Is Bounded Reasoning about Rationality Driven by Limited Ability?*

Amanda Friedenberg† Terri Kneeland‡

June 12, 2023

Abstract

Rationality and common belief of rationality (RCBR) is a standard benchmark in game theory. Yet, a body of experimental research points to departures from RCBR. These RCBR departures are typically viewed as an artifact of limits in the ability to engage in interactive reasoning, i.e., to reason through sentences of the form "I think, you think, I think, etc ..." We provide a novel identification strategy to test the hypothesis that RCBR departures are determined by limits in interactive reasoning. It benefits from not relying on auxiliary measures of "ability" or "sophistication" that can capture distinct concepts. We conduct an experiment based on this identification strategy and show that at least 60% of subjects have RCBR departures that are not an artifact of limited ability to engage in interactive reasoning. Moreover, the experiment provides insight into how subjects reason when they depart from RCBR. The results suggest that subjects' reasoning depends on certain natural heuristics.

The standard approach to game theory implicitly takes as given that players are strategically sophisticated. In epistemic game theory, this is traditionally formalized as rationality and common belief of rationality (RCBR): 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. While RCBR is an important theoretical benchmark, lab experiments have pointed to departures from RCBR. In particular, experiments suggest that there is "bounded reasoning about rationality." See Stahl and Wilson (1995), Costa-Gomes, Crawford, and Broseta (2001), and Costa-Gomes and Crawford (2006).[1]

Why do we observe such departures from RCBR? A common rationale appeals to the difficulties of engaging in interactive reasoning. For Ann to engage in RCBR, she must specify what she thinks about Bob's play, what she thinks Bob thinks about her play, etc. Thus, it requires Ann to reason through infinite-length sentences of the form "I think that you think that I think. . . ." Yet there is evidence from cognitive psychology that subjects are limited in their ability to engage in interactive

---

[*] We thank Adam Brandenburger, Vince Crawford, Julian Romero, Ryan Oprea, and Jim Powell for helpful conversations. Frederic Ebrard, Julen Ortiz De Zarate Pina, and Alex Ralphs provided excellent research assistance. The project was supported by ERC Grant SUExp - 801635. We also thank seminar audiences at Brown, Calgary, NYU, Princeton, SFU, and UCSD. The paper builds on insights from an earlier project with Willemien Kets. We thank Willemien for her contribution to that work and for allowing us to build on those insights.

[†] University of Michigan, amanda@amandafriedenberg.org, http://amandafriedenberg.org

[‡] University College London, t.kneeland@ucl.ac.uk, http://tkneeland.com

[1] Reasoning according to the level-k or cognitive hierarchy models need *not* reflect departures from RCBR. Section 6.A discusses this point. Papers mentioned here show departures from RCBR.

reasoning. (See, e.g., Kinderman, Dunbar, and Bentall, 1998; Stiller and Dunbar, 2007, among others.) In particular, subjects find it difficult to specify these sentences beyond a certain finite length—even if such sentences do not refer to the rationality of others.

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, consider two attorneys, Ann and Bob, who have a history of interacting on cases. Given this past interaction, Ann may not be prepared to believe that Bob is rational. Or, she may believe that Bob is rational, but may not believe that Bob believes that she is rational, etc. Such departures from RCBR are not driven by limited ability, but are nonetheless departures that would impact how Ann strategizes about the current case.

Are departures from RCBR driven by limitations on players' ability to engage in interactive reasoning? Or are there systematic departures from RCBR that cannot be explained by these limitations? The answer has important implications for predicted behavior and for how solution concepts are defined. If the bounds on reasoning about rationality are determined by limits in the ability to engage in interactive reasoning, then there is no room for players to reason beyond how they reason about rationality. (This is by definition.) However, if bounded reasoning about rationality is not determined by limits in the ability to engage in interactive reasoning, then players may well engage in forms of interactive reasoning distinct from reasoning about rationality. This leaves room for the possibility of novel solution concepts, which have not been explored in the literature. Our main contribution is to provide a definitive answer to this question. In so doing, we can learn whether it is or is not valuable to explore such alternative notions of reasoning (or solution concepts). A secondary contribution is to explore what such reasoning may look like; we do so in the context of what we will call a *heuristic beliefs model*, discussed below.

Let us start with the basic question: Are departures from RCBR driven by limitations on players' ability to engage in interactive reasoning? One might hope to address the question directly, by eliciting subjects' hierarchies of beliefs. However, the very act of attempting to elicit the subjects' hierarchies may cause them to engage in higher levels of interactive reasoning than they may otherwise do. In turn, this can suggest evidence that departures from RCBR are *not* driven by limited ability to engage in interactive reasoning, even if they are. This points to a central difficulty: An attempt to measure the ability to reason interactively may alter how players reason. As such, it is difficult to take a direct approach to address the question. Section 6.B discusses why other techniques from the literature—ones that also don't involve directly measuring the ability to reason interactively—do not suffice for the question raised here.

Instead, we take an indirect approach. We develop a conceptual framework that allows us to test the hypothesis that departures from RCBR are determined by limited ability to engage in interactive reasoning. Importantly, the test will allow us to do so by only making use of choice data—without having to rely on belief elicitation. Moreover, the test will provide insight into how players' do reason when their reasoning departs from RCBR.

The key innovation behind the test is the distinction between rational behavior and what we call strategic behavior. 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.[2] A player 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 that departs from subjective expected utility. (So, a rational player conforms to the "textbook notion of rationality," while a strategic player conforms to a "generalized notion of rationality.")

We will distinguish between reasoning about rationality and strategic reasoning. Loosely:

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

- *Strategic Reasoning:* Ann has a *strategic bound* of k if she is strategic, believes that Bob is strategic, believes that Bob believes that she is strategic and so on, up to the statement that includes the word "strategic" k times, but no further.

Because a rational player is strategic, Ann's strategic bound k must be at least as high as her rationality bound m—i.e., $k \geq m$.

Strategic reasoning requires an ability to engage in interactive reasoning. Thus, a subject's "ability bound" must be at least as high as her strategic bound. If a subject's strategic bound is strictly higher than her rationality bound—that is, if $k > m$—then the subject's ability bound is also strictly higher than her rationality bound. In this case, a departure from RCBR is not an artifact of an inability 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; it is not obvious what (if any) observable implications arise from strategic reasoning. In Sections 1-2, we point to observable implications in a particular class of games: permuted ring games (Kneeland, 2015). That is, in permuted ring games, we can identify the strategic bounds as distinct from the rationality bounds.

We conduct an experiment based on this identification strategy. The experiment contains two treatments, each based on a permuted ring game. The two ring games are equivalent from the perspective of RCBR, but are not equivalent from the perspective of strategic reasoning. In fact, players do systematically behave differently across these games. (See Section 4.2.) The two games illustrate that players do have a gap between their rationality bounds and their strategic bounds. In particular, 60% of subjects have a rationality bound that is strictly lower than their strategic bound. (Section 5.5 argues that the cross-treatment differences suggest this gap may well be

---

[2] This terminology is consistent with the use in epistemic game theory. See, e.g., Brandenburger, 2007, Dekel and Siniscalchi, 2014, and Battigalli, Friedenberg, and Siniscalchi, 2021.

higher.) Moreover, on average, subjects with a gap between their rationality and strategic bounds have higher average earnings than their no-gap counterparts. As we will discuss in Section 4.1, this may reflect a deliberate decision to not engage in common belief of rationality.

From the perspective of the identification strategy, we impose few assumptions on the nature of strategic reasoning. This is an important and deliberate feature of our analysis. A more restrictive model of strategic reasoning would have the benefit of yielding a tighter set of predictions; but, it would come at the cost of imposing auxiliary assumptions about reasoning and beliefs—assumptions that would require verification separate from the identification strategy. That said, there are specific assumptions about strategic reasoning that are, arguably, intuitive. Those assumptions take the following form: If Ann believes Bob is strategic, she assigns probability $p \in [0, 1]$ to Bob being rational and probability $(1 - p)$ to Bob playing according to a simple belief-independent heuristic. Moreover, the "heuristic part" of Ann's beliefs is driven by three rules-of-thumb. (These rules-of-thumb correspond to Bob having a theory that maximizes the maximum payoff, the minimum payoff, or the sum of payoffs.) We refer to this model as the *heuristic beliefs model*. The model provides refined predictions of behavior in our two treatments.

Ultimately, the data will either refute the heuristic beliefs model or provide evidence in favor of the model. We show that 88% of subjects play in accordance with the model. Moreover, under the heuristic beliefs model, we would expect the two treatments to differ in their distribution of rationality bounds in exactly the way that they do differ. Section 5.6 augments these results, by looking at ring game data from other studies; it points out that the data in those studies are also consistent with the heuristic beliefs model.

This last fact points to a more general takeaway. Consider a player with a rationality bound of m. If that bound is determined by limited ability to engage in interactive reasoning, then our prediction would be any m-undominated strategy—i.e., any strategy that survives m rounds of deletion of iterated strong dominance. (This follows from standard results, e.g., Tan and da Costa Werlang, 1988.) However, if not, the player may well engage in higher levels of strategic reasoning and that may rule out certain m-undominated strategies. This would suggest a refinement of the m-undominated strategies. In pointing to certain heuristic beliefs that are supported by the data, the results here take a first step toward understanding the nature of that refinement. So, the results can be seen as an integral step in showing how players strategically reason when they depart from the RCBR benchmark.

This paper fits into a growing literature that uses ring games to identify aspects of players' reasoning. (These games were introduced in Kneeland, 2015.) Importantly, this paper's identification strategy is novel and is used to address a distinct question from the literature. In particular, Kneeland's original use ruled out strategic reasoning in an attempt to identify the *exact* level of reasoning about rationality consistent with the data.[3] For our purposes, it suffices to identify the *maximum* level of reasoning about rationality consistent with the data. See Section 2.1. Our paper and identification strategy point out why ring games—as a tool—may be useful for identifying

---

[3] Specifically, Kneeland's exclusion restriction rules out non-rational strategic reasoning.

forms of interactive reasoning, even when the researcher is not attempting to identify an exact level of reasoning about rationality. In this sense, it provides a methodological contribution.

This paper also asks a question distinct from that in Sprenger and Zhao (2021). (See Section 6.D.) In particular, Sprenger and Zhao propose a specific model of strategic reasoning—one that is more restrictive than our heuristic beliefs model. That said, the rationale behind their model has features similar to the heuristics beliefs model. (Specifically, it allows Ann to either believe that Bob is rational or believe that Bob uses a heuristic based on the maximum sum of payoffs he can obtain.) Section 5.6 discusses the differences between the two models and how they are manifested in the data.

The paper proceeds as follows. Section 1 introduces the idea behind the identification strategy and an important comparative static. Section 2 describes the identification strategy more completely. Section 3-4 describe the experimental design and main results. Section 5 introduces the heuristic beliefs model and revisits the experimental results from the perspective of that model. Section 6 fills in some important discussions. Those include connections to level-k and cognitive hierarchy models (6.A), other potential—but ultimately lacking—approaches to identification (6.B), connections to the literature on reasoning across games (6.D) and how the work opens up new avenues for research (6.E).

1 An Illustrative Example

This section uses an example to highlight key features of the identification strategy. In particular, it illustrates how we can separately identify the strategic and rationality bounds.

Figure 1.1 describes 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. Write $(\mathrm{e}, \mathrm{f}_*)$ to denote that a player chooses action e in G and action f_* in G_*. We refer to such an action profile as a *strategy*. We now point to the important features of these games.

For Player 1 (P1, she), the payoff matrix given by G_* is a relabeling of the payoff matrix given by G. That is, there is a permutation $\Pi_1 : \{\mathrm{a}, \mathrm{b}, \mathrm{c}, \mathrm{d}\} \rightarrow \{\mathrm{a}_*, \mathrm{b}_*, \mathrm{c}_*, \mathrm{d}_*\}$ with

$$\Pi_1(\mathrm{a}) = \mathrm{d}_* \qquad \Pi_1(\mathrm{b}) = \mathrm{b}_* \qquad \Pi_1(\mathrm{c}) = \mathrm{a}_* \qquad \Pi_1(\mathrm{d}) = \mathrm{c}_*,$$

so that P1's row $\mathrm{e} \in \{\mathrm{a}, \mathrm{b}, \mathrm{c}, \mathrm{d}\}$ in G corresponds to P1's row $\Pi_1(\mathrm{e}) \in \{\mathrm{a}_*, \mathrm{b}_*, \mathrm{c}_*, \mathrm{d}_*\}$ in G_*. Moreover, the strategy $(\mathrm{a}, \Pi_1(\mathrm{a})) = (\mathrm{a}, \mathrm{d}_*)$ is the dominant strategy.

For Player 2 (P2, he), the payoff matrices in G and G_* are the same. P2 does not have a dominant strategy; in fact, each of P2's strategies is undominated. However, P2 has a unique *iteratively undominated* (*IU*) strategy—namely, $(\mathrm{b}, \mathrm{a}_*)$.

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 maximizes her subjective expected utility—i.e., she chooses a best response given her subjective belief about how P2 plays the game. Then, she would play the dominant strategy $(\mathrm{a}, \Pi_1(\mathrm{a})) = (\mathrm{a}, \mathrm{d}_*)$. Notice that, if she is rational,

P1's Payoffs

P2

		a	b	c	d
	a	17	15	18	16
P1	b	15	14	15	15
	c	6	4	14	8
	d	12	2	2	10

P2's Payoffs

P1

		a	b	c	d
	a	8	14	4	18
P2	b	16	4	2	10
	c	15	17	4	4
	d	14	6	20	10

(a) Figure G

P1's Payoffs

P2

		a_*	b_*	c_*	d_*
	a_*	6	4	14	8
P1	b_*	15	14	15	15
	c_*	12	2	2	10
	d_*	17	15	18	16

P2's Payoffs

P1

		a_*	b_*	c_*	d_*
	a_*	8	14	4	18
P2	b_*	16	4	2	10
	c_*	15	17	4	4
	d_*	14	6	20	10

(b) Figure G_*

Figure 1.1. A Two-Player Example

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 chooses 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 subjective expected utility. 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 $(\text{c}, \Pi_1(\text{c})) = (\text{c}, \text{a}_*)$.

This example points to the approach we take more generally. Because P1's payoffs in G_* are a relabeling of her payoffs in G, P1's strategic behavior varies systematically between the two games: Any theory that leads P1 to play e in G will lead P1 to play $\Pi_1(\text{e})$ in G_*. As a consequence, P1's *strategically optimal behavior* corresponds to the graph of the permutation Π_1—i.e., to the set

$$\text{Strategic}_1 = \{(\text{a}, \Pi_1(\text{a})), (\text{b}, \Pi_1(\text{b})), (\text{c}, \Pi_1(\text{c})), (\text{d}, \Pi_1(\text{d}))\}. \tag{1}$$

We exploit this fact below.[4]

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, thus, strategic). We distinguish between three scenarios.

First, suppose P2 reasons about rationality. By this, we mean that P2 *believes*—i.e., assigns probability 1 to the event—that P1 is rational. In this case, he must assign probability 1 to P1

---

[4] This argument implicitly assumes that either (i) P1's strategically optimal behavior does not depend on her beliefs about P2's play, or (ii) P1 has the same beliefs in G and G_*. Our identification strategy (Section 2) will not impose these assumptions exogenously. Instead, it will focus on an experimental design in which these assumptions are satisfied endogenously.

playing the dominant strategy $(\mathrm{a}, \Pi_1(\mathrm{a})) = (\mathrm{a}, \mathrm{d}_*)$. As P2 is rational, he chooses a best response to this belief; thus, he plays the IU strategy $(\mathrm{b}, \mathrm{a}_*)$. This is a *non-constant strategy* for P2—i.e., a strategy that varies across G and G_*.

Second, suppose P2 believes that P1 is strategic. Then, P2 has a *2-strategic belief*, i.e., a belief that assigns probability 1 to P1 choosing a strategy in Strategic_1 (Equation (1)). Notice, P2 may well believe that P1 is rational—after all, P1's dominant strategy is strategic. But, P2 need not. For instance, he may assign probability $4/5$ to P1 being rational and probability $1/5$ to the lucky-6 strategy $(\mathrm{c}, \Pi_1(\mathrm{c})) = (\mathrm{c}, \mathrm{a}_*)$. If P2 holds this 2-strategic belief, his best response is to play $(\mathrm{d}, \mathrm{a}_*)$. Notice, this too is a non-constant strategy for P2, but one that differs from the IU strategy.

Third, suppose P2 believes 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_*. Since P2's payoffs do not vary across the two games, his best response does not vary. That is, P2's best response is to play a *constant strategy*—i.e., $(\mathrm{a}, \mathrm{a}_*)$, $(\mathrm{b}, \mathrm{b}_*)$, $(\mathrm{c}, \mathrm{c}_*)$, or $(\mathrm{d}, \mathrm{d}_*)$.

Relationship Between Bounds 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 believes that the other player is not strategic, then she also believes that the other player is irrational.

However, a player's strategic bound may be strictly higher than her rationality bound: P2 may believe that P1 is strategic, even though he does not believe that P1 is rational. In this case, bounded reasoning about rationality is not entirely determined by limits in ability.[5] With this in mind, our question is: Does there exist a gap between the strategic and rationality bounds? We next describe how we identify such a gap.

Identification We seek a conservative estimate of the gap between the strategic and rationality bounds. As such, we seek to identify:

(i) The *maximum* level of reasoning about rationality consistent with observed behavior.

(ii) The *minimum* level of strategic reasoning consistent with observed behavior.

To identify these bounds, we assume that observed behavior is rational, in the sense that it is consistent with a player choosing a best response given her belief. As a consequence, we assume that observed 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. Table 1.1 summarizes the identification. We now explain.

The top row of Table 1.1 points to the identification of the rationality bound. Suppose we observe P2 play the IU strategy. In this case, we identify P2 as having a *rationality bound of 2*. His

---

[5] 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.

	Observed Behavior			
	IU	Non-IU Non-Constant 2-Strategic	Constant	Other
Rationality Bound	2	1	1	1
Strategic Bound	2	2	1	NC
Gap	No	Yes	No	NC

Table 1.1. Identification: Observed P2 Behavior
NC = non-classified

behavior is consistent with being rational and believing that P1 is rational. Notice, P2 may also play the IU strategy if he assigns probability strictly less than 1 to P1's rationality. For instance, it is a best response for P2 to play the IU strategy, if he assigns probability .95 to P1 being rational and probability .05 to P1 playing the lucky-6 strategy. But, because we seek the maximum level of reasoning about rationality consistent with observed behavior, we identify the rationality bound as 2. On the other hand, if we observe P2 play a non-IU strategy, we identify P2 as having a *rationality bound of 1*. His behavior is consistent with being rational, but inconsistent with both being rational and believing that P1 is rational.

The middle row of Table 1.1 points to the identification of the strategic bound. If we observe P2 play a constant strategy, we identify P2 as having a *strategic bound of 1*. Each constant strategy is consistent with P2 being rational and believing that P1 is *not* strategic. For instance, it is a best response for P2 to play $(\mathrm{c}, \mathrm{c}_*)$, if he assigns probability 1 to P1 playing $(\mathrm{b}, \mathrm{b}_*)$. This behavior—i.e., the constant strategy $(\mathrm{c}, \mathrm{c}_*)$—is also consistent with a rational P2 believing that P1 is strategic. (P1's action b is associated with the same payoff row as P1's action b_* and so $(\mathrm{b}, \Pi_1(\mathrm{b})) = (\mathrm{b}, \mathrm{b}_*)$ is also strategic.) But because we seek the minimum level of strategic reasoning consistent with observed behavior, we identify the strategic bound as 1.

Suppose, instead, we observe P2 play a non-constant strategy. The strategy is inconsistent with a rational P2 believing that P1 is not strategic. As such, the minimum level of strategic reasoning consistent with observed behavior is 2. If the strategy is also a best response to a 2-strategic belief, we identify P2 as having a *strategic bound of* 2.

Notice, there are non-constant strategies that are not a best response to a 2-strategic belief. As an example, $(\mathrm{a}, \mathrm{b}_*)$ is a non-constant strategy and, so, is not classified as having a strategic bound of 1. But, because $(\mathrm{a}, \mathrm{b}_*)$ is not a best response under any 2-strategic belief, it is also not classified as having a strategic bound of 2. If we were to observe this behavior, we would simply not classify a strategic bound. Section 6.C discusses this point.

The bottom row of Table 1.1 points to when we identify a gap between the rationality and strategic bounds. If P2 plays the IU strategy or a constant strategy, there is no gap identified. However, if we observe P2 play a non-IU strategy that is both non-constant and a best response under a 2-strategic belief, then we identify a gap.

Deliberate Choice or Errors? Our identification strategy rests on an implicit assumption that observed non-constant behavior is a consequence of deliberate choice by a rational P2. An alternate hypothesis is that such non-constant behavior is an artifact of errors or mistakes. We rule out this alternate error hypothesis by studying behavior in a variant of the original game. Refer to Figure 1.2. There are two changes relative to the original game in Figure 1.1. First, in the game G, two rows of P1 are swapped: b and d. Second, the permutation of P1 has changed; it is now

$$\Pi_1(\mathrm{a}) = \mathrm{d}_* \qquad \Pi_1(\mathrm{b}) = \mathrm{a}_* \qquad \Pi_1(\mathrm{c}) = \mathrm{b}_* \qquad \Pi_1(\mathrm{d}) = \mathrm{c}_*.$$

Importantly, the dominant strategy of P1 remains $(\mathrm{a}, \Pi_1(\mathrm{a})) = (\mathrm{a}, \mathrm{d}_*)$ and the payoff matrices of P2 have not changed. This implies that P2's IU strategy remains $(\mathrm{b}, \mathrm{a}_*)$.

P1's Payoffs

P2

		a	b	c	d
P1	a	17	15	18	16
	b	12	2	2	10
	c	6	4	14	8
	d	15	14	15	15

P2's Payoffs

P1

		a	b	c	d
P2	a	8	14	4	18
	b	16	4	2	10
	c	15	17	4	4
	d	14	6	20	10

(a) Figure G

P1's Payoffs

P2

		a_*	b_*	c_*	d_*
P1	a_*	12	2	2	10
	b_*	6	4	14	8
	c_*	15	14	15	15
	d_*	17	15	18	16

P2's Payoffs

P1

		a_*	b_*	c_*	d_*
P2	a_*	8	14	4	18
	b_*	16	4	2	10
	c_*	15	17	4	4
	d_*	14	6	20	10

(b) Figure G_*

Figure 1.2. A Two-Player Example: Changing the Permutation

Because the IU strategy remains unchanged, Figures 1.1 and 1.2 remain unchanged from the perspective of rationality and belief about rationality. If observed non-IU behavior were only an artifact of errors, we would expect to see the same distribution of play in Figures 1.1 and 1.2. To see this, take two examples of errors. First, suppose there is measurement error that arises from players mistakingly clicking on the wrong choice in the computer interface.[6] There is no reason to hypothesize that such mistakes would depend on the game; as such, we would expect to see the same distribution of play across Figures 1.1 and 1.2. Second, suppose there are mistakes made in decision-making. A prominent model in the literature is that players make errors according to a logistic best response. In that case, we would expect to see non-IU play. But, we would expect to see the same distribution of play across Figures 1.1 and 1.2 since P2's payoffs are unchanged across these games.

