb}) = (1, \mathrm{cb}, \mathrm{cb})$). Doing so does not change the conclusions: The minimal set of gap types that maximizes predictive success in our model is also the minimal set of types that maximizes predictive success in that model. To understand why, consider the case of a strategic bound $k = 4$ and note $O(1, 1, 1) = S(1, 1, 1)$ is a singleton that contains the IU strategy. Moreover, for each type t_4, $O(1, 1, 1) \subseteq O(t_4)$ and $S(1, 1, 1) \subseteq S(t_4)$. Thus, a model with both $(1, 1, 1)$ and type t_4 has the same predictive success as a type with only type t_4.[20] Under the alternate approach, there is a model with no gap types, which maximizes predictive success. However that set of types is not minimal. Because we impose minimality, we also explicitly include no gap types in the model.

Appendix D Section 6

The researcher observes an experimental dataset, corresponding to an observation for each subject $n = 1, \ldots, N$. An observation is an action profile for subject n, i.e., some $x^n = (x^n(1), \ldots, x^n(4))$, where $x^n(i) \in \{\mathrm{a}, \mathrm{b}, \mathrm{c}\} \times \{\mathrm{a}_*, \mathrm{b}_*, \mathrm{c}_*\}$ denotes the behavior in the role of Pi. The experimental dataset is given by $\boldsymbol{x} = (x^n)_{n=1}^N$.

Write $x^{\mathrm{rat}}(i)$ for the IU strategy in the role of Pi. (So, $x^{\mathrm{rat}}(1)$ is $(\mathrm{a}, \mathrm{c}_*)$, $x^{\mathrm{rat}}(2)$ is $(\mathrm{a}, \mathrm{b}_*)$, etc.) Let $\eta[x^n, i]$ be an indicator for whether $x^n(i)$ and $x^{\mathrm{rat}}(i)$ agree. So if $x^n(i) = x^{\mathrm{rat}}(i)$, then $\eta[x^n, i] = 1$ and if $x^n(i) \neq x^{\mathrm{rat}}(i)$, then $\eta[x^n, i] = 0$. Let $\kappa[x^n, i]$ be an indicator for whether $x^n(i)$ is a constant strategy. So, $\kappa[x^n, i] = 1$ if $x^n(i)$ is a constant strategy and $\kappa[x^n, i] = 0$ otherwise.

D.1 Random Choice Model

Write $\mathcal{M}^{\mathrm{RC}} = (T^{\mathrm{RC}}, \pi, \boldsymbol{\varepsilon})$ for a **Random Choice** (**RC**) model. It has three components: A set of type $T^{\mathrm{RC}} = \{r_1, r_2, r_3, r_4\}$, a probability distribution on T^{RC}, namely $\boldsymbol{\pi} = (\pi_1, \pi_2, \pi_3, \pi_4)$, and type-specific trembles $\boldsymbol{\varepsilon} = (\varepsilon_m)_{m=1}^4 \in (0,1)^4$. For a subject of type r_m, the probability of observing $x^n(i)$ is

$$p(x^n(i)|r_m, \varepsilon_m) = \begin{cases} (1-\varepsilon_m)^{\eta[x^n,i]} \left(\frac{\varepsilon_m}{8}\right)^{(1-\eta[x^n,i])} & \text{if } i \leq m \\ \left(\frac{1-\varepsilon_m}{3}\right)^{\kappa[x^n,i]} \left(\frac{\varepsilon_m}{6}\right)^{(1-\kappa[x^n,i])} & \text{if } i \geq m+1, \end{cases}$$

and the probability of observing x^n is

$$p(x^n|\boldsymbol{\varepsilon}) = \prod_{i=1}^{4} p(x^n(i)|r_m, \varepsilon_m). \tag{1}$$

To better understand what is involved, consider a subject who has a rationality and strategic bound of $m = 3$. Then, modulo trembles, the subject plays the IU strategy in the role of P1-P2-P3 and a constant strategy in the role of P4. With this in mind, in the role of Pi=P1,P2,P3, we will observe the IU strategy with probability $(1 - \varepsilon_i)$. The RC model further assumes that trembles are independent of the payoffs of the game. So, in the role of Pi=P1,P2,P3, we will observe any non-IU strategy with probability $\frac{\varepsilon_i}{8}$. Likewise, the RC model assumes that in the role of P4, each constant strategy is equally likely. So, in the role of P4, we will observe each constant strategy with probability $\frac{(1-\varepsilon_4)}{3}$ and each non-constant strategy with probability $\frac{\varepsilon_4}{6}$. The RC model assumes that these trembles are independent across player roles.

---

[20] An analogous argument applies to a subject whose strategic bound is 3. However, for a subject with a strategic bound of 2, we need not have that $O(1, \mathrm{cb}, \mathrm{cb}) \subseteq O(t_2)$ and $S(1, \mathrm{cb}, \mathrm{cb}) \subseteq S(t_2)$. (This is because of the discipline imposed on 2-strategic beliefs.) That said, this non-inclusion only holds for types t_2 will negative predictive success.

The likelihood of observing behavior x^n in the model $\mathcal{M}^{\mathrm{RC}}$ is

$$\mathcal{L}(x^n \mid \mathcal{M}^{\mathrm{RC}}) = \sum_{m=1}^{4} \left(\pi_m \times p(x^n|\boldsymbol{\varepsilon})\right).$$

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

$$\ln \mathcal{L}^*(\boldsymbol{x} \mid \mathcal{M}^{\mathrm{RC}}) = \sum_{n=1}^{N} \ln \mathcal{L}(x^n \mid \mathcal{M}^{\mathrm{RC}}).$$

We choose $\hat{\mathcal{M}}^{\mathrm{RC}}$ to maximize $\ln \mathcal{L}^*(\boldsymbol{x} \mid \mathcal{M}^{\mathrm{RC}})$. (Recall, when we do so, we include all—i.e., even irrational—observations. Thus, $N = 80$.) Table 6.1 describes RC model that maximizes the log-likelihood of observing the experimental dataset.

D.2 Adjusted Predictive Success

Let T^* be the PB model that maximizes predictive success (Table 5.2) and let $\hat{\mathcal{M}}^{\mathrm{RC}}$ be the RC model that maximizes the aggregate log-likelihood (Table 6.1). Note, T^* induces a set of predicted outcomes, viz. $O(T^*) = \bigcup_{t \in T^*} O(t)$. For each $x \in O(T^*)$, the likelihood of observing x if the data is generated by $\hat{\mathcal{M}}^{\mathrm{RC}}$ is given by $\mathcal{L}(x \mid \hat{\mathcal{M}}^{\mathrm{RC}})$. The likelihood of observing some outcome predicted by T^* is then

$$\overline{A}(T^* \mid \hat{\mathcal{M}}^{\mathrm{RC}}) := \sum_{x \in O(T^*)} \mathcal{L}(x \mid \hat{\mathcal{M}}^{\mathrm{RC}}).$$

The **Adjusted Predictive Success** of T^* given $\hat{\mathcal{M}}^{\mathrm{RC}}$ is

$$\overline{\mathcal{PS}}(T^* \mid \hat{\mathcal{M}}^{\mathrm{RC}}) = H(T^*) - \overline{A}(T^* \mid \hat{\mathcal{M}}^{\mathrm{RC}}).$$

When we compute the Adjusted Predictive Success we again include all 80 observations. (Refer to Footnote 19.)