Title
Essays on Strategic Sophistication

Author
Zhao, Wanqun

Publication Date
2021

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA SAN DIEGO

Essays on Strategic Sophistication

A dissertation submitted in partial satisfaction of the
requirements for the degree
Doctor of Philosophy

in

Management

by

Wanqun Zhao

Committee in charge:

Professor Anya Samek, Chair
Professor Charles Sprenger, Co-Chair
Professor Vincent Crawford
Professor Uri Gneezy
Professor Marta Serra Garcia
Professor Isabel Trevino

2021

The dissertation of Wanqun Zhao is approved, and it is acceptable in quality and form for publication on microfilm and electronically.

University of California San Diego

2021

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGEMENTS

I would like to thank my advisor, Prof. Anya Samek for introducing me to the field of behavioral and experimental economics. The time spent at BEE Lab has sparked my interest in the field and motivated me to go to graduate school to learn more about doing research. I would like to thank my advisor, Prof. Charles Sprenger for supporting me to pursue my interest and guiding me to become a better researcher. Without his trust and encouragement, I could not have reached this current level of success. I would like to thank Prof. Vincent Crawford for all the thought-provoking conversations and detailed feedback. I have benefited enormously from all the interactions. I also remain grateful for all my committee members who have guided me throughout my journey here. You are all special to me.

I would also like to thank my parents for supporting me in studying overseas. Special thanks to my husband Fangzhou, for helping me debugging my codes, and believing in me when I was discouraged. Thank you for being there for me all the time.

Chapter 1, in full, is a reprint of the material as it appears in Games, vol. 11(3), 2020. The dissertation author was the primary investigator and author of this paper.

Chapter 2, in part is currently being prepared for submission for publication of the material. It is coauthored with Sprenger, Charles. The dissertation author was the primary author of this chapter.

Chapter 3, in part is currently being prepared for submission for publication of the material.

VITA

2015	B. A. in Economics *with honor*, University of Wisconsin, Madison
2015	B. A. in Mathematics *with honor*, University of Wisconsin, Madison
2015-2021	Graduate Teaching Assistant, University of California San Diego
2021	Ph. D. in Management, University of California San Diego

PUBLICATIONS

Wanqun Zhao, "Cost of Reasoning and Strategic Sophistication", *Games*, vol. 11(3), 2020.

ABSTRACT OF THE DISSERTATION

Essays on Strategic Sophistication

by

Wanqun Zhao

Doctor of Philosophy in Management

University of California San Diego, 2021

Professor Anya Samek, Chair
Professor Charles Sprenger, Co-Chair

This dissertation investigates the persistence of strategic sophistication across different scenarios. In daily life, people often make decisions based on their beliefs about other people's actions, and sometimes engage in iterated steps of reasoning. Open questions exist for how contextual aspects influence strategic choices. When given additional information about their opponents, or the strategic environment, do people's strategic choices change

accordingly? This dissertation uses experimental methods to answer these questions. This work helps to inform foundational models of strategic decision making. It points to the significant effects of contextual variables on strategic behaviors that must be incorporated into models of choices.

Chapter 1

Cost of Reasoning and Strategic Sophistication

1.1 Introduction

The use of the level-k model has prevailed in the literature for characterizing people's initial responses in laboratory strategic games(Nagel, 1995; Stahl and Wilson, 1994). The model characterizes the player's systematic deviations from the Nash equilibrium using a bounded rational-type explanation. The level-0 type's action is assumed to be uniformly distributed over all actions (or in some cases, level-0 type's action is the most prominent action available), whereas the level-1 type has the best response to the expected action of the level-0 type. The level-2 type has the best response to the expected action of the level-1 type. The iterations follow this pattern, as the level-k type always has the best response to the actions of level-$k-1$ type. Such patterns of off-equilibrium play have been evidenced in many laboratory experiments. In Nagel's p-beauty contest game, Nagel found spikes that correspond to the first and second rounds of iterative best responses (Nagel, 1995). Stahl and Wilson found similar evidence of level-1 and level-2 types with 10 matrix games (Stahl and Wilson, 1994). Camerer et al. developed a cognitive hierarchy model (Camerer et al., 2004). Instead of holding a belief that all the other players are type k-1, level-k players in the cognitive hierarchy model assign a probability distribution over all the lower types. Many other studies used the level-k model to explain laboratory data (matrix game: Costa-Gomes et al. (2001); beauty contest game: Ho et al. (1998); Bosch-Domènech et al. (2002); Duffy and Nagel (1997); Grosskopf and Nagel (2008); sequential game: Ho and Su (2013); auction: Crawford and Iriberri (2007); Ivanov et al. (2010); Crawford, Costa-Gomes and Iriberri also provide a comprehensive literature review (Crawford et al., 2013)).