However, if non-IU behavior results from strategic reasoning, we may well see a different distribution of play across these games. For instance, suppose P2 assigns probability .4 to "P1 is

---

[6] This is the only measurement error that we can see arising. In the experiment itself, we minimize such measurement error by giving subjects the opportunity to review and revise their choices.

rational" and probability .6 to "P1 chooses the lucky-6 strategy." In Figure 1.1, P2 assigns .4 : .6 to $(\mathrm{a}, \mathrm{d}_*) : (\mathrm{c}, \mathrm{a}_*)$; but, in Figure 1.2, P2 assigns .4 : .6 to $(\mathrm{a}, \mathrm{d}_*) : (\mathrm{c}, \mathrm{b}_*)$. In the former case, P2's best response is $(\mathrm{d}, \mathrm{b}_*)$, while in the latter case it is $(\mathrm{d}, \mathrm{a}_*)$.

In sum, if the non-IU play only arises from errors—and not from strategic reasoning—we would expect to see the same distribution of play across Figures 1.1 and 1.2. However, if P2 is rational and believes that P1 is strategic, there may well be different distributions of play across Figures 1.1 and 1.2, even if there are *also* errors. The econometric analysis will show that the observed differences in empirical distributions cannot arise from a single true distribution of play. This means that the differences in distributions cannot solely be an artifact of errors—the two distributions must be fundamentally different.[7]

Assumptions on Strategic Reasoning So far, we have taken a broad view of strategic reasoning. This choice is deliberate. Our goal is to address the question: Is bounded reasoning about rationality determined by limits on the ability to engage in interactive reasoning? For the purpose of addressing this question, we want to make few assumptions about the nature of reasoning and beliefs that players hold. A restrictive model of strategic reasoning would involve auxiliary assumptions about reasoning and beliefs—assumptions that require verification separate from our identification strategy. That said, if a player's ability to reason about rationality is not determined by limits in ability, there is room for particular forms of strategic reasoning to systematically shape beliefs. Understanding the nature of such reasoning can inform new models in epistemic game theory and new solution concepts. With this in mind, we explore a natural class of assumptions, with an eye toward letting the data tell us whether or not these assumptions serve to shape players' beliefs.

Models of Heuristic Beliefs Suppose P2 does not believe that "P1 is rational," but does believe that "P1 is strategic." In this case, P2 believes that P1 has a theory for how to play the game. In principle, P2's model of P1's theory may be complex. But, it seems natural that P2's model is shaped by simple heuristics or rules-of-thumb that P1 can adopt.

Let's take a concrete example of one such a heuristic: P1 seeks to maximize the minimum payoff that she can receive. If P2 reasons that this is the *only* criterion that determines P1's strategic belief, then P2 would assign probability 1 to P1 choosing the strategy that is best according to that criterion—i.e., to P1 choosing the dominant strategy $(\mathrm{a}, \Pi_1(\mathrm{a}))$. If, instead, P2 thinks that this is only one criterion that can shape P1's beliefs, then P2 may not be prepared to assign probability 1 to the dominant strategy. However, if this heuristic is prominent in P2's model, then he should reason that strategies associated with higher minimum payoffs are more likely. Thus, in Figure 1.1, P2's belief should satisfy

$$\mathrm{Pr}_2(\mathrm{a}, \Pi_1(\mathrm{a})) \geq \mathrm{Pr}_2(\mathrm{b}, \Pi_1(\mathrm{b})) \geq \mathrm{Pr}_2(\mathrm{c}, \Pi_1(\mathrm{c})) \geq \mathrm{Pr}_2(\mathrm{d}, \Pi_1(\mathrm{d})), \tag{2}$$

---

[7] We are *not* saying that there are no errors: The empirical distributions may be different from the true distributions. But, the econometric analysis rules out the hypothesis that the two empirical distributions are realizations of a single true distribution—that the only difference between the distributions are errors.

since $(\mathrm{a}, \Pi_1(\mathrm{a}))$ has a minimum payoff of 15, $(\mathrm{b}, \Pi_1(\mathrm{b}))$ has a minimum payoff of 14, etc.

Of course, maximizing the minimum payoff is only one of many possible heuristics. This paper focuses on a class of three heuristics. The heuristics are based on whether P1 seeks to maximize the maximum, minimum, or sum of payoffs that she can receive. Our goal is to understand whether the class, as whole, is prominent in forming P2's strategic beliefs. That is, the goal is to either falsify or provide evidence in favor of the hypothesis that P2's beliefs are shaped by this class of heuristics. As such, the experiments are designed so that P1 ranks strategies the same way irrespective of which heuristic in the class P1 holds. That is, $(\mathrm{e}, \Pi_1(\mathrm{e}))$ has a higher maximum payoff than $(\mathrm{f}, \Pi_1(\mathrm{f}))$ if and only if $(\mathrm{e}, \Pi_1(\mathrm{e}))$ has a higher minimum payoff and a higher payoff sum than $(\mathrm{f}, \Pi_1(\mathrm{f}))$. In this case, we say that $(\mathrm{e}, \Pi_1(\mathrm{e}))$ *heuristically dominates* $(\mathrm{f}, \Pi_1(\mathrm{f}))$.

Observable Implications of Heuristic Beliefs Models We pointed out that, if a rational P2 believes that P1 is strategic, P2's behavior may vary across Figures 1.1 and 1.2. In the specific case in which P2 holds heuristic beliefs, we can provide a tighter prediction on P2's behavior. To better understand, it will be useful to understand how the differences between Figures 1.1 and 1.2 impact the nature of heuristic beliefs.

Recall, there are two differences between Figures 1.1 and 1.2. First, in the game G, two of P1's rows are flipped: b and d. Second, the permutation Π_1 changes. Flipping the rows in G impacts how actions in G are ranked according to the heuristics. In Figure 1.1, b heuristically dominates c which heuristically dominates d; in Figure 1.2, d heuristically dominates c which heuristically dominates b. Changing the permutation changes which actions in G_* correspond to the actions b, c, and d. Taken together, a player who holds heuristic beliefs has beliefs that satisfy

$$\mathrm{Pr}_2(\mathrm{a}, \mathrm{d}_*) \geq \mathrm{Pr}_2(\mathrm{b}, \mathrm{b}_*) \geq \mathrm{Pr}_2(\mathrm{c}, \mathrm{a}_*) \geq \mathrm{Pr}_2(\mathrm{d}, \mathrm{c}_*) \tag{3}$$

in Figure 1.1, but

$$\mathrm{Pr}_2(\mathrm{a}, \mathrm{d}_*) \geq \mathrm{Pr}_2(\mathrm{d}, \mathrm{c}_*) \geq \mathrm{Pr}_2(\mathrm{c}, \mathrm{b}_*) \geq \mathrm{Pr}_2(\mathrm{b}, \mathrm{a}_*) \tag{4}$$

in Figure 1.2. These different requirements of heuristic beliefs impacts which strategies are vs. are not a best response.

As an illustration, suppose that P2 assigns probability .85 to the heuristic best strategy, probability .1 to the heuristic second-best strategy, and probability .05 to the heuristic third-best strategy. While, in Figure 1.1, P2's best response is to play $(\mathrm{c}, \mathrm{a}_*)$, in Figure 1.2, P2's best response is to play the IU $(\mathrm{b}, \mathrm{a}_*)$. So, such a player would depart from IU in Figure 1.1, but not in Figure 1.2. If, instead, P2 assigns probability .75 : .15 : .1 to the heuristic best, second-best, and third-best strategies, P2's best response would depart from IU in both games—but would involve different behavior across the two games: In Figure 1.1, $(\mathrm{c}, \mathrm{a}_*)$ is the best response; in Figure 1.2, $(\mathrm{d}, \mathrm{a}_*)$ is the best response.

Notice, in each of these examples, P2 forms his heuristic belief systematically; that is, he assign the same probability to the heuristic-best vs. second-best, etc. Even in this case, P2's best responses differ across the two games. Payoffs in Figures 1.1 and 1.2 were chosen so that this is true more generally.[8] Figures 5.1-5.2 in Section 5 illustrate that the set of strategies that are a best response to some heuristic belief differs across the two figures—pointing to one observable implication of heuristic beliefs.

Differences in Identified Rationality Bounds There is a second observable implication of Heuristic Belief Models: We would expect to see more subjects play IU in Figure 1.2 than in Figure 1.1. To understand why, suppose that P2 holds a heuristic belief. Then, he will assign some probability $p > 0$ to P1 choosing the heuristic-best strategy; P1's heuristic-best strategy is P1's dominant strategy. (Thus, we can view P2 as holding a belief that assigns probability p to P1 being rational and probability $(1 - p)$ to P1 playing an invariant but irrational strategy.) Certainly, if $p = 1$, P2's IU strategy is a best response to such a belief. However, it will also be a best response if p is less than 1 but "high." Consider the minimum probability p that P2 can assign to the dominant strategy of P1 and have the IU strategy as a best response. This probability is smaller in Figure 1.2 than in Figure 1.1. (This feature is by design—that is, the payoffs in Figures 1.1 and 1.2 were chosen so that this property holds.) So, for instance, a rational P2 that always assigns probability .85 to P1's dominant strategy would not play IU in Figure 1.1, but must play IU in Figure 1.2.

Identified Gap The differences in the identified rationality bounds have consequences for the identification of a gap between the rationality and strategic bounds. To see this, suppose that players' *actual* rationality and strategic bounds do not vary across treatments. With heuristic beliefs, we would expect Figure 1.1 to be associated with a lower *identified* rationality bound relative to Figure 1.2. So the identified gap will be larger in Figure 1.1 relative to Figure 1.2.

Figures 1.1 vs. 1.2 serve different roles in our analysis. Both games allow us to identify a gap. However, the identified gap is a lower bound on the actual gap. We would expect Figure 1.2 to underestimate the gap relative to Figure 1.1. For this reason, we think of Figure 1.1 as better suited for identifying the gap. It will be part of, what we will call, an *Identification* treatment. While Figure 1.2 does allows us to identify a gap, its central role is to allow for the comparative statics discussed above—i.e., to generate a change in the game that allows us to address the question of deliberate choice and to investigate the heuristic beliefs model. So, it will be part of, what we will call, a *Comparative Static* treatment.

2 Identifying the Gap

This section lays out key assumptions in our identification of the rationality and strategic bounds. We do so in the context of a four-player game. Doing so has two benefits relative to the two-player game. First, it allows us to identify a gap between rationality and strategic bounds up to four levels of reasoning. Second, it allows us to use behavior across player roles to tighten the identification.

---

[8] This will not be—and cannot be—true for *every* heuristic belief. For instance, a belief that assigns probability 1 to P1's dominant strategy is a heuristic belief. To distinguish between deliberate choice and noise, the IU strategy is the same across the two games and, so, P2's best response will be the same.

P1's Payoffs

		P4			
		a	b	c	d
P1	a	17	15	18	16
	b	15	14	15	15
	c	6	4	14	8
	d	12	2	2	10

P2's Payoffs

		P1			
		a	b	c	d
P2	a	8	14	4	18
	b	16	4	2	10
	c	15	17	4	4
	d	14	6	20	10

P3's Payoffs

		P2			
		a	b	c	d
P3	a	12	14	7	20
	b	18	4	7	14
	c	8	16	2	6
	d	2	15	17	8

P4's Payoffs

		P3			
		a	b	c	d
P4	a	6	8	16	2
	b	20	12	12	6
	c	14	17	4	6
	d	8	2	15	18

(a) Figure G

P1's Payoffs

		P2			
		a$_*$	b$_*$	c$_*$	d$_*$
P1	a$_*$	6	4	14	8
	b$_*$	15	14	15	15
	c$_*$	12	2	2	10
	d$_*$	17	15	18	16

P2's Payoffs

		P1			
		a$_*$	b$_*$	c$_*$	d$_*$
P2	a$_*$	8	14	4	18
	b$_*$	16	4	2	10
	c$_*$	15	17	4	4
	d$_*$	14	6	20	10

P3's Payoffs

		P2			
		a$_*$	b$_*$	c$_*$	d$_*$
P3	a$_*$	12	14	7	20
	b$_*$	18	4	7	14
	c$_*$	8	16	2	6
	d$_*$	2	15	17	8

P4's Payoffs

		P3			
		a$_*$	b$_*$	c$_*$	d$_*$
P4	a$_*$	6	8	16	2
	b$_*$	20	12	12	6
	c$_*$	14	17	4	6
	d$_*$	8	2	15	18

(b) Figure G_*

Figure 2.1. Identification Game

Figure 2.1 describes two games, G and G_*. Each of the games has a *ring structure* (Kneeland, 2015): 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). Let us point to three important features of the games. First, P1's and P2's payoff matrices are as in Figure 1.1. So, for P1, the payoff matrix in G_* is a relabeling of the payoff matrix in G and P2 has the same payoff matrix across the two games. P3 and P4 also have the same matrices across the two games. Second, while P1 has a dominant strategy, all strategies are undominated for P2, P3, and P4. Third, the games are dominance solvable.

Figure 2.2 also describes two games with a ring structure. The only difference between Figure 2.1 and 2.2 is P1's payoff matrix. Whereas P1's payoffs in Figure 2.1 correspond to Figure 1.1, P1's payoffs in Figure 2.2 correspond to Figure 1.2. Thus, round-for-round, iterated dominance is the same across Figures 2.1 and 2.2. Given the connection to Figures 1.1-1.2, we refer to Figure 2.1 as the *Identification Game* and Figure 2.2 as the *Comparative Static Game*. Each subject will play one of these games, not both.

For the purpose of describing how the rationality and strategic bounds are identified, we will focus our discussion on the Identification Game (Figure 2.1). A subject who is assigned to the Identification Game plays both G and G_*, in each of the player roles. As such, an observation consist of a subject's behavior across eight games—that is, an observation is a profile $x = (x(1), x(2), x(3), x(4))$, where each $x(i) \in \{\mathrm{a}, \mathrm{b}, \mathrm{c}, \mathrm{d}\} \times \{\mathrm{a}_*, \mathrm{b}_*, \mathrm{c}_*, \mathrm{d}_*\}$ 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). Therefore, we can use the subjects' behavior across both games and 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.

The next subsections will elaborate on the identification strategy. But, Section 1 provided the basic insight: If we identify a subject as having a rationality bound of m, the subject plays IU in

P1's Payoffs

P4

		a	b	c	d
	a	17	15	18	16
P1	b	12	2	2	10
	c	6	4	14	8
	d	15	14	15	15

P2's Payoffs

P1

		a	b	c	d
	a	8	14	4	18
P2	b	16	4	2	10
	c	15	17	4	4
	d	14	6	20	10

P3's Payoffs

P2

		a	b	c	d
	a	12	14	7	20
P3	b	18	4	7	14
	c	8	16	2	6
	d	2	15	17	8

P4's Payoffs

P3

		a	b	c	d
	a	6	8	16	2
P4	b	20	12	12	6
	c	14	17	4	6
	d	8	2	15	18

(a) Figure G

P1's Payoffs

P2

		a_*	b_*	c_*	d_*
	a_*	12	2	2	10
P1	b_*	6	4	14	8
	c_*	15	14	15	15
	d_*	17	15	18	16

P2's Payoffs

P1

		a_*	b_*	c_*	d_*
	a_*	8	14	4	18
P2	b_*	16	4	2	10
	c_*	15	17	4	4
	d_*	14	6	20	10

P3's Payoffs

P2

		a_*	b_*	c_*	d_*
	a_*	12	14	7	20
P3	b_*	18	4	7	14
	c_*	8	16	2	6
	d_*	2	15	17	8

P4's Payoffs

P3

		a_*	b_*	c_*	d_*
	a_*	6	8	16	2
P4	b_*	20	12	12	6
	c_*	14	17	4	6
	d_*	8	2	15	18

(b) Figure G_*

Figure 2.2. Comparative Static Game

the role of Pi=P1,…,Pm—but not in the role of P$(m+1)$. If we identify a subject as having a strategic bound of m, the subject plays a non-constant strategy profile in the role of Pm and a constant strategy profile in the role of Pi for each $i > m$.

2.1 Identifying the Rationality Bounds

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 1 to the event that P$(i-1)$ is $(m-1)$*-rational.* Note, a subject is m-rational if and only if she satisfies *rationality and* $(m-1)^{th}$*-order belief of rationality* (R$(m-1)$BR). (See, e.g., Remark 5.1 in Battigalli, Friedenberg, and Siniscalchi, 2021 for this standard result.) Say the subject has a *rationality bound of* m if she is m-rational but not $(m+1)$-rational.

There is a tight connection between m-rationality and the IU strategies: 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 or, equivalently, m rounds of rationalizability. (See, e.g., Tan and da Costa Werlang (1988), Battigalli and Siniscalchi (2002), among others.) In the ring games given by G and G_*, an observation $x = (x(1), x(2), x(3), x(4))$ survives m rounds of iterated dominance if and only if $(x(1), \ldots, x(m))$ is IU. As such:

Identification (Rationality Bound)**.** *Identify* $x = (x(1), x(2), x(3), x(4))$ *as having a* rationality bound of m *if*

(i) $(x(1), \ldots, x(m))$ *is IU; and*

(ii) if $m = 1, 2, 3$*, then* $x(m+1)$ *does not survive IU.*

So, an observation is identified as having a rationality bound of m if the observed behavior is consistent with m-rationality and, when $m < 4$, the observed behavior is inconsistent with $(m+1)$-rationality. If a subject has a rationality bound of $m \leq 4$, then her behavior generates an observation with an identified rationality bound of $n \geq m$. But, if a subject is m-rational for $m > 4$, then her behavior generates an observation with an identified rationality bound of 4.

Importantly, 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 $x(4)$ is IU and $x(2)$ is not IU. Suppose, unlike the identification strategy here, we identified the rationality bound by only looking at observed behavior in a single player role. Then, we would use the fact that $x(4)$ is IU to incorrectly conclude that the subject's rationality bound is 4. However, the subject's behavior in the role of P2—namely, $x(2)$—does not survive two rounds of iterated dominance. As such, the observation x is inconsistent with a rational subject who assigns probability 1 to the event that "P1 is rational." As a consequence, the subject's behavior is inconsistent with 4-rationality. Thus, our identification strategy instead identifies this observation as having a rationality bound of 1.

2.2 Identifying the Strategic Bounds

Section 1 gave the basic idea for how we identify the strategic bound: We identify P2's strategic bound as 1 if we observe P2 play a constant strategy. We identify P2's strategic bound as 2 if his behavior is inconsistent with a strategic bound of 1 and, moreover, it is a best response under a 2-strategic belief. This is the approach that we take more generally.

When we identify a subject's reasoning, we will assume that she is rational—not simply strategic. Moreover, we assume that no subject's behavior (in any player role) is a result of indifference.[9] This implies that each subject chooses amongst pure strategies. In addition, we assume that each subject believes that others choose amongst pure strategies.[10]

With this in mind, we think of the subject as having a belief in each player role Pi about P$(i-1)$'s behavior across G and G_*. Write Pr_i for Pi's distribution on $\{\mathrm{a}, \mathrm{b}, \mathrm{c}, \mathrm{d}\} \times \{\mathrm{a}_*, \mathrm{b}_*, \mathrm{c}_*, \mathrm{d}_*\}$. So, $\mathrm{Pr}_i(\mathrm{e}, \mathrm{f}_*)$ is the probability that Pi assigns to P$(i-1)$ playing the strategy $(\mathrm{e}, \mathrm{f}_*)$. We must specify what it means for the subject to believe that the other players are not strategic vs. are strategic. For that, we introduce two identification assumptions: the Principle of Non-Strategic Reasoning and the Principle of Strategic Reasoning.

Principle of Non-Strategic Reasoning Consider the case in which a subject believes that others are not strategic. Section 1 introduced the basic idea: If Pi believes that P$(i-1)$ is not strategic, she reasons that P$(i-1)$ does not have a theory about how to play the game. As a consequence, she believes that P$(i-1)$'s behavior does not depend on specific parameters of the

---

[9] The data in this paper supports this assumption. This can be seen by applying the argument in Kneeland's (2015) footnote 20 to our dataset. Kneeland's data also supports this assumption. (Footnote 20 discusses rationalizable strategies, but the same argument applies to all strategies.)

[10] This will effectively follow from assumptions about strategic reasoning that we impose below. For clarity, we simply assume this from the onset.

game—including P$(i-1)$'s payoffs. Thus, Pi has the same belief about how P$(i-1)$ plays the game in both G and 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 others 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 others does not depend on the player role.[11]

Say a subject believes others are not strategic if she satisfies:

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

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

Principle of Strategic Reasoning Consider the case in which a subject believes that others are strategic. Section 1 introduced the basic idea: If Pi believes that P$(i-1)$ is strategic, she reasons that P$(i-1)$'s decisions are determined by his payoff matrix and potentially his beliefs about play. For instance, Pi may believe that P$(i-1)$ adopts a rule-of-thumb in which he always chooses a strategy that generates the highest maximum payoff—or, alternatively, that P$(i-1)$ adopts a rule-of-thumb that generates the highest minimum or sum of payoffs. Or, alternatively, Pi may believe that P$(i-1)$ adopts a rule-of-thumb whereby P$(i-1)$, first, chooses a strategy that could lead to a payoff of 6 if such an action exists and, second, if not, plays a best response given his subjective belief about the play of the game. Or, alternatively, Pi may believe that P$(i-1)$ is rational. In each of these cases Pi believes that P$(i-1)$ has some theory about how to play the game; in the latter two examples, she believes that P$(i-1)$'s rule-of-thumb depends on his beliefs.

More generally, when we think of a subject that believes the other players are strategic, we think of a subject that believes that other players choose a strategically optimal strategy. To understand which behavior is strategically optimal, notice the following: For each player Pi, there is a permutation Π_i of i's actions from G to G_* that preserves Pi's payoff matrix. P1's permutation was described on p. 5. For Pi=P2,P3,P4, the permutation is constant—i.e., mapping each $\Pi_i(\mathrm{e}) = \mathrm{e}_*$. 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 e in G would also serve as a rationale for playing the permuted $\Pi_i(\mathrm{e})$ 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 e in G is the same as the probability she assigns to $\Pi_{(i-1)}(\mathrm{e})$ in G_*). Put differently, in these cases, Pi's strategic optimality is invariant to the permutation of payoff-equivalent action labels.

---

[11] This is a reasonable assumption in the experiment, where subjects do not observe the identity of others.

With this in mind, say **Pi is strategic** if, whenever Pi has a constant belief, she plays a strategy $(\mathrm{e}, \Pi_i(\mathrm{e})) \in \{(\mathrm{a}, \Pi_i(\mathrm{a})), (\mathrm{b}, \Pi_i(\mathrm{b})), (\mathrm{c}, \Pi_i(\mathrm{c})), (\mathrm{d}, \Pi_i(\mathrm{d}))\}$. To better understand this definition, suppose that Pi is strategic. If Pi has different beliefs across G and G_*, we cannot draw conclusions about her behavior. However, if she has the same beliefs across G and G_*, any of her *invariant strategies*—i.e., strategies of the form $(\mathrm{e}, \Pi_i(\mathrm{e}))$—is strategically optimal. Moreover, in that case, the invariant strategies are the only strategies that are strategically optimal, provided Pi is not indifferent.

Say a subject believes that others are strategic if, in each player role Pi, she satisfies:

> **Principle of Strategic Reasoning:** Suppose Pi believes that P$(i-1)$ has a constant belief. If $\mathrm{Pr}_i(\mathrm{e}, \mathrm{f}_*) > 0$, then $(\mathrm{e}, \mathrm{f}_*) = (\mathrm{e}, , \Pi_{(i-1)}(\mathrm{e}))$.

The Principle of Strategic Reasoning captures the idea that Pi believes P$(i-1)$ chooses a strategically optimal strategy and is not indifferent between any such strategies. Thus, if Pi believes that "P$(i-1)$ has the same beliefs about P$(i-2)$'s behavior across G and G_*," then the principle requires Pi to believe that P$(i-1)$ plays an invariant strategy.

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_*. For instance, P2 may assign probability $1/2$ to P1 being rational and probability $1/2$ to P1 adopting the lucky-6 rule-of-thumb. In that case, $\mathrm{Pr}_2(\mathrm{a}, \mathrm{d}_*) = \mathrm{Pr}_2(\mathrm{c}, \mathrm{a}_*) = 1/2$. The Principle of Strategic Reasoning implicitly requires that P2 has the same belief about the nature of P1's strategic optimality criterion across G and G_*. So, for instance, P2 cannot assign probability 1 to P1 playing a best response in G and probability 1 to P1 playing the lucky-6 strategy in G_*. If P2 had such a belief, he would believe that P1's theory of how to play the game changes across G and G_*, despite the fact that the two games are payoff-equivalent (up to the permutation of action labels).[12]

Strategic Bounds Say a subject is *1-strategic* if, in each player role, she is strategic. Say a subject is *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 1 to the event) that other subjects are $(k-1)$-strategic. A subject has a *strategic bound of 1* if she is 1-strategic and satisfies the Principle of Non-Strategic Reasoning. Inductively, a subject has a *strategic bound of* k if she is k-strategic and believes that other players have a strategic bound of $(k-1)$.

We will inductively identify the strategic bounds. (Proposition A.1 in Appendix A.1 establishes that this identification corresponds to the minimal strategic bound consistent with observed behavior.) We begin with a strategic bound of $k = 1$.

---

[12] Take another example: Suppose P3 believes that P2 maximizes her expected payoffs and that P2 assigns probability $1/2 : 1/2$ to $(\mathrm{a}, \mathrm{a}_*) : (\mathrm{d}, \mathrm{d}_*)$. 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 P2 uses to resolve this indifference in G gets translated into G_*. For instance, P3 may reason that P2 resolves this indifference in G by choosing the action with the highest minimum 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 minimum payoff (i.e., a_*). If not, P2 would effectively be using a different notion of strategic optimality across the two games.

Identification (Strategic Bound 1)**.** *Identify* $(x(1), x(2), x(3), x(4))$ *as having a* strategic bound of 1 *if there exists a constant belief* Pr *on* $\{\mathrm{a}, \mathrm{b}, \mathrm{c}, \mathrm{d}\} \times \{\mathrm{a}_*, \mathrm{b}_*, \mathrm{c}_*, \mathrm{d}_*\}$ *so that 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 each of $x(1), x(2), x(3)$, and $x(4)$ is a best response under the same constant belief. In the role of P1, the observed behavior $x(1)$ must be non-constant: The dominant strategy $(\mathrm{a}, \mathrm{d}_*)$ is the only strategy that can be a best response to any belief. In the role of Pi=P2,P3,P4, the observed behavior $x(i)$ must be constant: Pi=P2,P3,P4 has the same payoff matrix in G and G_* and so only a constant strategy can be a unique best response to a constant belief. (The uniqueness requirement reflects the assumption that no behavior is an artifact of indifference.) Thus, $x \in \{(\mathrm{a}, \mathrm{d}_*)\} \times \{(\mathrm{a}, \mathrm{a}_*), (\mathrm{b}, \mathrm{b}_*), (\mathrm{c}, \mathrm{c}_*), (\mathrm{d}, \mathrm{d}_*)\}^3$.

Suppose that we have identified strategic bounds of $1, \ldots, k-1$, where $4 > k-1$. We now identify a strategic bound of k.

Identification (Strategic Bound $k \geq 2$)**.** *Identify* $(x(1), x(2), x(3), x(4))$ *as having a* strategic bound of $k \geq 2$ *if it has not been identified as having a strategic bound of* $j < k$ *and the following hold:*

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

(ii) $x(2)$ *is a unique best response under a* 2*-strategic belief; and*

(iii) for each j *with* $4 \geq j > k$, $x(j)$ *is constant.*

To better understand the identification, focus on a subject who has a strategic bound of 2. This is a subject who, in each player role Pi believes that "P$(i-1)$ is strategic and believes P$(i-2)$ is non-strategic." (This uses the fact that she believes that others are strategic and have a strategic bound of 1.) Recall, if P$(i-1)$ believes that "P$(i-2)$ is non-strategic," then P$(i-1)$ has the same belief about P$(i-2)$'s behavior across G and G_*. Thus, the Principle of Strategic Reasoning says that, in the role of Pi, this subject must believe that P$(i-1)$'s behavior is invariant to permuting equivalent action labels. That is, Pi's belief—namely, Pr_i—must satisfy

$$\mathrm{Pr}_i(\mathrm{a}, \Pi_{i-1}(\mathrm{a})) + \mathrm{Pr}_i(\mathrm{b}, \Pi_{i-1}(\mathrm{b})) + \mathrm{Pr}_i(\mathrm{c}, \Pi_{i-1}(\mathrm{c})) + \mathrm{Pr}_i(\mathrm{d}, \Pi_{i-1}(\mathrm{d})) = 1. \tag{5}$$

For Pi=P2, this is a 2-strategic belief and so $x(2)$ must be a unique best response under a 2-strategic belief. For Pi = P3, P4, this is a constant belief. (In that case, each $(\mathrm{e}, \Pi_{i-1}(\mathrm{e})) = (\mathrm{e}, \mathrm{e}_*)$.) Since $x(i) = x(3), x(4)$ must be a unique best response and Pi's matrix is the same across G and G_*, it follows that $x(i) = x(3), x(4)$ is constant.

Note, there are observations in $\{(\mathrm{a}, \mathrm{d}_*)\} \times \{(\mathrm{a}, \mathrm{a}_*), (\mathrm{b}, \mathrm{b}_*), (\mathrm{c}, \mathrm{c}_*), (\mathrm{d}, \mathrm{d}_*)\}^3$ that cannot be identified as having a strategic bound of 1; these are observations $(x(1), x(2), x(3), x(4))$ for which there is no single constant belief Pr where each $x(i)$ is a unique best response under Pr. (See Table D.1.) Those observations are constant in player role Pi=P2,P3,P4. They can be identified with a strategic bound of 2—but only if their behavior in the role of P2 is uniquely optimal under a

2-strategic belief. More generally, an observation is identified as having a strategic bound of 2 if and only if (i) $x(1)$ is the dominant strategy; (ii) $x(2)$ is a unique best response under a 2-strategic belief; and (iii) $x(3), x(4)$ are constant. Many of these observations have $x(2)$ non-constant. (See Tables D.2 and D.4.)

An observation is identified as having a strategic bound of 3 (resp. 4) if and only if (i) $x(1)$ is the dominant strategy; (ii) $x(2)$ is a unique best response under a 2-strategic belief; and (iii) $x(3)$ is non-constant and $x(4)$ is constant (resp. (iii) $x(4)$ is non-constant). Tables D.5-D.6 provide the observations with an identified strategic bound of 3-4.

2.3 Gap Between the Rationality and Strategic Bounds

The identified rationality bound must be less than the identified strategic bound. (See Lemma A.6.) We seek to identify whether there is a gap between the identified rationality and identified strategic bounds.

Identification (Gap versus No Gap)**.**

(i) Identify $(x(1), x(2), x(3), x(4))$ *as having a* gap *if the observation is identified with a rationality bound* $m \geq 1$ *and a strategic bound of* $k > m$*.*

(ii) Identify $(x(1), x(2), x(3), x(4))$ *as having* no gap *if the observation is identified with a rationality bound* $m \geq 1$ *and a strategic bound of* $k = m$*.*

Notice, if an observation is identified as having a strategic bound of 1 or a rationality bound of 4, then the observation is identified as having no gap.

To see the rationale, consider a rational subject who has a strategic bound of k. If the subject's behavior generates an observation with an identified gap, then the subject's rationality bound must be strictly less than k. (See Lemma A.7.) So, the subject's rationality bound cannot be determined by limitations in the ability to engage in interactive reasoning.

Finally, notice that, if the observation is identified as no gap, the behavior may still have been generated by a rational subject whose rationality bound is strictly lower than her strategic bound. (The identified rationality bound may be strictly higher than the subject's rationality bound and the identified strategic bound may be strictly lower than the subject's strategic bound.) That is, the identified gap is a conservative estimate of the gap.

3 Experimental Design

The experiment was conducted online using the ELFE subject pool at UCL.[13] We used ORSEE (Greiner, 2015) for recruitment and collected data from 295 undergraduate subjects. The program

---

[13] We ran three lab sessions in March 2020, just before in-person labs closed. The results from our online experiment are in line with the results from the lab sessions. Because the online experiment is a different medium, we do not include the data from the in-person sessions.

used in the experiment was written in oTree (Chen, Schonger, and Wickens, 2016). Appendix C provides screenshots of the instructions and an example game.

Each subject was assigned to one of two treatments: the Identification (IDENT) Treatment or the Comparative Static (CS) Treatment. IDENT (resp. CS) corresponds to the games in Figure 2.1 (resp. Figure 2.2). Within a subject's assigned treatment, the subject played G and G_* in each of the four player roles (P1, P2, P3, and P4). As such, each subject played eight games. The order of the games was random. Subjects were required to spend at least 90 seconds on each of the games. After making their choices in all games, subjects were given the opportunity to revise their choices. There was no feedback throughout the play of the games and the revision.

The treatment was randomized at the subject level within each session—i.e., within a given experimental session, some subjects were assigned to IDENT and others were assigned to CS. At the end of the experimental session, subjects were randomly and anonymously matched with three other participants in the same treatment. One of the eight games was selected for payment; the same game was selected for the four subjects matched together. Subjects were paid based on their action and the actions of their randomly matched counterparts. Subjects received the GBP value of their payoff in the selected game.

Because the experiment was conducted online, we implemented several features important to preserve anonymity, decrease dropout rates, increase attention, and ensure that subjects understood the experiment.[14] We highlight several of these features. First, we adopted a Zoom protocol similar to that used by EBEL at UCSB. (Appendix C describes the protocol.) Second, we had a completion fee of 3.5 GBP. Third, we implemented an incentivized quiz; a screenshot of the quiz can be found in Appendix C. Subjects who answered all the quiz questions correctly on the first try received a bonus of 3 GBP. Moreover, subjects who did not answer the quiz correctly within three tries were automatically assigned to a low-stakes version of the game (played against a computer).[15] Fourth, on each page of the experiment, subjects could both reveal the instructions and anonymously ask questions via a chat box. That is, they did not have to navigate between the experimental interface and the Zoom app to either see instructions or chat with the experimenter.

On average, subjects earned 15 GBP plus the completion fee. In addition 89% of subjects answered all the quiz questions correctly on the first try and, so, also earned the quiz fee. Subjects were paid by bank transfer using Wise.

---

[14] *Ex ante*, we could not ensure that there would be no dropout. As such, we had to have an experimental design that was robust to dropout. This is why we matched subjects into groups of four after subjects completed the play of the game. (Note, this does not affect the incentives of participants.) If we end up with a session in which a treatment does not have a multiple of four, we complete the incomplete group with behavior from subjects in a complete group.

[15] Only one subject was assigned to the low-stakes game. Because the incentives in that game were quite different from Figures 2.1-2.2, that subject's data is omitted from the dataset.

	Rationality Bound			
Strategic Bound	RB=1	RB=2	RB=3	RB=4
SB=1	39	–	–	–
SB=2	36	21	–	–
SB=3	17	17	4	–
SB=4	40	46	3	44

Table 4.1. Gap Between Identified Rationality and Identified Strategic Bounds

	Rationality			
Strategic	RB=1	RB=2	RB=3	RB=4
SB=1	15	–	–	–
SB=2	20	12	–	–
SB=3	12	8	3	–
SB=4	27	24	3	16

(a) Identification Treatment (IDENT)

	Rationality			
Strategic	RB=1	RB=2	RB=3	RB=4
SB=1	24	–	–	–
SB=2	16	9	–	–
SB=3	5	9	1	–
SB=4	13	22	0	28

(b) Comparative Static Treatment (CS)

Table 4.2. Gap Between Identified Rationality and Identified Strategic Bounds: Treatment

4 Experimental Results

We collected data for 295 subjects. Of those, 147 subjects were assigned to IDENT and 148 subjects were assigned to CS. Thus, there are 295 observations of the form $x_i = (x_i(1), x_i(2), x_i(3), x_i(4))$. We use the approach laid out in Section 2 to assign each x_i a rationality bound (if possible) and a strategic bound (if possible).

Of the 295 subjects, 27 chose a dominated strategy and, thus, we cannot assign a rationality bound to those subjects. Of the remaining 268 subjects, one chose strategies that are inconsistent with our assumptions about strategic reasoning and, thus, we cannot assign a strategic bound to that subject. As such, 28 subjects fall outside the purview of our analysis. So, our analysis focuses on the behavior of the remaining 267 subjects.

4.1 The Gap

Table 4.1 reports the interaction between the identified rationality and strategic bounds. Observations with an identified gap are those that fall on the off-diagonal. Notice, 159 observations are identified as having a gap. This constitutes 60% of the observations. Table 4.2 highlights that the prevalence of an identified gap is not a treatment-specific effect. In IDENT, 67% of the observations are identified as having a gap, and in CS, 51% of the observations are identified as having a gap. (When we discuss the heuristic beliefs model, we return to discuss differences in these distributions.)

The identified gap reflects a stark contrast between the distribution of identified rationality and strategic bounds. The majority of subjects have a low identified rationality bound: 49% have a rationality bound of 1, and 32% have a rationality bound of 2. By contrast, 50% of subjects have an identified strategic bound of 4.

On average, subjects identified as having a gap fare better relative to their no-gap counterparts.

	Identified Strategic Bound			
	SB=1	SB=2	SB=3	SB=4
No Gap	14.89	15.70	14.06	13.31
Gap	-	15.63	15.44	14.61

Table 4.3. Average Expected Payoffs

To illustrate this, in each player role, we compute each subject's (i.e., observation's) expected earnings given the empirical distribution of observed behavior. (This empirical distribution includes the behavior of all subjects—both classified and non-classified.) We take the average across player roles to compute each subject's average expected earnings. Table 4.3 reports the average earnings across observations identified with a strategic bound of k and a gap (resp. no gap).

Subjects identified with both a strategic bound of $k \geq 3$ and a gap outperform their no-gap counterparts. In fact, for subjects with an identified strategic bound of $k \geq 3$, we can reject the hypothesis that the true average payoffs of the no-gap subjects are greater than or equal to the true average payoffs of their gap counterparts. (A 1-sided two-sample t-test returns a p-value of .0001 for $k = 3$ and .0000 for $k = 4$. We take 'reject' to mean 'reject at the 5% significance level.') However, we cannot reject the hypothesis that, when the identified strategic bound is 2, average payoffs are the same across gap and no-gap subjects. (The 2-sided two-sample t-test returns a p-value=.6123.)

The fact that gap subjects outperform their no-gap counterparts is of interest. Bounds on reasoning about rationality are often interpreted as limits on the ability to engage in interactive reasoning. However, such bounds may instead reflect a deliberate decision to *not* believe that other players are rational (or reason about rationality). The fact that the gap subjects outperform their no-gap counterparts may further indicate that some subjects are capable of reasoning about rationality at higher levels, but simply choose not to do so.

4.2 Distribution of Play

Identifying the strategic bound rests on an implicit assumption that observed non-constant behavior is a consequence of deliberate choice—and not *only* an artifact of errors. Section 1 (p. 9) provided a method for how we can reject the null hypothesis that non-constant non-IU behavior is an artifact of errors: If P2's behavior were only an artifact of errors, then we would expect to observe the same distribution of P2 play across the two treatments. Thus, we can reject the null hypothesis if we observe a different distribution of P2 play in IDENT vs. CS. This very same argument applies equally to the distribution of P3's and P4's behavior. By contrast, rational subjects with a strategic bound of $k \geq 2$ *may* induce different distributions of play in the roles of Pi=P2,. . .,Pk.

With this in mind, we look at the distribution of Pi's play for subjects with an identified strategic bound $k \geq i$. Figure 4.1 illustrates that, in each player role Pi=P2,P3,P4, these empirical distributions are different across treatments. In particular, in each player role, there is substantially more IU play in CS vs. IDENT. Conversely, in each player role, there is some non-constant strategy

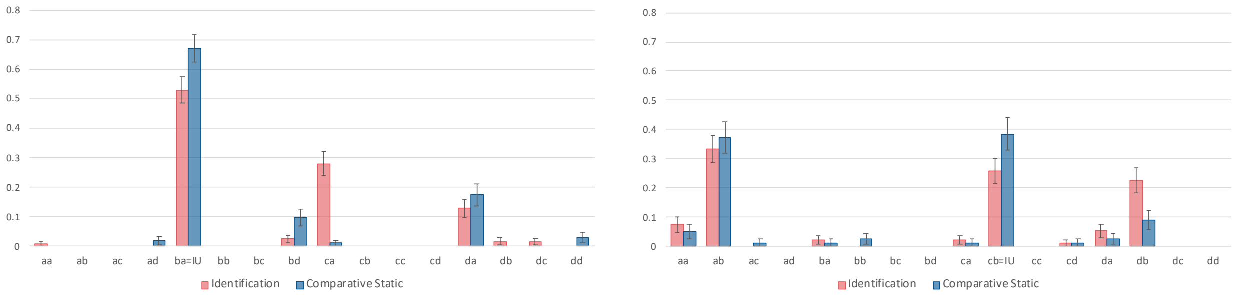

(a) P2 Distribution Given SB ≥ 2

(b) P3 Distribution Given SB ≥ 3

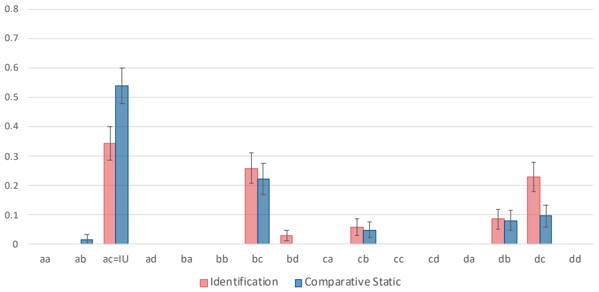

(c) P4 Distribution Given SB ≥ 4

Figure 4.1. Pi Distribution of Play Given SB $\geq i$: Error bars represent bootstrapped standard errors

$(\mathrm{e}, \mathrm{f}_*)$ for which there is substantially more $(\mathrm{e}, \mathrm{f}_*)$ play in IDENT vs. CS. For instance, in the role of P2, 28% of subjects play $(\mathrm{c}, \mathrm{a}_*)$ in IDENT, but only 1% of subjects play it in CS. Similarly, for $(\mathrm{d}, \mathrm{b}_*)$ in the role of P3 (23% in IDENT vs. 9% in CS) and $(\mathrm{d}, \mathrm{c}_*)$ in the role of P4 (23% in IDENT vs. 10% in CS).

Figure 4.1 highlights how the empirical distributions differ across player roles. It leaves open the possibility that the empirical distributions are generated by a single underlying distribution. That is, it leaves open the possibility that the distributions differ only in terms of errors.[16] A chi-square test of homogeneity allows us to test the hypothesis that the empirical distributions differ only in terms of errors—that the two *true* distributions are, in fact, the same. We implement a conservative application of this test, one that makes it difficult to reject the hypothesis of homogeneity.[17] This allows us to reject the hypothesis that the *true* distribution of P2's play (conditional on an identified strategic bound of $k \geq 2$) is independent of the treatment. (The p-value is .0001.) The same conservative application of the chi-square test of homogeneity does not allow us to reject the hypothesis that the *true* distribution of P3's (resp. P4's) play conditional on an identified strategic bound of $k \geq 3$ (resp. $k = 4$) is independent of the treatment. (The p-values are .6575 and .8405,

---

[16] Note a subtle distinction: While the distribution of errors is expected to be the same across treatments, the realization of errors may well be different across treatments.

[17] In each player role, there are strategies that are not played in either treatment. These zero-frequency strategies cause a problem in computing the chi-square test statistic. The statistic itself can be corrected for by effectively ignoring zero-frequency strategies. A conservative approach is to, nonetheless, take the number of degrees of freedom to be one less than the number of strategies, i.e., 15; this is what we report in the text. A more permissive approach would be to take the number of degrees of freedom to be one less than the number of strictly-positive-frequency strategies. Under this permissive approach, the p-values are .0000, .1978, and .1396 respectively.

respectively.) However, the fact that we can reject this null for *some* player allows us to reject the hypothesis that the non-constant non-IU behavior is an artifact of errors.

Two features of the cross-treatment distributions of play provide auxiliary evidence that the data is generated by strategic reasoning. First, Figure 4.1 highlights the fact that there are many non-constant strategy profiles that are simply not played. If the non-constant data were generated by errors, we might expect to see more of these strategies played; that is, we might expect some of these non-constant strategies to be played even if only by a few subjects. (We would expect to see these actions played if the dataset were large.) By contrast, as the discussion of the heuristic beliefs model (Section 5) will drive home, the observed pattern is natural if the data is generated by strategic reasoning.

Second, consider a subject whose *actual* strategic bound is k. If the subject were assigned to both treatments then, in player roles $\mathrm{P}i > \mathrm{P}k$, she should play the same strategy across both. As a consequence, for subjects with an identified strategic bound of k, we should observe the same distribution of play across treatments. Indeed, this is what we observe. Moreover, these distributions are (almost) degenerate. Subjects with an identified strategic bound of 1 play $(\mathrm{d}, \mathrm{d}_*)$ in the role of P2, $(\mathrm{a}, \mathrm{a}_*)$ in the role of P3, and $(\mathrm{b}, \mathrm{b}_*)$ in the role of P4; this behavior is irrespective of the treatment. Almost all subjects with an identified strategic bound of $k = 2, 3$ play analogously in the role of $\mathrm{P}i > \mathrm{P}k$.[18]

5 Heuristic Belief Model

The experimental results highlight that a substantial fraction of subjects have a gap between their identified strategic and identified rationality bounds. This implies that their ability to reason about rationality is not determined by limits in ability. As a consequence, there is room for particular forms of strategic reasoning to systematically shape beliefs. This section introduces a natural class of assumptions on strategic reasoning, with an eye toward letting the data tell us whether or not these assumptions serve to shape players' beliefs.

5.1 Key Assumptions

The heuristic beliefs model captures the idea that, if a player's strategic reasoning departs from reasoning about rationality, then her reasoning is shaped by certain natural heuristics. This is captured by two interrelated assumptions: simple-strategic beliefs and ordered-heuristic beliefs.

To better understand these assumptions, note that strategically optimal behavior can depend on both payoffs and beliefs. Rationality is one such form of strategic optimality. But, there are many forms of strategic optimality that depend on payoffs and not on beliefs. For instance, the lucky-6 rule-of-thumb depends on payoffs but not on beliefs. When a subject's notion of strategic

---

[18] If the identified strategic bound is strictly lower than the actual strategic bound, then subjects with the same identified strategic bound may well play differently across treatments. The fact that subjects play the same across treatments suggests that the identified and actual strategic bounds may coincide.

optimality does not depend on her beliefs, we think of the subject as adopting a rule-of-thumb or heuristic.

Simple-strategic beliefs is an assumption that strategic reasoning is shaped by a belief that either other players are rational or they adopt a belief-independent heuristic. That is, Pi believes that P$(i-1)$'s strategically optimal but irrational behavior depends *only* on P$(i-1)$'s payoff matrix (and not his beliefs). So, if Pi believes that P$(i-1)$ is strategic, she believes that P$(i-1)$ is either rational or P$(i-1)$ plays an invariant strategy. With this in mind, write

$$\mathrm{Inv}_i = \{(a, \Pi_i(a)), (b, \Pi_i(b)), (c, \Pi_i(c)), (d, \Pi_i(d))\}$$

for the set of invariant strategies of Pi.

Assumption 5.1. *Pi has a **simple-strategic belief*** Pr_i *if* $\mathrm{Pr}_i(\mathrm{e}, \mathrm{f}_*) > 0$ *implies that either* $(\mathrm{e}, \mathrm{f}_*)$ *is rational or* $(\mathrm{e}, \mathrm{f}_*) \in \mathrm{Inv}_{i-1}$.

Under the assumption of simple-strategic beliefs, a player's belief can be decomposed into a "rational part" and an "invariant (but irrational) part." In the example of Section 1, we implicitly assumed that P2 has a simple-strategic belief. (See footnote 4.)

Ordered-heuristic beliefs pertain to the relative weight that a subject assigns to invariant but irrational strategies. Section 1 explained the idea: If P$(i-1)$ prefers $(f, \Pi_{i-1}(f))$ to $(e, \Pi_{i-1}(e))$ according to a class of natural heuristics, then Pi should assign a higher weight $(f, \Pi_{i-1}(f))$ vs. $(e, \Pi_{i-1}(e))$. (Refer back to p. 11 for the rationale.)

With this in mind, write $(f, \Pi_i(f)) \geq_i^{\max} (e, \Pi_i(e))$ (resp. $(f, \Pi_i(f)) \geq_i^{\min} (e, \Pi_i(e))$) if the maximum (resp. minimum) payoff that Pi can achieve by choosing f is at least as high as the maximum (resp. minimum) payoff that Pi can achieve by choosing e. Likewise, write $(f, \Pi_i(f)) \geq_i^{\mathrm{sum}} (e, \Pi_i(e))$ if the sum of Pi's payoffs associated with f is at least as high as the sum of Pi's payoffs associated with e. Say $(f, \Pi_i(f))$ **heuristic dominates** $(e, \Pi_i(e))$ for Pi if $(f, \Pi_i(f)) \geq_i^{\max} (e, \Pi_i(e))$, $(f, \Pi_i(f)) \geq_i^{\min} (e, \Pi_i(e))$, and $(f, \Pi_i(f)) \geq_i^{\mathrm{sum}} (e, \Pi_i(e))$; in that case, write $(f, \Pi_i(f)) \geq_i^{\square} (e, \Pi_i(e))$. Note, the games were designed so that, for each Pi=P1,P2,P3, $\geq_i^{\max}$, $\geq_i^{\min}$, $\geq_i^{\mathrm{sum}}$ are complete orders that coincide. So, for Pi=P1,P2,P3, $\geq_i^{\square}$ is a complete order.[19]

Assumption 5.2. *Pi satisfies **ordered-heuristic beliefs** if Pi has a simple-strategic belief* Pr_i *and the following holds: If* $(f, \Pi_{i-1}(f))$ *heuristic-dominates* $(e, \Pi_{i-1}(e))$ *and* $(e, \Pi_{i-1}(e))$ *is irrational, then* $\mathrm{Pr}_i(f, \Pi_{i-1}(f)) \geq \mathrm{Pr}_i(e, \Pi_{i-1}(e))$.

The assumption of heuristic beliefs requires that, if Pi assigns positive probability to an irrational strategy $(e, \Pi_{i-1}(e))$ and $(f, \Pi_{i-1}(f))$ has a higher maximum/minimum/sum than $(e, \Pi_{i-1}(e))$, then Pi assigns higher probability to $(f, \Pi_{i-1}(f))$ over $(e, \Pi_{i-1}(e))$. We refer to a belief that satisfies this condition as an *ordered-heuristic belief.* (Note, an ordered-heuristic belief is, by definition, a simple-strategic belief.)

---

[19] More precisely, $\geq_i^{\max}$ and $\geq_i^{\mathrm{sum}}$ are complete and strict orders that coincide. Moreover, for Pi=P1,P2,P3, $(f, \Pi_i(f)) \geq_i^{\max} (e, \Pi_i(e))$ implies $(f, \Pi_i(f)) \geq_i^{\min} (e, \Pi_i(e))$. So, for P$i$=P1,P2,P3, $\geq_i^{\square}$ is a complete strict order that coincides with $\geq_i^{\max} = \geq_i^{\mathrm{sum}}$. Note, $(f, \Pi_4(f)) \geq_4^{\max} (e, \Pi_4(e))$ does not imply $(f, \Pi_4(f)) \geq_4^{\min} (e, \Pi_4(e))$; as such, $\geq_4^{\square}$ is not a complete order. However, because only P1 cares about P4's behavior and we restrict attention to rational observations, this will not matter.

Let us pick up on two comments made in Section 1 about the ordered-heuristic beliefs assumption. First, the assumption rests on three rules of thumb: maximum, minimum, and sum. We think that each of these rules of thumb is reasonable. For Pi=P1,P2,P3, the relations $\geq_i^{\max}$, $\geq_i^{\min}$, $\geq_i^{\text{sum}}$ coincide; as such, we need not "choose" between the criteria. (Again, this feature is by experimental design.) Instead, we seek to understand whether or not they jointly impact strategic reasoning. Second, we also don't rule out the possibility that other rules of thumb may be at play. In fact, that very possibility is why we allow $\Pr_i(e, \Pi_{i-1}(e)) > 0$ even if $(f, \Pi_{i-1}(f))$ heuristic dominates $(e, \Pi_{i-1}(e))$. Instead, we view these rules of thumb as being prominent in how players engage in strategic reasoning. Third, we do not insist that these rules of thumb are, in fact, an important determinant of beliefs. Instead, we seek to allow the data to weigh in on this question. As we will explain, it does so in two ways: first, in terms of the cross-treatment distribution of play and, second, in term of the cross-treatment distribution of rationality bounds.

5.2 Heuristic Reasoning and Bounds

Say a subject is *1-heuristic* if, in each player role, she is strategic. Say a subject is *2-heuristic* if she is 1-heuristic and, in each player role, (i) she has an ordered-heuristic belief, and (ii) she believes that, if a subject is rational, the subject is not indifferent between any two strategies. For $k \geq 3$, a subject is *k-heuristic* if she is $(k-1)$-heuristic and believes that others are $(k-1)$-heuristic. She has a *heuristic bound of 1* if she is 1-heuristic and satisfies the Principle of Non-Strategic Reasoning. Inductively, a subject has a *heuristic bound of* k if she is k-heuristic and believes that others have a heuristic bound of $(k-1)$.

If a subject is k-heuristic, the subject is also k-strategic; likewise, if a subject has a heuristic bound of k, the subject also has a strategic bound of k. (See Lemmata B.1-B.2.) The key is that, if a subject has a simple-strategic belief, then she must believe that others satisfy the Principle of Strategic Reasoning. So, this new model of strategic reasoning—the **Heuristic Beliefs (HB) model**—is nested in the original model of strategic reasoning—what we now refer to as the **Strategic Beliefs (SB) model**.

We identify an observation $(x(1), x(2), x(3), x(4))$ as having a heuristic bound of 1 if we identify it as having a strategic bound of 1. To identify an observation as having a heuristic bound of $k \geq 2$, it will be useful to inductively define i-heuristic beliefs (an analogue to 2-strategic beliefs). Set $\text{HB}_1 = \{(\text{a}, \text{d})\}$—i.e., the set that contains the dominant strategy for P1. Call $\Pr_i$ a *i-heuristic belief* if it is an ordered-heuristic belief of Pi that assigns probability one to $\text{HB}_{i-1} \cup \text{Inv}_{i-1}$. Set HB_i as the set of strategies of Pi that are a unique best response under an i-heuristic belief.

Identification. *Identify* $(x(1), x(2), x(3), x(4))$ *as having a heuristic bound of* $k \geq 2$ *if it has not been identified as having a heuristic bound of* $j < k$ *and the following hold:*

(i) for each $j = 1, \ldots, k$, $x(j) \in \text{HB}_j$*; and*

(ii) for each $j > k$, $x(j)$ *is a constant strategy.*

An observation $(x(1), x(2), x(3), x(4))$ is identified as having a heuristic bound of k if and only if the observation is consistent with a heuristic bound of k and there is no other $j < k$ so that the observation is consistent with a heuristic bound of j. Proposition B.1 in the Appendix shows that this identification corresponds to the minimal heuristic bound consistent with observed behavior, provided a k-heuristic subject satisfies $(k-1)^{\text{th}}$-order belief of non-indifference. (Note, by definition, a 2-heuristic subject believes non-indifference; this new assumption is an analogue for higher-orders.)

If an observation is identified as having a heuristic bound of k, it is identified as having a strategic bound of k. (See Lemma B.3.) But, an observation can be identified as having a strategic bound of k even if it is not identified as having any heuristic bound. So, the HB model provides tighter predictions than the SB model. (See Tables D.2-D.6 in the Appendix.) Despite the tighter set of predictions, 88% of rational observations are consistent with the HB model, in that they are identified as having some heuristic bound. We next discuss treatment differences that we would expect—and observe—under the HB model.

5.3 Distribution of Play

The HB model has implications for the expected distributions of play: If a rational subject has a heuristic bound of $k \geq i \geq 2$, then the subject's behavior in the role of Pi lies in HB_i. Since a 2-heuristic belief depends on the treatment, the set HB_2 varies based on the treatment. As a consequence, for each $k \geq 3$, the sets HB_k will vary by treatment.

Distribution of P2's Play To better understand the consequences for P2's play, consider a rational subject with a heuristic bound of $k \geq 2$. In the role of P2, the subject plays a strategy that is a unique best response to a 2-heuristic belief. We point out two properties of this belief. First, because the belief is a simple-strategic belief, it assigns probability $p \in [0, 1]$ to the dominant strategy $(\text{a}, \Pi_1(\text{a})) = (\text{a}, \text{d}_*)$ and probability $(1 - p)$ to the irrational but invariant $\{(\text{b}, \Pi_1(\text{b})), (\text{c}, \Pi_1(\text{c})), (\text{d}, \Pi_1(\text{d}))\}$. This irrational invariant set differs across treatments since Π_1 differs across treatments. Second, because the belief is an ordered-heuristic belief, it assigns weight in accordance with P1's heuristic dominance order. The heuristic dominance order differs across treatments—above and beyond the permutation changing across treatments—since the rows b and d are flipped for P1.

To make this more concrete, suppose this belief of P2 assigns probability zero to the "heuristic worst" strategy—i.e., the strategy that is "worst" according to P1's heuristic dominance order. So, in IDENT, $\text{Pr}_2(\text{d}, \Pi_1(\text{d})) = \text{Pr}_2(\text{d}, \text{c}_*) = 0$, and in CS, $\text{Pr}_2(\text{b}, \Pi_1(\text{b})) = \text{Pr}_2(\text{b}, \text{a}_*) = 0$. Under this assumption, we can draw P2's belief in a two-dimensional simplex. Figure 5.1 depicts this simplex; any point in the large triangle specifies the probability that P2 assigns to each of the remaining invariant strategies. For both treatments, the x-axis represents the probability that P2 assigns to the rational (a, d_*). This is also the probability that P2 assigns to the "heuristic best" strategy. For both treatments, the y-axis represents the probability that P2 assigns to the

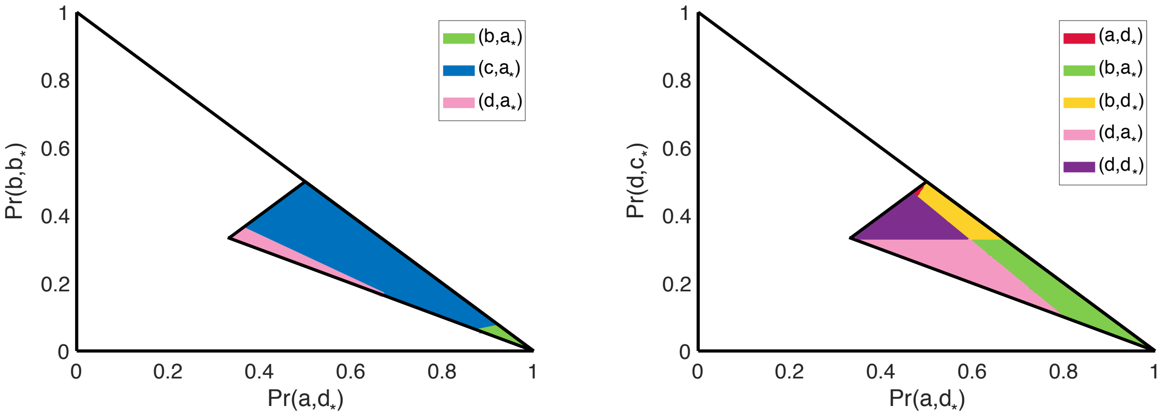

(a) Identification (IDENT): $\Pr(\mathrm{d},\mathrm{c}_*) = 0$

(b) Comparative Static (CS): $\Pr(\mathrm{b},\mathrm{a}_*) = 0$

Figure 5.1. P2's Beliefs and Best Responses

"heuristic second-best" strategy and the distance from the diagonal represents the probability that P2 assigns to the "heuristic third-best" strategy. So, the vertex $(1, 0)$ corresponds to P2 assigning probability 1 to the IU strategy; the vertex $(0, 1)$ corresponds to P2 assigning probability 1 to the "heuristic second-best;" and the vertex $(0, 0)$ corresponds to P2 assigning probability 1 to the "heuristic third-best." Importantly, the treatments differ in what is the "heuristic second-best" (resp. "heuristic third-best") strategy: The second-best (resp. third-best) is $(\mathrm{b}, \Pi_1(\mathrm{b})) = (\mathrm{b}, \mathrm{b}_*)$ (resp. $(\mathrm{c}, \Pi_1(\mathrm{c})) = (\mathrm{c}, \mathrm{a}_*)$) in IDENT and $(\mathrm{d}, \Pi_1(\mathrm{d})) = (\mathrm{d}, \mathrm{c}_*)$ (resp. $(\mathrm{c}, \Pi_1(\mathrm{c})) = (\mathrm{c}, \mathrm{b}_*)$) in CS.

The ordered-heuristic beliefs assumption adds the requirement that the probability assigned to the x-dimension must be weakly higher than the probability assigned to the y-dimension which, in turn, must be weakly higher than the probability assigned to the diagonal. So, the set of beliefs consistent with the ordered-heuristic beliefs assumption corresponds to the small triangles in Figure 5.1 (outlined in black). While these are each ordered-heuristic beliefs, they assign different probabilities to P1's strategies since P1's matrix has changed. This has implications for P2's best response: In each treatment, P2 may have an ordered-heuristic belief that assigns probability p (resp. q) to the heuristic-best strategy (resp. heuristic-second-best strategy) and, despite this, P2 may have a different best response across the treatments. This can be seen from the colored areas within the triangles, which represent P2's best responses. In particular, there are points in the triangle at which the best response differs across treatments. Moreover, the set of all best responses also differs across treatments: In IDENT, the strategies that are a unique best response under some ordered-heuristic belief are $(\mathrm{b}, \mathrm{a}_*)$, $(\mathrm{c}, \mathrm{a}_*)$, and $(\mathrm{d}, \mathrm{a}_*)$; in a sense, $(\mathrm{c}, \mathrm{a}_*)$ is the most prominent among these. In CS, the strategies that are a unique best response are $(\mathrm{a}, \mathrm{d}_*)$, $(\mathrm{b}, \mathrm{a}_*)$, $(\mathrm{b}, \mathrm{d}_*)$, $(\mathrm{d}, \mathrm{a}_*)$, and $(\mathrm{d}, \mathrm{d}_*)$.

The message remains analogous, even if P2 holds an ordered-heuristic belief that assigns positive probability to the "heuristic worst" outcome. Figure 5.2 in the Appendix varies the probability assigned to the "heuristic worst" outcome. [20] In each treatment, the set of non-constant strategies

---

[20] The probability that P2 assigns to the "heuristic third-best" strategy must be at least as high as the probability

that are a best response under some ordered-heuristic belief is the same as in Figure 5.1. (The only addition is that the constant $(\mathrm{d}, \mathrm{d}_*)$ can be a best response to an ordered-heuristic belief in IDENT.) So, put together, these figures depict HB_2.

The predictions accord with the observed distributions of P2 play, conditional on being identified as having a strategic bound of $k \geq 2$. Refer to Figure 4.1a. First, we observe that more subjects play $(\mathrm{c}, \mathrm{a}_*)$ in IDENT (28%) vs. CS (1%); the strategy $(\mathrm{c}, \mathrm{a}_*)$ is a best response (under a 2-heuristic belief) only in IDENT. Second, we observe that more subjects play the IU strategy, $(\mathrm{b}, \mathrm{a}_*)$, in CS (67%) vs. IDENT (53%); the IU strategy is a "more prominent" best response in CS. More generally, strategies that are played often are in HB_2. But, more notably, the strategies that are not in HB_2 are rarely played—this is the case despite the fact that about $3/4$ of the feasible strategies are not in HB_2. In IDENT (resp. CS), only 6% (resp. 1%) of the strategies played are not in HB_2.

Distribution of P3's and P4's Play Now turn to behavior in the role of Pj, where $j \geq 3$. To do so, consider a rational subject with a heuristic bound of $k \geq j$. Since Pj has a simple-strategic belief, any strategy that gets strictly positive probability must be a strategy that is either rational for $\mathrm{P}(j-1)$ or invariant for $\mathrm{P}(j-1)$. Since the subject has a heuristic bound of j, Pj also believes that the rational strategies of $\mathrm{P}(j-1)$ are contained in HB_{j-1}. As such, in the role of Pj, the subject must assign probability one to

$$\mathrm{HB}_{j-1} \cup \mathrm{Inv}_{j-1} = \mathrm{HB}_{j-1} \cup \{(\mathrm{a}, \mathrm{a}_*), (\mathrm{b}, \mathrm{b}_*), (\mathrm{c}, \mathrm{c}_*), (\mathrm{d}, \mathrm{d}_*)\}.$$

Now notice that, because HB_2 varies by treatment, the beliefs of such a P3 will vary by treatment. As such HB_3 varies by treatment. An analogous argument implies that HB_4 varies by treatment. The sets HB_3 and HB_4 can be read from Tables D.5-D.6.

The predictions accord with the observed distributions of play. Begin by focusing on subjects with an identified strategic bound of $k \geq 3$ and their behavior in the role of P3. There are three notable features of the distribution. First, we observe that more subjects play $(\mathrm{d}, \mathrm{b}_*)$ in IDENT vs. in CS. This corresponds to the fact that $(\mathrm{d}, \mathrm{b}_*)$ is in HB_3 for IDENT but not CS. Moreover, in IDENT, there is exactly one degenerate belief under which the strategy $(\mathrm{d}, \mathrm{b}_*)$ is a best response—a belief that assigns probability one to $(\mathrm{c}, \mathrm{a}_*)$. As pointed out above, $(\mathrm{c}, \mathrm{a}_*)$ is a best response under a 2-heuristic belief in IDENT—indeed, it is a best response that is often played in IDENT. Thus, it is intuitive that we would observe $(\mathrm{d}, \mathrm{b}_*)$ played in IDENT. Second, we observe that more subjects play the IU strategy $(\mathrm{c}, \mathrm{b}_*)$ in CS; as we will explain in the next subsection, this is natural given that, for P2, IU is a "more prominent" best response in CS. Third, strategies that are not in HB_3 are not played often. (Only 11% of the strategies played are not in HB_3.)

The picture is analogous for P4 behavior of subjects with an identified strategic bound of 4. First, we observe more subjects play $(\mathrm{d}, \mathrm{c}_*)$ in IDENT vs. in CS. Just like for P3, there is only one degenerate belief under which the strategy $(\mathrm{d}, \mathrm{c}_*)$ is a best response—a belief that assigns

---

assigned to the "heuristic worst." So, increasing the probability of the "heuristic worst" strategy shrinks the heuristic beliefs triangle and moves it to the interior of the simplex. For any ordered-heuristic belief, the probability of the "heuristic worst" strategy is at most .25.

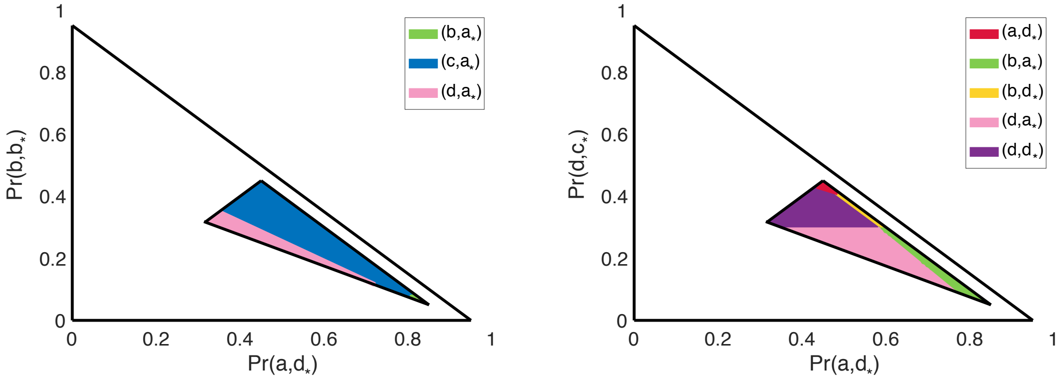

(a) IDENT: $\Pr(\mathrm{d}, \mathrm{c}_*) = .05$ (b) CS: $\Pr(\mathrm{d}, \mathrm{c}_*) = .05$

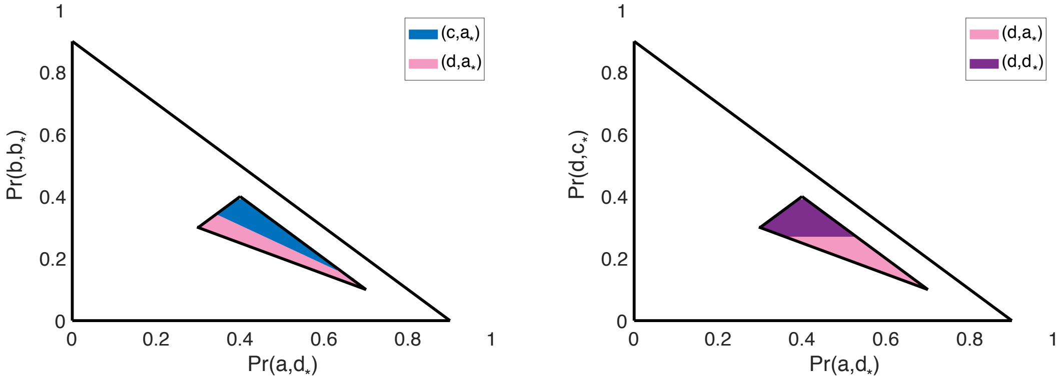

(c) IDENT: $\Pr(\mathrm{d}, \mathrm{c}_*) = .10$ (d) CS: $\Pr(\mathrm{d}, \mathrm{c}_*) = .10$

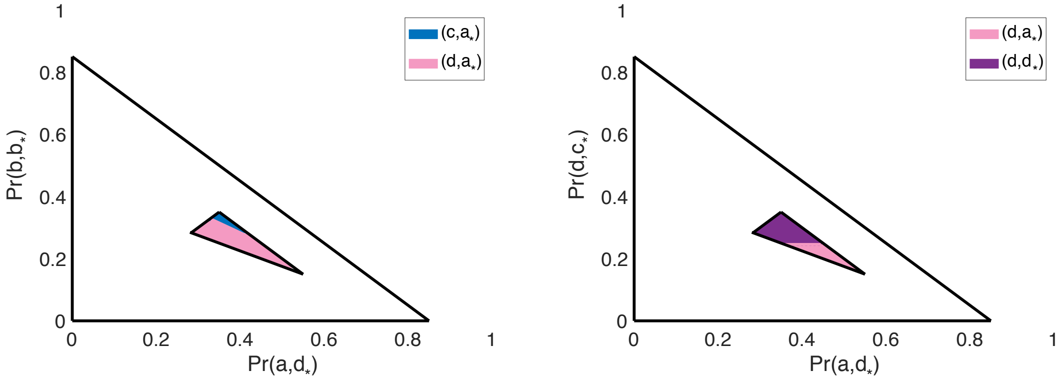

(e) IDENT: $\Pr(\mathrm{d}, \mathrm{c}_*) = .15$ (f) CS: $\Pr(\mathrm{d}, \mathrm{c}_*) = .15$

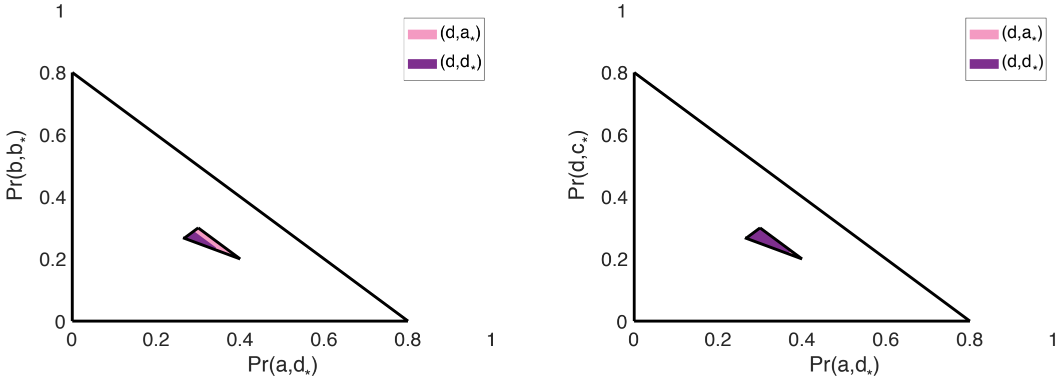

(g) IDENT: $\Pr(\mathrm{d}, \mathrm{c}_*) = .20$ (h) CS: $\Pr(\mathrm{d}, \mathrm{c}_*) = .20$

Figure 5.2. P2 Beliefs and Best Responses

	Rationality Bound			
Treatment	RB=1	RB=2	RB=3	RB=4
Identification (IDENT)	53%	31%	4%	11%
Comparative Static (CS)	46%	31%	1%	22%

Table 5.1. Rationality bounds ($p = .0391$)

probability one to $(\mathrm{d}, \mathrm{b}_*)$. This is a 3-heuristic belief in IDENT but not in CS. (The strategy $(\mathrm{d}, \mathrm{c}_*)$ is a best response under a 4-heuristic belief; however, any such belief must be non-degenerate assigning positive probability to both HB_3 and Inv_3.) Second, we observe more IU behavior in CS. Third, strategies that are not in HB_4 are not played often. (Only 5% of the strategies played are not in HB_4.)

5.4 Identified Rationality Bounds

Assume that subjects' *actual* rationality bounds do not vary across treatments. Under the HB model, we would expect IDENT to be associated with lower *identified* rationality bounds relative to CS. Section 1 (p. 12) pointed to the basic idea. We now expand.

Consider a rational subject whose heuristic bound is $k \geq 2$. First focus on behavior in the role of P2. Because P2 has an ordered-heuristic belief, the belief can be written as a vector (p, q, r), where p is the probability of the "heuristic best" strategy, q is the probability of the "heuristic second-best" strategy, and r is the probability of the "heuristic third-best" strategy. So, P2 assigns probability p to P1's dominant strategy. Certainly, if $p = 1$, P2's IU strategy is a best response to such a belief. However, it will also be a best response if p is less than 1 but "sufficiently high."

Suppose P2 forms his 2-heuristic belief systematically across treatments—that is, P2 always assigns the same probability $p : q : r$ to the heuristic best: second-best: third-best strategies. (Of course, what those strategies are change with the treatments.) Refer to Figures 5.1-5.2 and note that the IU strategy is a best response in IDENT, only if it is also a best response in CS. Thus, if P2's 2-heuristic belief is determined systematically across treatments, then P2 would play the IU strategy in IDENT only if P2 would also play the IU strategy in CS. (Even if P2 only has nearby beliefs, a similar conclusion would hold.) But the converse does not hold: For the IU strategy to be a best response in IDENT, P2 must assign $p \geq \frac{121}{158}$ to the heuristic-best strategy. However, in CS, the IU strategy *can* be a best response if P2 only assigns probability $p \geq \frac{80}{137}$ to the heuristic-best strategy.

The analysis of P3 and P4 yields analogous cross-treatment differences. The key is that a 3- or 4-heuristic P3 equates "P2 is rational" with "P2 plays a best response to a 2-heuristic belief." If such a P3 thinks that P2 forms his 2-heuristic belief systematically across treatments, the probability that P3 assigns to the event that "P2 plays the IU strategy" will be higher in CS vs. IDENT. Thus, P3 is more prone to play IU in CS vs. IDENT. And similarly for P4.

Indeed, this is what we observe in the data. Table 5.1 implies that the empirical CDF of identified rationality bounds in CS first-order stochastically dominates the empirical CDF of identified

	Strategic Bound			
Treatment	SB=1	SB=2	SB=3	SB=4
Identification (IDENT)	11%	23%	16%	50%
Comparative Static (CS)	19%	20%	12%	50%

(a) Strategic Bounds ($p = .2245$)

	Heuristic Bound			
Treatment	HB=1	HB=2	HB=3	HB=4
Identification (IDENT)	12%	23%	15%	50%
Comparative Static (CS)	21%	22%	11%	47%

(b) Heuristic Bounds ($p = 0.2558$)

Table 5.2. Empirical Distributions of Bounds by Treatment

rationality bounds in IDENT. IDENT is associated with a larger (resp. smaller) fraction of subjects with an identified rationality bound of 1 (resp. 4). A chi-square test of homogeneity allows us to reject the hypothesis that the *true* distributions of (identified) rationality bounds are the same. (The associated p-value is .0391.)

5.5 Identified Gap

While the identified rationality bounds vary by treatment, we cannot rule out the possibility that the true distributions of identified strategic bounds are the same across treatments. Refer to Table 5.2a, which depicts the empirical distribution of identified strategic bounds. (A chi-square test of homogeneity is associated with a p-value of .2245, meaning that we cannot reject the hypothesis of homogeneity.) Referring to Table 5.2b, a similar message pertains to the identified heuristic bounds.

This suggests that, if subjects hold heuristic beliefs, we would expect to see a larger identified gap in IDENT relative to CS. Indeed, that was the message of Table 4.2. In IDENT, 67% of observations were identified as having a gap, and in CS, 51% of observations were identified as having a gap. These results further suggest that the identified gap in CS may be an underestimate of the actual gap.

5.6 Other Ring Game Data

To ensure that our findings are robust, it is useful to look at the implications of the HB model in light of Kneeland's (2015) and Sprenger and Zhao's (2021) data. Both papers feature permuted ring games in which players choose one of three actions. In our discussion, we order the players so that, as in this paper, P1 has a dominant strategy; P2's payoffs depend only on P1's and P2's behavior; etc.

In Kneeland's data, 99% of rational subjects have behavior consistent with the SB model and 81% have behavior consistent with the HB model. Of the classified subjects, 38% are identified as

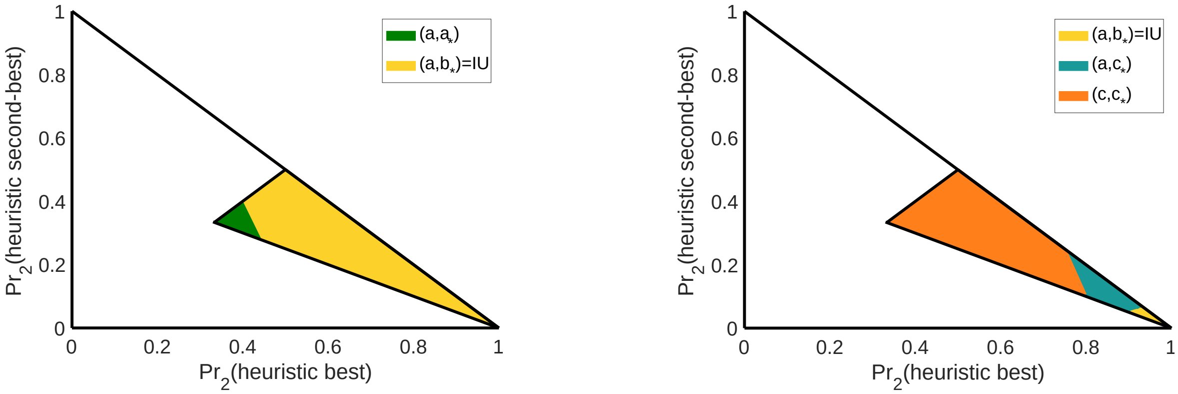

(a) Kneeland's (2015) Game

(b) Sprenger and Zhao's (2021) Modified Game

Figure 5.3. P2's Beliefs and Best Responses

having a gap between their strategic and rationality bounds. This number is notably lower than the cross-treatment 60% identified here (or the 67% identified in IDENT). However, viewed through the lens of the HB model, this is what we would expect.

To see this, refer to Figure 5.3a. It illustrates P2's best responses under any 2-heuristic belief, in Kneeland's games. Notice, the only non-constant behavior that is a best response to a 2-heuristic belief is P2's IU strategy. This implies that, under the HB model, we would also expect to see more IU behavior in the role of P3 and P4, even if subjects deviate from 3- and 4-rationality. In fact, in each player role Pi, the only non-constant strategy consistent with the HB model is the IU strategy for Pi. Therefore, we would expect subjects to have a higher identified rationality bound in Kneeland, even if their *actual* rationality bounds are the same as in our dataset. As such, we would expect a smaller identified gap in Kneeland.

The full dataset from Sprenger and Zhao is not publicly available, but we can still infer the behavior within a given player role. Their (within player role) data appears consistent with the HB model, in that at least 90% of subjects play strategies consistent with the HB model. As an example of how this consistency arises, refer to Figure 5.3b. The only non-constant non-IU strategy that is a best response to a 2-heuristic belief is $(\mathrm{a}, \mathrm{c}_*)$, and indeed 40% of subjects play $(\mathrm{a}, \mathrm{c}_*)$ in the role of P2.

It is worth noting that Sprenger and Zhao consider an alternate restricted model of strategic reasoning: the *focal beliefs model*. That model is related to but distinct from the HB model. Under the focal beliefs model, a player either assigns probability 1 to rationality or assigns probability 1 to the other player adopting a heuristic that is based on maximizing the sum (subject to a no-zero requirement). Interestingly, the focal beliefs model does not predict *any* non-constant non-IU play, something observed in the data. Moreover, the focal beliefs model does not predict our central comparative static (pp. 8-10).

Table 5.3 compares the rates of IU play in the role of P2, in Kneeland and Sprenger and Zhao. Notice, the rate is high (76%) in Kneeland and low (18%) in Sprenger and Zhao. When viewed through the lens of the HB model, this is exactly what we would expect. To see this, consider a rational subject whose heuristic bound is $k \geq 2$. Suppose she forms her 2-heuristic

Kneeland (2015)	76%
Comparative Static Treatment	57%
Identification Treatment	48%
Sprenger and Zhao (2021)	< 18%

Table 5.3. Percentage of Rational Subjects that Play IU in P2

belief systematically across treatments—that is, in the role of P2, she always assigns the same probability $p : q$ to the heuristic best : second-best strategies. Refer to Figure 5.3a and notice that the IU strategy is a best response in Sprenger and Zhao only if it is a best response in Kneeland. Thus, the subject plays the IU strategy in the role of P2 in Sprenger and Zhao only if she does the same in Kneeland.

We get a more complete picture by comparing this IU play in Kneeland and Sprenger and Zhao to that in IDENT and CS. Refer to Table 5.3: P2's IU play in Kneeland is higher than in CS and P2's IU play in Sprenger and Zhao is lower than in IDENT. Because these experiments vary in the number of actions, the comparison of behavior across experiments is more involved. Nonetheless, we can understand this difference in IU play in light of the HB model. For each treatment (or experiment) T, let $\underline{p}^T$ be the infimum of the set of $p \in [0, 1]$ so that the following holds: There is a 2-heuristic belief $\Pr_2$ where (i) P2 assigns probability p to P1's IU strategy, and (ii) in treatment T, P2's strict best response under $\Pr_2$ is his IU strategy. Thus, $\underline{p}^T$ is the minimum probability that P2 can assign to the heuristic best strategy of P1 (i.e., the rational strategy) and still have the IU strategy as a best response. The key is that

$$\underline{p}^{SZ} > \underline{p}^{\text{IDENT}} > \underline{p}^{\text{CS}} > \underline{p}^{K},$$

where SZ denotes the Sprenger and Zhao experiment and K denotes the Kneeland experiment. (To see this, refer to Figures 5.1, 5.2, and 5.3.) This suggests that, if subjects systematically determine the probability that they assign to the heuristic best (or rational) strategy, then we should see more IU behavior in Kneeland than in CS and more IU behavior in IDENT than in Sprenger and Zhao. Indeed, this is what we observe.

6 Discussion

6.A Level-k and Cognitive Hierarchy: The level-k model (Costa-Gomes, Crawford, and Broseta, 2001) and the cognitive hierarchy models (Camerer, Ho, and Chong, 2004) are often motivated by limitations in the players' ability to engage in interactive reasoning. 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 best-response bounds. Given the motivation, at first glance, it might seem as though that literature should have something to say about the question in the current paper. However, this is not the case.

This paper is concerned with departures from RCBR—a concept that is both conceptually and behaviorally distinct from the level-k and cognitive hierarchy models. Conceptually, level-k behavior is consistent with unlimited ability to engage in interactive reasoning. Moreover, the epistemic foundations for level-k behavior differ, in subtle ways, from m-rationality. See Brandenburger, Friedenberg, and Kneeland (2020) for an analysis that backs up these assertions. Further, the level-k and cognitive hierarchy models also allow for behavior that is consistent with RCBR. In particular, in some settings, the level-0 behavior is rationalizable. As a consequence, behavior that is viewed as level-k for some finite k will, in fact, be consistent with RCBR. Prominent examples include coordination games and the 11-20 games of Arad and Rubinstein (2012) and Alaoui and Penta (2016).[21]

Moreover, because of the literature's motivation, it has not been concerned with identifying ability bounds as distinct from best response bounds. Discussion 6.B points to some notable exceptions. However, as we will highlight there, the literature's measurement of ability is distinct from the concept here—i.e., the ability to engage in interactive reasoning.

6.B Other Potential Approaches: We argued that it is difficult to measure the ability to reason interactively without altering how players reason. That said, one might hope for another approach: one based on varying whether a subject plays against a more vs. less "sophisticated" subject pool—e.g., graduate vs. undergraduate, high vs. low Raven test score, computer vs. human, grandmaster vs. student, etc. (The technique was pioneered in Palacios-Huerta and Volij (2009), Agranov, Potamites, Schotter, and Tergiman (2012) and Georganas, Healy, and Weber (2015).) Perhaps if a subject plays differently when playing against a more vs. less sophisticated group, this should be taken as evidence that players engage in different levels of interactive reasoning across those two groups. As a consequence, in one case—say, the less sophisticated case—there is evidence that departures from RCBR are not driven by limits in ability.

This conclusion is premature—if not outright incorrect. Engaging in different levels of interactive reasoning is neither necessary nor sufficient for subjects to behave differently when playing against different populations. This is a familiar lesson from epistemic game theory: Subjects can have different beliefs about how different groups play the game, even if they engage in the same level of interactive reasoning about those groups. And, conversely, a subject may engage in different levels of interactive reasoning across groups, even if the subject has the same belief about play across the groups.[22] This is not only true, in principle, but also true for the experiments studied in the literature.

Recent work by Alaoui, Janezic, and Penta (2020) takes a different approach to varying subjects' level of "strategic sophistication." Importantly, their paper cannot address the question we pose:

---

[21] Arad and Rubinstein (2012) argue that the 11-20 game is so simple that, in the context of that game, there is unlimited ability to engage in interactive reasoning. Even if that were prima facie obvious, their paper would not allow us to conclude that the bound on ability is strictly higher than the maximum number m consistent with m-rationality—after all, in the 11-20 game, all behavior is consistent with RCBR.

[22] There is a second issue with this conclusion: It is not obvious that these notions of strategic sophistication are correlated with the ability to engage in interactive reasoning. Fe, Gill, and Prowse (2021) makes this point in a different context. In particular, they highlight that different measures of "sophistication" (in their context, measures of theory-of-mind vs. measures of cognitive ability) may be conceptually distinct and lead to different behavior. Thus, empirical work is needed to establish if and when they are "substitutable."

Is bounded reasoning about rationality driven by limited ability? To understand why, note that Alaoui, Janezic, and Penta's new approach keeps the subject pool fixed but varies the training a given subject receives in "how to solve" the game. That is, the treatment changes a subject's own training, but does not vary the training of the pool of players the subject plays against. (In both cases, the pool is untrained.) The idea is that a change in behavior post treatment suggests a change in the subject's ability to reason interactively. However, if the comparative static clearly points to a variation in the subject's ability to reason interactively, then it must be silent on whether there is variation in the subject's rationality bound. If the rationality bounds change in a way that matches a change in ability, then the subject's behavior would be consistent with no gap; if the rationality bounds did not vary, the subject's behavior would be consistent with a gap.[23]

6.C Strategic Bounds: Definition We pointed out that there are non-constant strategies of P2 that are not a best response under a 2-strategic belief and we do not classify associated observations as having (any) strategic bound. These observations are a best response under a belief of P2 which does not believe that "P1 is strategic;" but, they are not a best response under a belief of P2 that believes "P1 is not strategic."

More generally, if Pi believes that "P$(i-1)$ is not strategic" then Pi does not believe that "P$(i-1)$ is strategic." But the converse does not hold. So, if we were to adopt a definition of strategic bounds based, instead, on "not believing P$(i-1)$ is strategic," we would classify observations as we do, but also admit observations with higher strategic bounds. The approach we take is more conservative. (At the same time, there is only one subject whose behavior is consistent with a more lenient definition of strategic bounds.)

6.D Reasoning Across Games Table 5.1 pointed out that the identified rationality bounds differ across treatments. We argued that, under the HB model, this is natural: Even if players have the same rationality bounds across treatments, their identified rationality bounds may well differ across treatments.

This points to a notable feature of the level-k and cognitive hierarchy literatures: In those models, "levels of reasoning" is not a portable concept. (See, e.g., Georganas, Healy, and Weber, 2015.) That is, subjects who are characterized as being level-k reasoners in the context of one game may be characterized as being level-j reasoners in the context of another game, for some $j \neq k$. But, this behavior is also consistent with subjects whose identified rationality bounds change across games, even if though their actual rationality bounds do not. Thus, m-rationality may well be a

---

[23] Alaoui, Janezic, and Penta draw this conclusion because they think of ability as arising from a level-k model. It is worth noting that, for a minimally more permissive proxy of ability, it is not clear that subjects' ability bounds do vary across treatments. Because Alaoui, Janezic, and Penta focus on a variant of the 11-20 game, their experiments cannot rule out that the training itself may change the subjects' beliefs about how the game is played: This is so even if, both before and after the training, the subjects reason according to common belief of rationality (and so have unlimited ability to reason interactively). There is good reason to believe the training may change beliefs in this way. In their version of the 11-20 game, the entire strategy set is IU. (Specifically, 20 is a best response under a non-degenerate belief.) As such, under common belief of rationality, players may assign positive probability to any strategy. The training implicitly suggests that 20 is not played and, so, may change beliefs.

portable concept. This can be explored in future research, since the game parameters impact the relationship between the identified and actual rationality bounds.

6.E A Route to New Solution Concepts In the Introduction, we pointed out that, if bounded reasoning about rationality is not determined by limits in the ability to engage in interactive reasoning, then players may well engage in forms of strategic reasoning distinct from reasoning about rationality. The HB model highlights one such form of strategic reasoning. The experimental results suggest that certain rules-of-thumb may well shape players beliefs, in so far as they depart from reasoning about rationality. But, out of intellectual cautiousness, we stop short of offering the model as a novel solution concept.

There are many important questions about the nature of strategic reasoning that are simply not addressed by the HB model. For instance, we designed the experiments so that a class of heuristics coincide; the experiment cannot speak to whether a specific heuristic within that class drives heuristic reasoning. Likewise, we took no stand on how players reason across games that differ by more than a relabeling of actions. (For that reason, we took no stand on how players reason across player roles.) These—and, no doubt, other considerations—deserve further analyses. We see the results here as opening up a literature.

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., K. A. Janezic, and A. Penta (2020). Reasoning about others' reasoning. *Journal of Economic Theory 189*, 105091.

Alaoui, L. and A. Penta (2016). Endogenous depth of reasoning. *The Review of Economic Studies 83*(4), 1297–1333.

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., A. Friedenberg, and M. Siniscalchi (2021). Epistemic game theory: Reasoning about strategic uncertainty. unpublished.

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

Brandenburger, A. (2007). The power of paradox: some recent developments in interactive epistemology. *International Journal of Game Theory 35*(4), 465–492.

Brandenburger, A., A. Friedenberg, and T. Kneeland (2020). Two approaches to iterated reasoning in games. `http://u.arizona.edu/~afriedenberg/levelk.pdf`.

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, D. L., M. Schonger, and C. Wickens (2016). otree—an open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance 9*, 88–97.

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.

Dekel, E. and M. Siniscalchi (2014). Epistemic game theory. *Handbook of Game Theory 4*.

Fe, E., D. Gill, and V. Prowse (2021). Cognitive skills, strategic sophistication, and life outcomes. `https://web.ics.purdue.edu/~gill53/FeGillProwse_CogSkills.pdf`.

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

Greiner, B. (2015). Subject pool recruitment procedures: organizing experiments with orsee. *Journal of the Economic Science Association 1*(1), 114–125.

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.

Palacios-Huerta, I. and O. Volij (2009). Field centipedes. *American Economic Review 99*(4), 1619–35.

Sprenger, C. and W. Zhao (2021). Predictive accuracy for measures of higher order rationality - an experiment using ring game.

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.

For Online Publication Only

Appendix A The Identification Strategy

It will be convenient to introduce notation that will be used throughout the appendices. Given a (compact metric) set Y, write $\Delta(Y)$ for the set of probability measures on Y.

Denote the set of strategies by $S = \{\mathrm{a}, \mathrm{b}, \mathrm{c}, \mathrm{d}\} \times \{\mathrm{a}_*, \mathrm{b}_*, \mathrm{c}_*, \mathrm{d}_*\}$ and the set of observations by $X = S \times S \times S \times S$. As in Section 5 of the main text, we will write

$$\mathrm{Inv}_i = \{(\mathrm{a}, \Pi_i(\mathrm{a})), (\mathrm{b}, \Pi_i(\mathrm{b})), (\mathrm{c}, \Pi_i(\mathrm{c})), (\mathrm{d}, \Pi_i(\mathrm{d}))\}$$

Write $\mathcal{SB}^k$ for the **set of observations with an identified strategic bound of** k. So, when $k \geq 2$, an observation $x = (x(1), x(2), x(3), x(4)) \in \mathcal{SB}^k$ if and only if (i) $x(1) = (\mathrm{a}, \mathrm{d}_*)$; (ii) $x(2)$ is a unique best response under a 2-strategic belief; and (iii) for each $i \in \{j : 4 \geq j > k\}$, $x(i)$ is constant.

A.1 Identification of the Strategic Bounds

In this appendix, we draw a connection between behavior of a rational subject that has a strategic bound of k and an observation $x \in \mathcal{SB}^k$. In doing so, we equate a *rational* subject with a subject that plays a unique best response given her first-order belief (i.e., belief about observations). (This fits with the earlier assumption that a rational subject is not indifferent between any two strategies.) We will show that an observation $x \in \mathcal{SB}^k$ if and only if (i) a rational subject with a strategic bound of k would play x, and (ii) a rational subject with a strategic bound of $j < k$ would not play x.

The result is immediate for $k = 1$. So we focus on showing the result for $k \geq 2$. In particular, we will show the following:

Proposition A.1. *For each $k \geq 2$ the following hold:*

(i) If a rational subject with a strategic bound of k chooses an observation x, then $x \in \mathcal{SB}^j$ for some $j \leq k$.

(ii) If $x \in \mathcal{SB}^k$, then there exists a rational subject with a strategic bound of k that chooses x.

Proposition A.1 serves to characterize the behavior of a rational subject who has a strategic bound of k. Part (i) argues that a necessary condition is that the behavior has an identified strategic bound of $j \leq k$. Part (ii) gives sufficiency—an observation with an identified strategic bound of k must be consistent with the behavior of a rational subject who has a strategic bound of k.

Proof of Necessity Proposition A.1(i) will follow immediately from the following Lemmata:

Lemma A.1. *Consider a rational subject who has a strategic bound of $k \geq 2$. In the role of P2, the subject plays a strategy that is a unique best response to a 2-strategic belief.*

Lemma A.2. *Fix some j, k with $4 \geq j > k \geq 2$. Consider a rational subject who has a strategic bound of k. In the role of Pj, the subject plays a constant strategy.*

To show these results, it will be useful to, in turn, have two auxiliary lemmata.

Lemma A.3. *Let Pi = P1, P3, P4. If Pi is strategic and believes that P$(i-1)$ has a strategic bound of 1, then Pi has a constant belief.*

Proof. Suppose 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." By the Principle of Strategic Reasoning, Pi's belief satisfies $\mathrm{Pr}_i(\mathrm{d}, \mathrm{e}_*) > 0$ only if $\Pi_{(i-1)}(\mathrm{d}) = \mathrm{e}_*$. Since $i - 1 \neq 1$, Pi has a constant belief. ☐

Lemma A.4. *If P1 is strategic and believes that P4 has a strategic bound of 2, then P1 has a constant belief.*

Proof. Suppose P1 believes that P4 has a strategic bound of 2. Then, P1 believes that "P4 is 2-strategic and believes that P3 has a strategic bound of 1." By Lemma A.3, P1 believes that "P4 is 1-strategic and has a constant belief." Applying the Principle of Strategic Reasoning, P1 believes that P4 plays a constant strategy. ☐

Proof of Lemma A.1. Since the subject has a strategic bound of $k \geq 2$, in the role of P2, the subject satisfies the Principle of Strategic Reasoning. It suffices to show that P2 believes P1 has a constant belief: If so, the Principle of Strategic Reasoning implies that P2's belief assigns probability one to

$$\mathrm{Inv}_1 = \{(\mathrm{a}, \Pi_1(\mathrm{a})), (\mathrm{b}, \Pi_1(\mathrm{b})), (\mathrm{c}, \Pi_1(\mathrm{c})), (\mathrm{d}, \Pi_1(\mathrm{d}))\},$$

i.e., P2's belief is a 2-strategic belief.

For $k = 2$, this follows immediately from the fact that the subject believes that the player in the role of P1 satisfies the Principle of Non-Strategic Reasoning. For $k = 3$ (resp. $k = 4$), P2 believes that P1 has a strategic bound of 2 (resp. 3). As such, P2 believes that "P1 believes P4 has a strategic bound of 1 (resp. 2)." So, by Lemma A.3 (resp. Lemma A.4), P2 believes that P1 has a constant belief. ☐

Proof of Lemma A.2. Since the subject has a strategic bound of $k \geq 2$, the subject satisfies the Principle of Strategic Reasoning. It suffices to show that, if $j > k$, Pj believes P$(j-1)$ has a constant belief: If so, the Principle of Strategic Reasoning implies that Pj's belief assigns probability one to

$$\mathrm{Inv}_{j-1} = \{(\mathrm{a}, \Pi_{j-1}(\mathrm{a})), (\mathrm{b}, \Pi_{j-1}(\mathrm{b})), (\mathrm{c}, \Pi_{j-1}(\mathrm{c})), (\mathrm{d}, \Pi_{j-1}(\mathrm{d}))\}.$$

Note, Pj's belief is then a constant belief, as $j - 1$ is either 2 or 3. Since a rational subject has a unique best response to her belief, the subject plays a constant strategy.

Since $4 \geq j > k$, k is either 2 or 3. If $k = 2$, the claim follows immediately from the fact that the subject believes the player in the role of P$(j-1)$ satisfies the Principle of Non-Strategic Reasoning. If $k = 3$, Pj believes that P$(j-1)$ has a strategic bound of 2. As such, Pj believes that "P$(j-1)$ believes that P$(j-2)$ has a strategic bound of 1 (resp. 2)." So, by Lemma A.3, Pj believes that P$(j-1)$ has a constant belief. ☐

Proof of Sufficiency To show Proposition A.1(ii), we need to lay out certain epistemic formalisms. We use the minimal formalism necessary, introducing first-, second-, and higher-order beliefs, as they are needed. We don't redefine our concepts within those formalisms; we will be careful with words in a way that we think is clear.

Two factors are important in how we define these higher-order beliefs. First, we think of a subject as having hierarchies of beliefs about the observed behavior of a *single* other subject. This is for notational simplicity. Second, beliefs are defined over observations (instead of specifying beliefs in each player role about strategies played). This is because we will want to think about a subject with a heuristic bound of 2 believing that the other subject has the same belief independent of the player role.

It is be convenient to define **first-order belief maps** for observations with an identified strategic bound of $k \geq 2$. Note, a first-order belief is a belief on $X^1 = X$.

Definition A.1. *Fix $k \geq 2$. For each i, define a map $f_i^k : \mathcal{SB}^k \to \Delta(S)$ so that, for each $x = (x(1), x(2), x(3), x(4)) \in \mathcal{SB}^k$, the following hold:*

- *If $i = 2$, $f_i^k(x)$ is a 2-strategic belief under which $x(2)$ is a unique best response.*
- *If $i \in \{1\} \cup \{j : j > k\}$, $f_i^k(x)$ is a constant belief under which $x(i)$ is a unique best response.*
- *If $i \in \{j : k \geq j > 2\}$, $f_i^k(x)$ is a belief under which $x(i)$ is a unique best response.*

The first-order belief map *for k is a map $f^k : \mathcal{SB}^k \to \Delta(X)$ where, for each observation $x = (x(1), x(2), x(3), x(4)) \in \mathcal{SB}^k$, $f^k(x) = f_2^k(x) \otimes f_3^k(x) \otimes f_4^k(x) \otimes f_1^k(x)$.*[24]

A second-order belief is a belief in $X^2 = X^1 \times \Delta(X^1)$. To simplify notation, we think of a third-order belief as a belief about $X^3 = X^1 \times \Delta(X^2)$—i.e., specifying a belief about observations and other subjects' beliefs about "observations and first-order beliefs." (We can read off the second-order belief via marginalization.) And, similarly, we think of a fourth-order belief as a belief about $X^4 = X^1 \times \Delta(X^3)$.

It will suffice to focus on degenerate higher-order beliefs. Thus, given a belief $\nu \in \Delta(X^k)$, we will write δ_ν to indicate the belief in $\Delta(\Delta(X^k))$ that assigns probability one to ν. To construct the relevant degenerate beliefs, it is convenient to fix two objects: First, let $\Pr \in \Delta(S)$ be a constant belief and let $\overrightarrow{\Pr} \in \Delta(X)$ be the associated product measure $(\Pr \otimes \Pr \otimes \Pr \otimes \Pr)$. Second, for each $k \in \{2, 3\}$, fix $y^k \in \mathcal{SB}^k$.

Now, inductively define maps $h^k : \mathcal{SB}^k \to \Delta(X^k)$.

- For each $x \in \mathcal{SB}^2$, $h^2(x) = f^2(x) \otimes \delta_{\overrightarrow{\Pr}}$.
- Assume h^k has been defined for $k \in \{2, 3\}$. For each $x \in \mathcal{SB}^{k+1}$, let $h^{k+1}(x) = f^{k+1}(x) \otimes \delta_{h^k(y^k)}$.

The following Lemma establishes Proposition A.1(ii).

---

[24] We adopt this ordering so that f^k is indeed a distribution on observations $(y(1), y(2), y(3), y(4))$.

Lemma A.5. *Let $k \geq 2$. For each observation $x \in \mathcal{SB}^k$,*

- *(i) x is a unique best response under $f^k(x)$, and*
- *(ii) $(x, h^k(x))$ has a strategic bound of k.*

Proof. By construction, each $x \in \bigcup_{k \geq 2} \mathcal{SB}^k$ is a unique best response under $f^k(x)$. So, we focus on showing that each $(x, h^k(x))$ has a strategic bound of k.

$k = 2$: Fix $x = (x(1), x(2), x(3), x(4)) \in \mathcal{SB}^2$. First observe that, for each i with $f_i^2(x)$ constant, $x(i) \in \text{Inv}_i$. Thus, $(x, h^2(x))$ is 1-strategic. Moreover, by construction, each $f_i^2(x)$ assigns probability one to Inv_{i-1}. Thus, $(x, h^2(x))$ is 2-strategic. Moreover, since $h^2(x)$ assigns probability one to the other player having the same constant belief Pr across player roles, $(x, h^2(x))$ also satisfies the Principle of Non-Strategic Reasoning. Thus, $(x, h^2(x))$ has a heuristic bound of 2.

$k \geq 3$: Assume the result holds for $k \in \{2, 3\}$. Fix some $x = (x(1), x(2), x(3), x(4)) \in \mathcal{SB}^{k+1}$. First observe that, for each i with $f_i^{k+1}(x)$ constant, $x(i) \in \text{Inv}_i$. Moreover, $h^{k+1}(x)$ assigns probability one to some k^{th}-order belief $h^k(y^k)$, where $y \in \mathcal{SB}^k$. So, by the induction hypothesis, $(x, h^{k+1}(x))$ has a heuristic bound of $k + 1$. ☐

A.2 Identifying the Gap

Lemma A.6. *If an observation has an identified strategic bound of k, then the identified rationality bound must be $m \leq k$.*

Proof. Fix an observation $(x(1), x(2), x(3), x(4)) \in \mathcal{SB}^k$. Suppose, contra hypothesis, that the observation has an identified rationality bound of $m > k$. Then, $x(k + 1)$ is the non-constant IU strategy, contradicting $(x(1), x(2), x(3), x(4)) \in \mathcal{SB}^k$. ☐

Lemma A.7. *Fix a rational subject who has a strategic bound of k.*

- *(i) If the subject's behavior generates an observation with an identified rationality bound of m, then $m \leq k$.*
- *(ii) If the subject's behavior generates an observation with an identified gap, then the subject's rationality bound must be strictly lower than k.*

Proof. Since the subject is rational and has a strategic bound of k, the subject's behavior generates an observation with an identified strategic bound of $j \leq k$. (See Proposition A.1(i).) By Lemma A.6, the identified rationality bound must be $m \leq j$. Thus, $m \leq k$, establishing part (i).

Now observe that, if the subject's behavior generates an identified gap, then the subject's identified rationality bound must be $m < 4$. As a consequence, the subject's rationality bound is some $n \leq m$. Put together, $n \leq m < j \leq k$, establishing part (ii). ☐

Appendix B Heuristic Beliefs Model

B.1 Heuristic Beliefs vs. Strategic Beliefs

Lemma B.1. *If a subject is k-heuristic, the subject is k-strategic.*

Proof. For $k = 1$, the claim is by definition. Fix a subject that is 2-heuristic. Then the subject is 1-strategic and, in each player role Pi, the subject has simple strategic belief Pr_i. So, if $\mathrm{Pr}_i(\mathrm{e}, \mathrm{f}_*) > 0$, then $(\mathrm{e}, \mathrm{f}_*)$ is either rational or invariant. As such, to show that the subject satisfies the Principle of Strategic Reasoning, it suffices to show that, if Pi believes that P$(i-1)$ has a constant belief and $\mathrm{Pr}_i(\mathrm{e}, \mathrm{f}_*) > 0$ for a rational $(\mathrm{e}, \mathrm{f}_*)$, then $(\mathrm{e}, \mathrm{f}_*)$ is invariant. This follows from the fact that Pi believes that, if P$(i-1)$ is rational, then he is not indifferent between any two strategies.

Now suppose the claim holds for $k \in \{2, 3\}$. Fix a $(k+1)$-heuristic subject. By the induction hypothesis, the subject is k-strategic and believes that others are k-strategic. Thus, the subject is $(k+1)$-strategic. □

Lemma B.2. *If a subject has a heuristic bound of k, the subject has a strategic bound of k.*

Proof. For $k = 1$, the claim is by definition. Assume the claim holds for $k \geq 1$, where $3 \geq k$. Fix a subject with a heuristic bound of $(k+1)$. By Lemma B.1, the subject is $(k+1)$-strategic. Moreover, by the induction hypothesis, the subject believes that the others have a strategic bound of k. Thus, the subject also has a strategic bound of $(k+1)$. □

Lemma B.3. *If an observation is identified as having a heuristic bound of k, it is identified as having a strategic bound of k.*

Proof. An observation is identified as having a heuristic bound of 1 if and only if it is identified as having a strategic bound of 1. As a consequence, if an observation is identified as having a heuristic bound of 2, it is also identified as having a strategic bound of 2. If it is identified as having a strategic bound of 2 but not identified as having a heuristic bound of 2, it must be that the behavior in the role of P2 is not a unique best response under a 2-heuristic belief. As a consequence, the observation won't be identified as having a heuristic bound of 3 or 4. So, if the observation is identified as having a heuristic bound of 3, it must be identified as having a strategic bound of 3. Repeating the argument gives the conclusion for $k = 3$ and $k = 4$. □

B.2 Identification of Heuristic Bounds

Let $\mathcal{HB}^k$ be **the set of observations with an identified heuristic bound of** k. So, $x \in \mathcal{HB}^1$ if and only if $x \in \mathcal{SB}^1$. Moreover, for $k \geq 2$, $\mathcal{HB}^k = \prod_{i \leq k} \mathrm{HB}_i \times \prod_{i > k} \mathrm{Inv}_i$.

To define the non-indifference condition, it is convenient to introduce the following definition: Say Pi is 1-**rational** if she plays a best response given a belief about P$(i-1)$'s play of the game.[25]

---

[25] Notice, here we do not include a requirement that the best response is unique.

Say Pi is m**-rational** if she plays a best response given a belief that assigns probability 1 to the event that "P$(i-1)$ is $(m-1)$-rational."

Say a subject satisfies **first-order belief of non-indifference** if, in each player role Pi, she believes:

> "If P$(i-1)$ is 1-rational, he is not indifferent between any two strategies."

Inductively, a subject satisfies k^{th}**-order belief of non-indifference** if, in each player role Pi, she believes:

> "If P$(i-1)$ is $(k-1)$-rational, he satisfies $(k-1)^{th}$-order belief of non-indifference."

In what follows, we assume that, if a subject is k-heuristic, the subjects satisfies $(k-1)^{\text{th}}$-order belief of non-indifference. If this holds for all $k \in \{1, 2, 3, 4\}$, we say the **belief of non-indifference assumption** holds.

Proposition B.1. *Fix $k \geq 2$.*

(i) Suppose that the belief of non-indifference assumption holds. If a rational subject with a heuristic bound of k chooses an observation x, then $x \in \mathcal{HB}^j$ for some $j \leq k$.

(ii) If $x \in \mathcal{HB}^k$, then there exists a rational subject with a heuristic bound of k that both satisfies belief of non-indifference and chooses x.

Proposition B.1 characterizes the behavior of a rational subject with a heuristic bound of k. Part (i) says that a necessary condition is that the subject's behavior has an identified heuristic bound of $j \leq k$. Part (ii) establishes the converse. Any observation with an identified heuristic bound of k can be generated by the behavior of a rational subject with a heuristic bound of k.

Proof of Necessity To prove part (i), we make use of the following Lemma.

Lemma B.4. *Suppose that the belief of non-indifference assumption holds. Fix $k \geq j \geq 2$. If a rational subject is k-heuristic then, in the role of Pj, the subject plays some strategy $x(j) \in \mathrm{HB}_j$.*

Proof. The proof is by induction on j. For $j = 2$, the proof is immediate from the fact that k-heuristic implies 2-heuristic and no subject is indifferent between any two actions. So suppose the claim holds for $j \in \{2, 3\}$ and let $k \geq (j+1)$. Note, a k-heuristic subject satisfies j^{th}-order belief of non-indifference and believes that others are j-heuristic. So, by the induction hypothesis, a k-heuristic subject believes that the rational strategies, in the role of Pj, are the strategies in HB_j. Since Inv_j is the set of constant strategies, a k-heuristic subject has a $(j+1)$-heuristic belief in the role of P$(j+1)$. Since no subject is indifferent between any two actions, a k-heuristic subject plays some strategy $x(j+1)$ that is a unique best response to a $(j+1)$-heuristic belief. ☐

Proof of Proposition B.1(i). Immediate from Lemma B.4, Lemma B.2, and Lemma A.2. ☐

Proof of Sufficiency To show Proposition B.1(ii), we need, once again, to introduce hierarchies of beliefs. We follow the approach in Appendix A.1, making suitable modifications. We often use the same notation for adjacent concepts. We do this to draw parallels. No confusion should result.

Definition B.1. *Fix $k \geq 2$. For each i, define a map $f_i^k : \mathcal{HB}^k \to \Delta(S)$ so that, for each $x = (x(1), x(2), x(3), x(4)) \in \mathcal{HB}^k$, the following hold:*

- *If $i \in \{2, \ldots, k\}$, $f_i^k(x)$ is a i-heuristic belief under which $x(i)$ is a unique best response.*
- *If $i \in \{1\} \cup \{j > k\}$, $f_i^k(x)$ is constant ordered-heuristic belief under which $x(i)$ is a unique best response.*

The ***first-order belief map for*** *k is a map $f^k : \mathcal{HB}^k \to \Delta(X)$ where, for each observation $x = (x(1), x(2), x(3), x(4)) \in \mathcal{HB}^k$, $f^k(x) = f_2^k(x) \otimes f_3^k(x) \otimes f_4^k(x) \otimes f_1^k(x)$.*

Note, for each k and each $x \in \mathcal{HB}^k$, $f^k(x)$ is a distribution on X. For each k, we can provide a uniform bound on the supports of distributions $f^k(x)$. In particular, set

$$\mathrm{BSupp}^k = \begin{cases} \mathrm{Inv}_1 \times \mathrm{Inv}_2 \times \mathrm{Inv}_3 \times \mathrm{Inv}_4 & \text{if } k = 2 \\ \mathrm{Inv}_1 \times (\mathrm{HB}_2 \cup \mathrm{Inv}_2) \times \mathrm{Inv}_3 \times \mathrm{Inv}_4 & \text{if } k = 3 \\ \mathrm{Inv}_1 \times (\mathrm{HB}_2 \cup \mathrm{Inv}_2) \times (\mathrm{HB}_3 \cup \mathrm{Inv}_3) \times \mathrm{Inv}_4 & \text{if } k = 4. \end{cases}$$

For each $k \geq 2$ and each $x \in \mathcal{HB}^k$, $\mathrm{Supp} f^k(x) \subseteq \mathrm{BSupp}^k$. Moreover, for each $k \geq 2$, $\mathcal{HB}^{k-1} \subseteq \mathrm{BSupp}^k$.

With this in mind, for each $k \geq 2$, construct a mapping $g^k : \mathrm{BSupp}^k \to \mathcal{HB}^{k-1}$ that satisfies the following two propreties: First, for each $x \in \mathcal{HB}^{k-1} \subseteq \mathrm{BSupp}^k$, $g^k(x) = x$. Second, for each $x = (x(1), x(2), x(3), x(4)) \in \mathrm{BSupp}^k \backslash \mathcal{HB}^{k-1}$, $g^k(x) = (y(1), y(2), y(3), y(4))$ where $y(i) = x(i)$ if $x(i) \in \mathrm{HB}_i$. (Observe that this can be done, since $\mathcal{HB}^{k-1} = \prod_{i<k} \mathrm{HB}_i \times \prod_{i \geq k} \mathrm{Inv}_i$.)

We also fix a constant belief $\Pr \in \Delta(S)$ under which each Pi has a unique best response. Let $\overrightarrow{\Pr} \in \Delta(X)$ be the associated product measure $(\Pr \otimes \Pr \otimes \Pr \otimes \Pr)$.

Now, inductively define maps $h^k : \mathcal{HB}^k \to \Delta(X^k)$ as follows. First, for each $x \in \mathcal{HB}^2$, $h^2(x) = f^2(x) \otimes \delta_{\overrightarrow{\Pr}}$. Assume that $h^k : \mathcal{HB}^k \to \Delta(X^k)$ has been defined for $k \in \{2, 3\}$. For each $x \in \mathcal{HB}^{k+1}$, let $h^{k+1}(x) = \mu \in \Delta(X \times \Delta(X^k))$ that satisfies the following:

$$\mu(y, \nu) = \begin{cases} f^{k+1}(x)(y) \times \delta_\nu & \text{if } y \in \mathrm{Supp} f^{k+1}(x) \text{ and } \nu = h^k(g^{k+1}(y)) \\ 0 & \text{otherwise.} \end{cases}$$

Note that $f^{k+1}(x) \in \Delta(X \times \Delta(X^k))$.

Lemma B.5. *Fix $k \geq 2$. For each $x \in \mathcal{HB}^k$, $(x, h^k(x))$ is rational and has a heuristic bound of k.*

Proof. Begin by fixing $x = (x(1), x(2), x(3), x(4)) \in \mathcal{HB}^2$. First observe that each $x(i)$ is a unique best response to the invariant belief $f_i^2(x)$. Thus, $(x, h^2(x))$ is rational and 1-heuristic. Moreover, each $f_i^2(x)$ is an ordered-heuristic belief. And, since $h^2(x)$ assigns probability one to the other playing Pr independent of the player role, $h^2(x)$ believes that, if the other subject plays a best response, that best response is unique. This establishes that $(x, h^2(x))$ is 2-heuristic. Finally, since $h^2(x)$ assigns probability one to the other player having the same constant belief Pr across player roles, $(x, h^2(x))$ also satisfies the Principle of Non-Strategic Reasoning. Thus, $(x, h^2(x))$ has a heuristic bound of 2.

Now assume the claim holds for $k \in \{2, 3\}$ and fix some $x \in \mathcal{HB}^{k+1}$. Note that x is a unique best response under $f^{k+1}(x)$. So, $(x, h^{k+1}(x))$ is rational. With this, we focus on showing that $(x, h^{k+1}(x))$ has a heuristic bound of k.

Begin by observing that, if $f_i^{k+1}(x)$ is a constant belief, $x(i)$ must be invariant. (This uses the fact that x is a unique best response under $f^{k+1}(x)$.) So, $(x, h^{k+1}(x))$ is 1-heuristic.

Now, let $\mu \equiv h^{k+1}(x)$. It suffices to show that, if $\mu(y, \nu) > 0$ then (y, ν) has a heuristic bound of k. Since $\mu(y, \nu) > 0$, $y = (y(1), y(2), y(3), y(4)) \in \text{Supp} f^{k+1}(x)$ and $\nu = h^k(g^{k+1}(y))$. If $y \in \mathcal{HB}^k$, the result follows from the induction hypothesis. If $y \notin \mathcal{HB}^k$, then there exists $z = (z(1), z(2), z(3), z(4))$ so that $g^k(y) = z$. Note, $(z, h^k(z))$ is rational and has a heuristic bound of k. (This follows from the induction hypothesis.) We use this fact to show that $(y, h^k(z))$ has a heuristic bound of k.

First observe that $(y, h^k(z))$ must be 1-heuristic. To see this, note that $y(1), y(3), y(4)$ are invariant and $y(2) = z(2)$. So, using the fact that $(z, h^k(z))$ is 1-heuristic, $(y, h^k(z))$ must also be 1-heuristic. From this and the fact that $(z, h^k(z))$ has a heuristic bound of k, it follows that $(y, h^k(z))$ also has a heuristic bound of k. ☐

Let $S(i)$ be the set of strategies in the role of Pi. A basic induction argument establishes the following:

Lemma B.6. *Fix $k \geq 2$ and $x \in \mathcal{HB}^k$. If $h^k(x)$ assigns strictly positive probability to $(y, \nu) \in X \times \Delta(X^{k-1})$ and (y, ν) is $(k-1)$-rational in the role of Pi, then $y(i)$ is a strict best response under $f_i^{k-1}(y) = \text{marg}_{S(i)}\nu$.*

Proof of Proposition B.1(ii) . Immediate from Lemmata B.5 and B.6. ☐

Appendix C Experimental Instructions

Zoom Meeting: Set-Up A recurring Zoom meeting was created with the following features. A waiting room was enabled and participants could not join before host. Participants' video was set to 'on,' but participants were muted upon entry. The Zoom chat feature was set to 'chat with host only.' In addition, participants could not rename themselves or provide meeting reactions and non-verbal feedback. Experimenter names were anonymized to "Main Experimenter" and "Experimental Assistant X" (where X took on one of several letters).

Zoom Meeting: Check-in Experimental Assistants allowed subjects into the Zoom meeting one at a time. They ensured that subjects had their video on and checked the subjects' names. Then, they changed the subject's name to a pre-specified number and put the subject back into the waiting room to await the start of the experiment. (A Powerpoint slide explained this to the subjects.) Messages went out regularly to subjects in the waiting room, to ensure that they understood that the experiment would soon start and that they would be required to have their video on throughout.

Zoom Meeting: Start of the Experiment After check-in was completed, subjects were allowed into the Zoom session. At that time, the Main Experimenter read the preliminary instructions described below. During that time, Experimental Assistants distributed links to the experiment. When all links were distributed, the experiment began. At this time, the Zoom meeting was set to "subjects cannot unmute themselves."

Preliminary Instructions The following is the text for the preliminary instructions read at the start of the session.

- Can you hear me? Please nod if you can. Thank you. If you cannot, please check your audio and volume.
- Thank you for agreeing to participate in today's ELFE experiment.
- You will soon receive a link to the experiment in the Zoom chat. It is important that you do not click on the link until I tell you to do so: If you click too early, it will likely take longer for everyone involved.
- Let me begin by telling you some important features about the experiment.
- It is important that you give the experiment your full attention. For the duration of the experiment, please do not use your mobile phone or engage in other activities. In addition, please keep your video camera on at all times and remain visible at your computer. If your video camera is not on, we will need to remove you from today's session.
- The experiment will begin with a consent form. Please read it carefully and be sure that you understand what you are consenting to. After that, you will receive a series of instructions for the experiment followed by a quiz that is intended to ensure that you have understood the instructions. It is important that you read the instructions carefully so that you can complete the quiz. If you answer the quiz correctly on the first try, you will earn an additional quiz fee described in the experiment. But, even if you don't, it is important that you answer the quiz correctly within three tries. If you do, you will have the opportunity to earn considerably more money during the experiment. The instructions will explain this further. Punchline: You have every incentive to read the instructions carefully.
- Once you are admitted to the experiment, you will have 75 minutes to complete it. You will be able to advance through the experiment at your own pace. On each screen, there will

be a blue chat icon in the lower left-hand corner. If you have any questions throughout the experiment, you can use that to chat with one of the experimenters.

- After everyone has completed the experiment, your payment information will appear on the screen. At that point, I will make an announcement on Zoom that the experiment is over and you are free to leave the Zoom meeting. If you leave the Zoom meeting before the experiment is over, you may not be paid for your participation in the experiment.
- While you will be able to advance through the experiment at your own pace, the experiment will only "end" after everyone in this session has completed the experiment. In order to receive payment, you will need to wait until the end, after everyone has completed the session, and the Zoom is dismissed. Expect this to take the full session length. Punchline: There is no incentive to finish the experiment quickly. Expect to be here the full session length. Feel free to carefully study the questions, so that your earnings are high.
- Before we begin, are there any questions? If so, please type it into the Zoom chat box and send it to myself—the Main Experimenter.
- Please wait quietly until I tell you to start the experiment. Again, please do not click the link until I tell you to do so. Thank you.

Screenshot of Instructions Below are screenshots of the instruction. Note, the instructions point to a "highlighting feature" that the subjects can use throughout the game to help them keep track of other players' payoffs. The instructions require the subject to use the highlighting feature before they can move on to a subsequent screen. (For the purposes of producing the screenshots on paper, they are scaled down. Subjects saw larger fonts.)

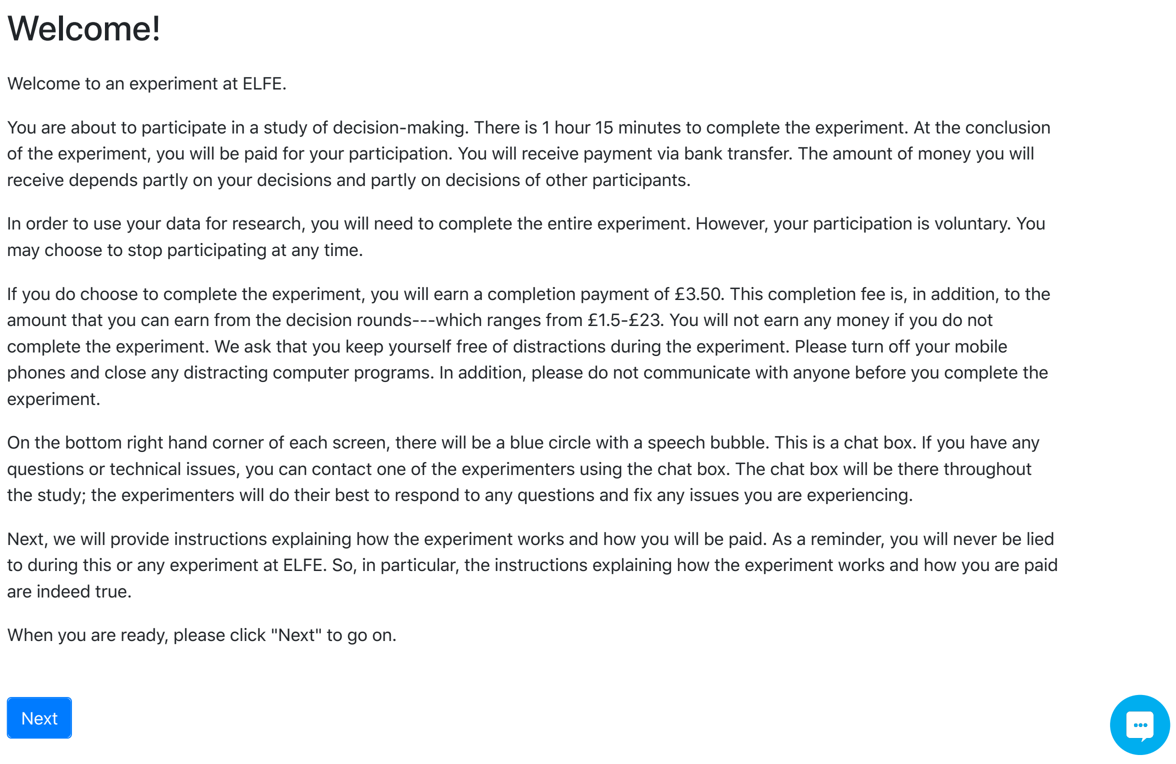

Welcome!

Welcome to an experiment at ELFE.

You are about to participate in a study of decision-making. There is 1 hour 15 minutes to complete the experiment. At the conclusion of the experiment, you will be paid for your participation. You will receive payment via bank transfer. The amount of money you will receive depends partly on your decisions and partly on decisions of other participants.

In order to use your data for research, you will need to complete the entire experiment. However, your participation is voluntary. You may choose to stop participating at any time.

If you do choose to complete the experiment, you will earn a completion payment of £3.50. This completion fee is, in addition, to the amount that you can earn from the decision rounds---which ranges from £1.5-£23. You will not earn any money if you do not complete the experiment. We ask that you keep yourself free of distractions during the experiment. Please turn off your mobile phones and close any distracting computer programs. In addition, please do not communicate with anyone before you complete the experiment.

On the bottom right hand corner of each screen, there will be a blue circle with a speech bubble. This is a chat box. If you have any questions or technical issues, you can contact one of the experimenters using the chat box. The chat box will be there throughout the study; the experimenters will do their best to respond to any questions and fix any issues you are experiencing.

Next, we will provide instructions explaining how the experiment works and how you will be paid. As a reminder, you will never be lied to during this or any experiment at ELFE. So, in particular, the instructions explaining how the experiment works and how you are paid are indeed true.

When you are ready, please click "Next" to go on.

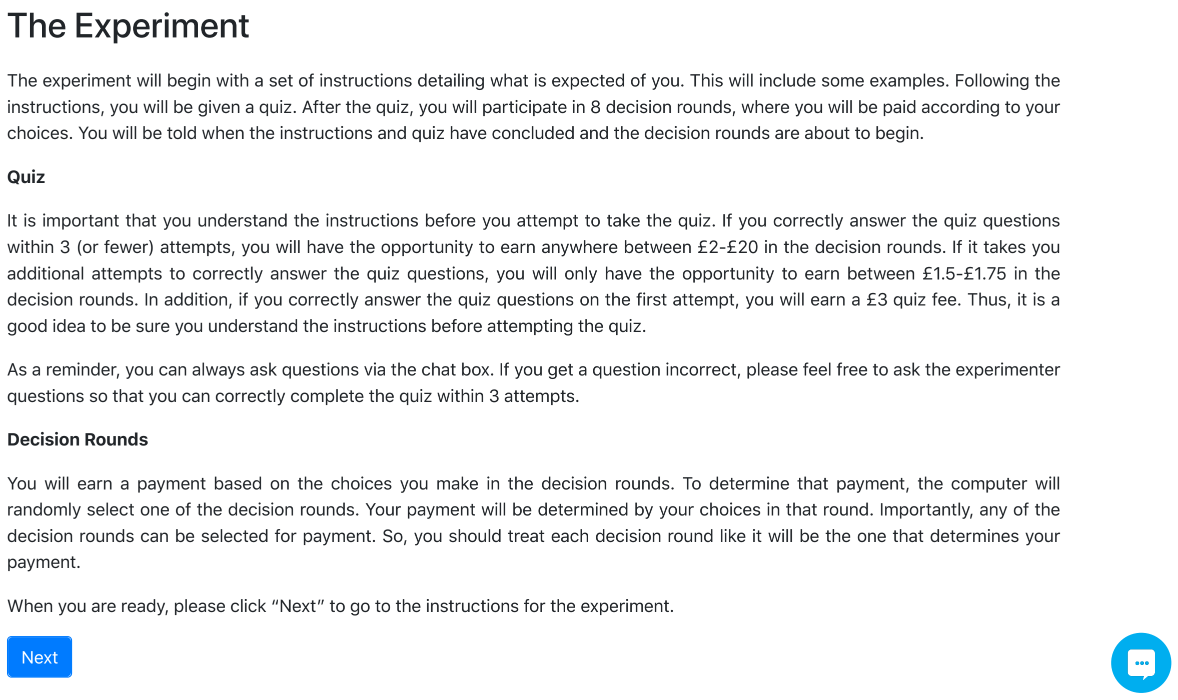

The Experiment

The experiment will begin with a set of instructions detailing what is expected of you. This will include some examples. Following the instructions, you will be given a quiz. After the quiz, you will participate in 8 decision rounds, where you will be paid according to your choices. You will be told when the instructions and quiz have concluded and the decision rounds are about to begin.

Quiz

It is important that you understand the instructions before you attempt to take the quiz. If you correctly answer the quiz questions within 3 (or fewer) attempts, you will have the opportunity to earn anywhere between £2-£20 in the decision rounds. If it takes you additional attempts to correctly answer the quiz questions, you will only have the opportunity to earn between £1.5-£1.75 in the decision rounds. In addition, if you correctly answer the quiz questions on the first attempt, you will earn a £3 quiz fee. Thus, it is a good idea to be sure you understand the instructions before attempting the quiz.

As a reminder, you can always ask questions via the chat box. If you get a question incorrect, please feel free to ask the experimenter questions so that you can correctly complete the quiz within 3 attempts.

Decision Rounds

You will earn a payment based on the choices you make in the decision rounds. To determine that payment, the computer will randomly select one of the decision rounds. Your payment will be determined by your choices in that round. Importantly, any of the decision rounds can be selected for payment. So, you should treat each decision round like it will be the one that determines your payment.

When you are ready, please click “Next” to go to the instructions for the experiment.

Instructions

You will be randomly matched with three other participants: Participant 2, Participant 3 and Participant 4. Both you and each of these participants will face 8 decision rounds; in each decision round, you will each make a choice. During the experiment, you will not learn any information about the choices made by other participants and they will not learn any information about the choices you make.

We now describe the decision problem that you will face in each round. You must choose 1 of 4 actions: a, b, c, or d. The three participants you are matched with will also choose 1 of 4 actions. Your earnings will depend on both the action you choose and the action that Participant 2 chooses. The table titled "Your Earnings" (left-hand table below) represents the possible earnings you can receive. Your action determines the row of the table and Participant 2's action determines the column of the table. So you choose amongst actions a, b, c, or d and Participant 2 chooses amongst actions e, f, g, or h. The cell that corresponds to this combination of actions will determine your earnings.

Your Earnings

		Participant 2's action			
		e	f	g	h
Your action	a	1	2	3	4
	b	5	6	7	8
	c	9	10	11	12
	d	13	14	15	16

Participant 2's Earnings

		Participant 3's action			
		i	j	k	l
Participant 2's action	e	17	18	19	20
	f	21	22	23	24
	g	25	26	27	28
	h	29	30	31	32

Participant 3's Earnings

		Participant 4's action			
		m	n	o	p
Participant 3's action	i	33	34	35	36
	j	37	38	39	40
	k	41	42	43	44
	l	45	46	47	48

Participant 4's Earnings

		Your action			
		a	b	c	d
Participant 4's action	m	49	50	51	52
	n	53	54	55	56
	o	57	58	59	60
	p	61	62	62	64

You can easily see the earnings you will get by clicking on rows and columns of the earnings tables. Clicking on a row action will highlight the associated row and show all the possible earnings you can get if you choose that row. Clicking on a column action will highlight the associated column and show all the possible earnings you can get if Participant 2 chooses that column. By highlighting both a row and column, you can see a darker red box that shows what you would get if you choose that row and Participant 2 chooses that column. At any time, you can unhighlight a row or column by clicking on it again. You can also highlight multiple rows or columns at the same time.

Try out clicking the row where you choose a and Participant 2 chooses e. There, you would earn £1, as illustrated by the darker red box. Now try instead clicking the row where you choose b and Participant 2 chooses f. Now you would earn £6.

The earnings tables also show earnings of Participants 2, 3, and 4. Participant 2 can choose one action from e, f, g, h; Participant 3 can choose an action from i, j, k, l; and Participant 4 can choose an action from m, n, o, p. Notice, Participant 2's earnings depend upon the action they choose and the action Participant 3 chooses; Participant 3's earnings depend upon the action they choose and the action Participant 4 chooses; Participant 4's earnings depend upon the action they choose and the action that you choose.

Notice, you can click the rows and columns in each Participant's earning table to see what they will earn. For instance, let's look at Participant 3's earnings table. Try clicking the row where Participant 3 chooses j and the column where Participant 4 chooses m. In that case, Participant 3 will get £37.

Importantly, the earnings tables will differ from round-to-round. You should always look carefully at the earnings tables at the beginning of each round to determine your earnings for that round.

You will be required to spend at least 90 seconds on each round. You may spend more time on each round, if you wish.

The “Next” button will appear after 2 minutes. Please use this time to carefully read the instructions. When you understand these instructions, please click “Next” to go to final remarks on the experiment. You can always come back to this page by clicking the button called "Instructions."

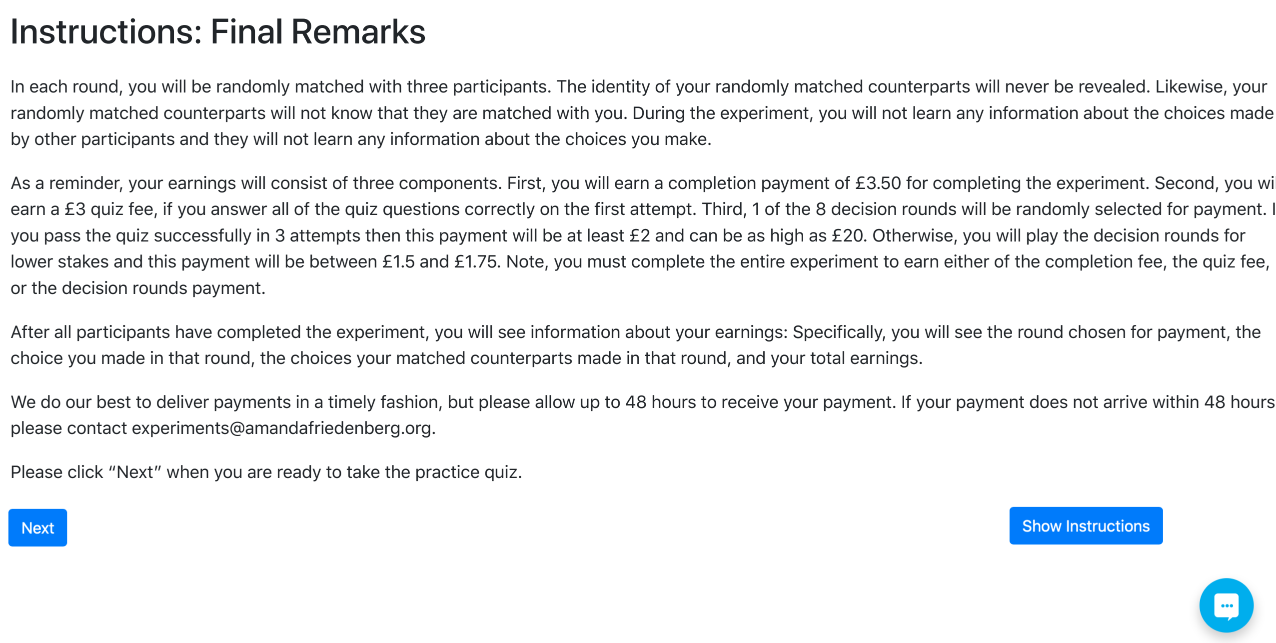

Instructions: Final Remarks

In each round, you will be randomly matched with three participants. The identity of your randomly matched counterparts will never be revealed. Likewise, your randomly matched counterparts will not know that they are matched with you. During the experiment, you will not learn any information about the choices made by other participants and they will not learn any information about the choices you make.

As a reminder, your earnings will consist of three components. First, you will earn a completion payment of £3.50 for completing the experiment. Second, you will earn a £3 quiz fee, if you answer all of the quiz questions correctly on the first attempt. Third, 1 of the 8 decision rounds will be randomly selected for payment. If you pass the quiz successfully in 3 attempts then this payment will be at least £2 and can be as high as £20. Otherwise, you will play the decision rounds for lower stakes and this payment will be between £1.5 and £1.75. Note, you must complete the entire experiment to earn either of the completion fee, the quiz fee, or the decision rounds payment.

After all participants have completed the experiment, you will see information about your earnings: Specifically, you will see the round chosen for payment, the choice you made in that round, the choices your matched counterparts made in that round, and your total earnings.

We do our best to deliver payments in a timely fashion, but please allow up to 48 hours to receive your payment. If your payment does not arrive within 48 hours, please contact experiments@amandafriedenberg.org.

Please click “Next” when you are ready to take the practice quiz.

Screenshot of Quiz Below is screenshots of the quiz.

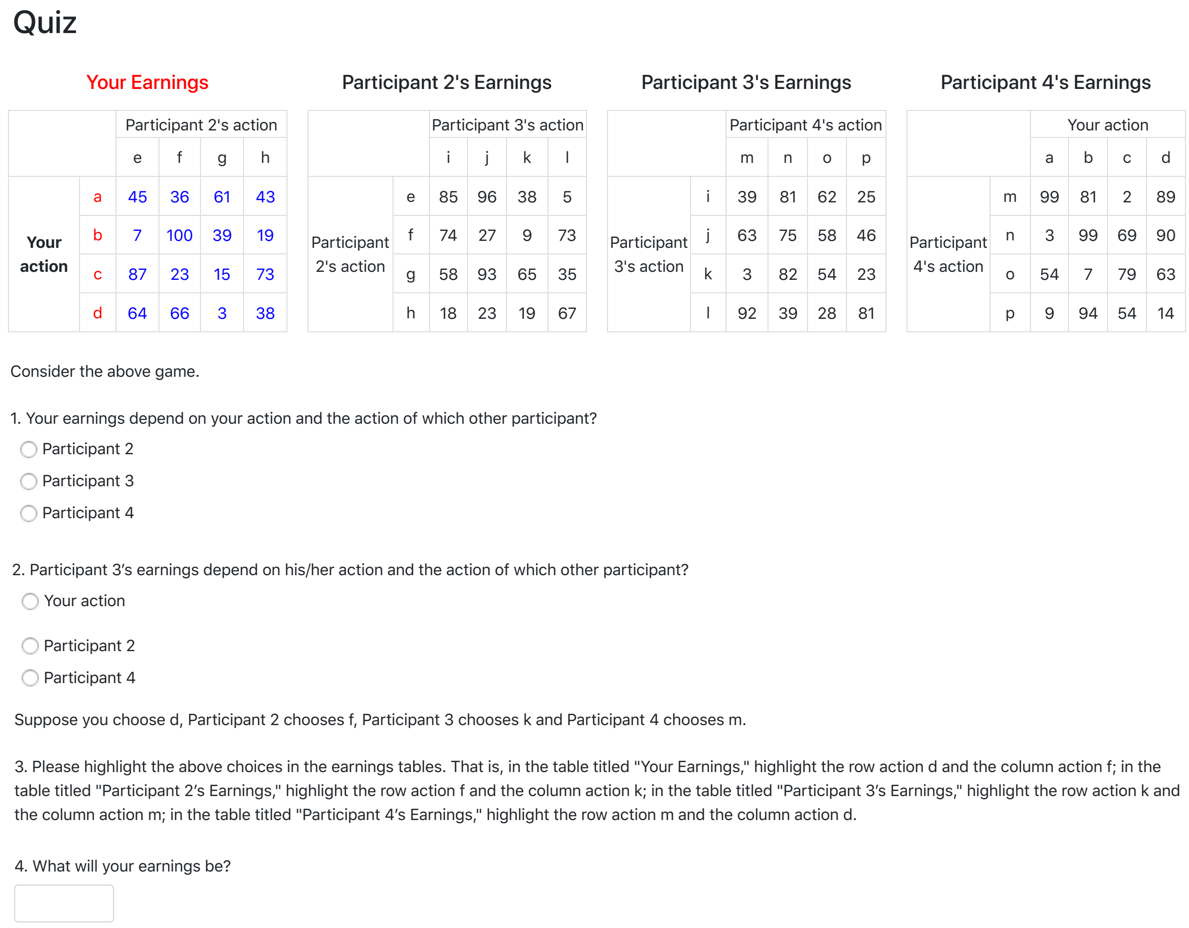

Quiz

Your Earnings

		Participant 2's action			
		e	f	g	h
Your action	a	45	36	61	43
	b	7	100	39	19
	c	87	23	15	73
	d	64	66	3	38

Participant 2's Earnings

		Participant 3's action			
		i	j	k	l
Participant 2's action	e	85	96	38	5
	f	74	27	9	73
	g	58	93	65	35
	h	18	23	19	67

Participant 3's Earnings

		Participant 4's action			
		m	n	o	p
Participant 3's action	i	39	81	62	25
	j	63	75	58	46
	k	3	82	54	23
	l	92	39	28	81

Participant 4's Earnings

		Your action			
		a	b	c	d
Participant 4's action	m	99	81	2	89
	n	3	99	69	90
	o	54	7	79	63
	p	9	94	54	14

Consider the above game.

1. Your earnings depend on your action and the action of which other participant?

○ Participant 2
○ Participant 3
○ Participant 4

2. Participant 3’s earnings depend on his/her action and the action of which other participant?

○ Your action

○ Participant 2
○ Participant 4

Suppose you choose d, Participant 2 chooses f, Participant 3 chooses k and Participant 4 chooses m.

3. Please highlight the above choices in the earnings tables. That is, in the table titled "Your Earnings," highlight the row action d and the column action f; in the table titled "Participant 2’s Earnings," highlight the row action f and the column action k; in the table titled "Participant 3’s Earnings," highlight the row action k and the column action m; in the table titled "Participant 4’s Earnings," highlight the row action m and the column action d.

4. What will your earnings be?

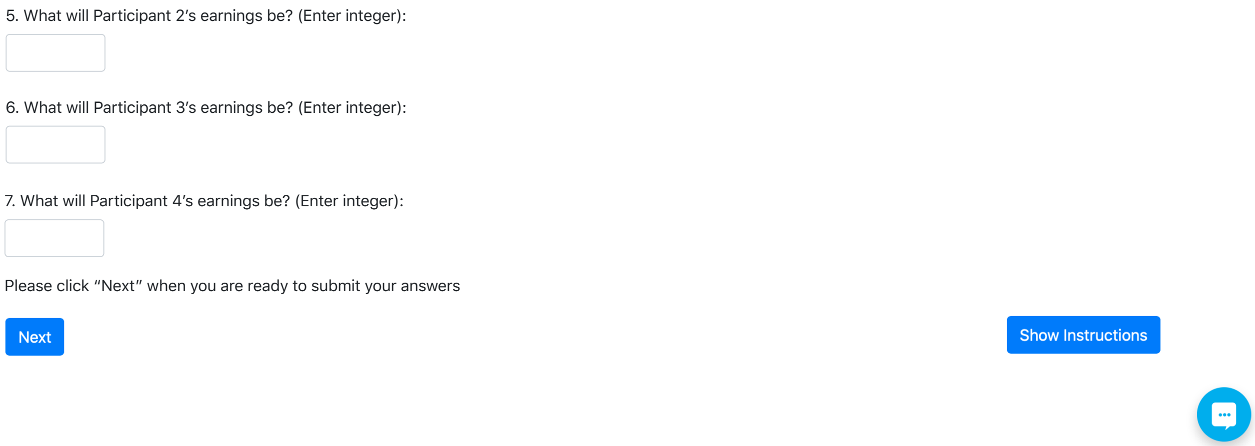

5. What will Participant 2’s earnings be? (Enter integer):

6. What will Participant 3’s earnings be? (Enter integer):

7. What will Participant 4’s earnings be? (Enter integer):

Please click “Next” when you are ready to submit your answers

Screenshot of Example Game Below is a screenshot of Treatment 1, Game G, role P4.

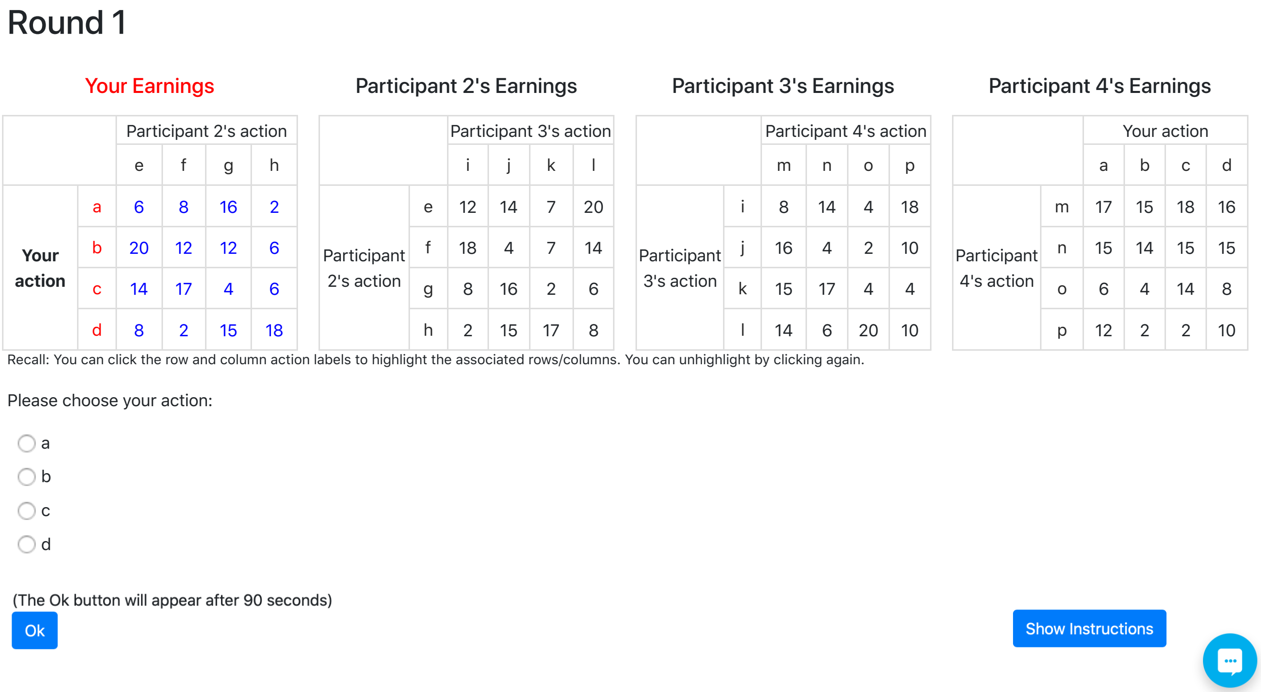

Round 1

Your Earnings

		Participant 2's action			
		e	f	g	h
Your action	a	6	8	16	2
	b	20	12	12	6
	c	14	17	4	6
	d	8	2	15	18

Participant 2's Earnings

		Participant 3's action			
		i	j	k	l
Participant 2's action	e	12	14	7	20
	f	18	4	7	14
	g	8	16	2	6
	h	2	15	17	8

Participant 3's Earnings

		Participant 4's action			
		m	n	o	p
Participant 3's action	i	8	14	4	18
	j	16	4	2	10
	k	15	17	4	4
	l	14	6	20	10

Participant 4's Earnings

		Your action			
		a	b	c	d
Participant 4's action	m	17	15	18	16
	n	15	14	15	15
	o	6	4	14	8
	p	12	2	2	10

Recall: You can click the row and column action labels to highlight the associated rows/columns. You can unhighlight by clicking again.

Please choose your action:

○ a
○ b
○ c
○ d

(The Ok button will appear after 90 seconds)

Appendix D Additional Tables and Figures

P1	P2	P3	P4
(a, d_*)	(a, a_*)	(a, a_*), (d, d_*)	(b, b_*), (c, c_*),(d, d_*)
(a, d_*)	(a, a_*)	(b, b_*)	(b, b_*), (d, d_*)
(a, d_*)	(b, b_*)	(a, a_*), (b, b_*)	(b, b_*), (d, d_*)
(a, d_*)	(c, c_*)	(a, a_*)	(b, b_*), (c, c_*)
(a, d_*)	(c, c_*)	(b, b_*)	(b, b_*)
(a, d_*)	(c, c_*)	(c, c_*)	(c, c_*)
(a, d_*)	(c, c_*),(d, d_*)	(d, d_*)	(b, b_*), (c, c_*)

Table D.1. Observations Identified as Strategic/Heuristic Bound of 1
For Both IDENT and CS

P1	P2	P3	P4
(a, d_*)	(a, a_*)	(a, a_*), (d, d_*)	(a, a_*)
(a, d_*)	(a, a_*)	(b, b_*)	(a, a_*), (c, c_*)
(a, d_*)	(a, a_*)	(c, c_*)	(a, a_*), (b, b_*), (c, c_*), (d, d_*)
(a, d_*)	(c, c_*)	(a, a_*)	(a, a_*), (d, d_*)
(a, d_*)	(c, c_*)	(b, b_*)	(a, a_*), (c, c_*), (d, d_*)
(a, d_*)	(c, c_*)	(c, c_*)	(a, a_*), (b, b_*), (d, d_*)
(a, d_*)	(c, c_*), (d, d_*)	(d, d_*)	(a, a_*), (d, d_*)
(a, d_*)	(d, d_*)	(a, a_*), (b, b_*), (c, c_*)	(a, a_*), (b, b_*), (c, c_*), (d, d_*)

Table D.2. Observations Identified as Strategic Bound of 2: Constant P2
Shaded is identified in both IDENT and CS; non-shaded is only identified in IDENT.

P1	P2	P3	P4
(a, d_*)	(d, d_*)	(d, d_*)	(a, a_*), (d, d_*)
(a, d_*)	(d, d_*)	(a, a_*), (b, b_*), (c, c_*)	(a, a_*), (b, b_*), (c, c_*), (d, d_*)

Table D.3. Observations Identified as Heuristic Bound of 2: Constant P2
For Both IDENT and CS

	P2		P3		P4	
	IDENT	CS	IDENT	CS	IDENT	CS
SB Model	$(a, c_*), (a, d_*), (b, a_*), (b, d_*), (c, a_*), (d, a_*), (d, b_*), (d, c_*)$	$(a, c_*), (a, d_*), (b, a_*), (b, d_*), (c, a_*), (c, b_*), (c, d_*), (d, a_*), (d, c_*)$	$(a, a_*),(b, b_*),(c, c_*),(d, d_*)$		$(a, a_*),(b, b_*),(c, c_*),(d, d_*)$	
HB Model	$(b, a_*), (c, a_*), (d, a_*)$	$(a, d_*), (b, a_*), (b, d_*), (d, a_*)$	$(a, a_*),(b, b_*),(c, c_*),(d, d_*)$		$(a, a_*),(b, b_*),(c, c_*),(d, d_*)$	

Table D.4. Identified as Strategic/Heuristic Bound of 2: Non-Constant P2

	P2		P3		P4	
	IDENT	CS	IDENT	CS	IDENT	CS
SB Model	$(a, a_*), (a, c_*), (a, d_*), (b, a_*), (b, d_*), (c, a_*), (c, c_*), (d, a_*), (d, b_*), (d, c_*), (d, d_*)$	$(a, c_*), (a, d_*), (b, a_*), (b, d_*), (c, a_*), (c, b_*), (c, c_*), (c, d_*), (d, a_*), (d, c_*), (d, d_*)$	$(a, b_*), (a, c_*), (a, d_*), (b, a_*), (b, c_*), (b, d_*), (c, a_*), (c, b_*), (c, d_*), (d, a_*), (d, b_*), (d, c_*)$	$(a, b_*), (a, c_*), (a, d_*), (b, a_*), (b, c_*), (b, d_*), (c, a_*), (c, b_*), (c, d_*), (d, a_*), (d, b_*), (d, c_*)$	$(a, a_*),(b, b_*),(c, c_*),(d, d_*)$	
HB Model	$(b, a_*), (c, a_*), (d, a_*), (d, d_*)$	$(a, d_*), (b, a_*), (b, d_*), (d, a_*), (d, d_*)$	$(a, b_*), (c, b_*),(d, a_*),(d, b_*)$	$(a, b_*), (b, a_*), (c, a_*),(c, b_*)$	$(a, a_*), (b, b_*), (c, c_*), (d, d_*)$	

Table D.5. Identified as Strategic/Heuristic Bound of 3

	P2		P3		P4	
	IDENT	CS	IDENT	CS	IDENT	CS
SB Model	$(a, a_*), (a, c_*), (a, d_*), (b, a_*), (b, d_*), (c, a_*), (c, c_*), (d, a_*), (d, b_*), (d, c_*), (d, d_*)$	$(a, c_*), (a, d_*), (b, a_*), (b, d_*), (c, a_*), (c, b_*), (c, c_*), (c, d_*), (d, a_*), (d, c_*), (d, d_*)$	$(a, a_*), (a, b_*), (a, c_*), (a, d_*), (b, a_*), (b, b_*), (b, c_*), (b, d_*), (c, a_*), (c, b_*), (c, c_*), (c, d_*), (d, a_*), (d, b_*), (d, c_*), (d, d_*)$	$(a, a_*), (a, b_*), (a, c_*), (a, d_*), (b, a_*), (b, b_*), (b, c_*), (b, d_*), (c, a_*), (c, b_*), (c, c_*), (c, d_*), (d, a_*), (d, b_*), (d, c_*), (d, d_*)$	$(a, b_*), (a, c_*), (a, d_*), (b, a_*), (b, c_*), (b, d_*), (c, a_*), (c, b_*), (c, d_*), (d, a_*), (d, b_*), (d, c_*)$	$(a, b_*), (a, c_*), (a, d_*), (b, a_*), (b, c_*), (b, d_*), (c, a_*), (c, b_*), (c, d_*), (d, a_*), (d, b_*), (d, c_*)$
HB Model	$(b, a_*), (c, a_*), (d, a_*), (d, d_*)$	$(a, d_*), (b, a_*), (b, d_*), (d, a_*), (d, d_*)$	$(a, a_*), (a, b_*), (c, b_*),(d, a_*),(d, b_*)$	$(a, a_*), (a, b_*), (b, a_*), (c, a_*),(c, b_*)$	$(a, c_*), (b, c_*), (d, b_*), (d, c_*)$	$(a, b_*),(a, c_*), (b, c_*), (c, b_*),(d, b_*), (d, c_*)$

Table D.6. Identified as Strategic/Heuristic Bound of 4