However, although the level-k model has proven its usefulness in characterizing initial responses for many laboratory games, its predictive power remains ambiguous because (1) it is often used posteriorly to classify a player's type given their actions and (2) the model lacks components related to individual characteristics that could help identify different types of players. It is important to understand how certain levels are reached for each individual, as it is a starting point for the discussion of the model's predictive power. Alaoui and Penta developed a framework called the endogenous depth of reasoning (EDR) model to explain what may happen in a player's head when they encounter a given strategic situation (Alaoui and Penta, 2016). The EDR model captures individual characteristics by introducing cost of reasoning, which is determined both by the strategic environment and by a player's endogenous cognitive ability. The model includes game-specific characteristics by introducing the benefit of reasoning through payoffs of the games. Lastly, the model allows a clear separation of cognitive bounds and behavioral levels observed in games by introducing higher-order beliefs. Such separation makes room for individual adjustments of *k*-levels in different strategic environments. As a result, a level-1 action observed from a game does not necessarily classify the player as a level-1 player. Instead, such action can be a product of the player's cost and benefit analysis and his belief about his opponents.

The EDR model provides a plausible starting point to study the persistence of the level-*k* model. However, as individuals have heterogeneous costs of reasoning and belief systems in all kinds of strategic situations, it is hard to conduct direct comparisons across games to test whether the behavioral *k*-levels follow the EDR model's predictions. In this paper, I use Costa-Gomes and Crawford's two-person guessing games (henceforth CGC06) and cognitive load to create different strategic environments (Costa-Gomes and Crawford, 2006). By controlling cognitive load, I create a standard for the cost of reasoning for all the subjects. Although individual cognitive ability may still have an effect, by using a within-subject experimental design, the individual effect will no longer impact the comparisons of strategic levels across games for the same subject. The revelation of information about the strategic environment is also carefully manipulated to clearly control the subject's belief space. The goal was to test whether the EDR model provides directional predictions about the changes on k-levels across games for any given subject. Alaoui and Penta tested the benefit part of their model using the 11–20 money request game with altered bonus rewards (Alaoui and Penta, 2016; Arad and Rubinstein, 2012b). To the best of my knowledge, this was the first paper to provide experimental tests of the EDR model by introducing different strategic environments with controlled cost and belief space.

With the 18 two-person guessing games in the experiment, the results suggest that the subject's behavioral levels systematically vary across the games. Subjects are mostly responsive to the changes in the strategic environment. Their directional changes in behavioral levels can be predicted by the EDR model when they are more cognitively capable or their opponent is less cognitively capable. An inherent cognitive bound exists for the subjects in different strategic environments. When comparing a subject's behavioral levels across all the games while providing the same amount of cognitive resources, their behavioral levels rarely exceed their cognitive bound level for that strategic environment.

A few other papers also studied the correlation of individual k-levels with cognitive ability. Allred et al. investigated the effects of cognitive load on strategic sophistication (Allred et al., 2016). In their experiments, they asked the subjects to perform a memorization task of either a three- or nine-digit binary number concurrently with strategic games such as beauty contest, 11–20, and 10 matrix games. They found that subjects with high loads (i.e., nine-digit number) were less capable of computing best responses, especially for the beauty contest game. They were also aware of their strategic disadvantages. The net result of cognitive load depended on the specific strategic context. Burnham et al. used a standard psychometric test to measure the cognitive abilities of their subjects, and correlated the test results with subjects' performances in a p-beauty contest game (Burnham et al., 2009). They found a negative correlation between cognitive test scores and entries in the beauty contest game, indicating that subjects with higher cognitive ability tend to be more strategically sophisticated in such games. Gill and Prowse used a 60-question non-verbal Raven test to assign subjects into high- and low-cognitive-ability groups (Gill and Prowse, 2016). They asked the subjects to play a p beauty contest game for 10 rounds, and found that subjects in the high-cognitive-ability group converged to equilibrium faster. These studies provided some evidence of the correlation of individual k-levels with cognitive ability or carefully controlled cognitive tasks. In my experiment, I used memorization tasks to manipulate the cost of reasoning for the subjects in the context of a two-person guessing game. According to Allred et al., higher cognitive load negatively affects a subject's ability to calculate the best responses in this type of guessing games (Allred et al., 2016). To attain a higher level of strategic sophistication, players have to exert more effort to combat the effects of cognitive load; therefore, the cost of reasoning increases with cognitive load in this strategic situation. Every subject experienced both the low and high cognitive loads at some point during the experiment, so they were fully aware of the additional cost of reasoning that was added by these memorization tasks. As a result, their cost of reasoning and their belief about their opponent's cost of reasoning can be quantified by the cognitive load.

The stability of k-levels is an important aspect in the level-k model literature. Stahl and Wilson used twelve normal-form games to estimate the player's level (Stahl and Wilson, 1995). They found that using a relatively low threshold, 35 out of 48 subjects could be classified as stable across games. Fragiadakis et al. asked the subjects to repeat their decisions in a series of two-person guessing games to subsequently best respond to their past actions (Fragiadakis et al., 2016). They found that only 40% of the subjects who were able to replicate the decisions could be classified as a known behavioral type. A few works mentioned the predictive power of strategic sophistication. Arad and Rubinstein used a multidimensional Colonel Blotto game to observe subject's multidimensional iterative reasoning process (Arad and Rubinstein, 2012a). They found that subjects with a higher level of reasoning in the 11–20 money request game also seem to have more rounds of iterative reasoning in this game.

Perhaps the most closely related work to this paper is Georganas, Healy, and Weber's 2015 paper (Georganas et al., 2015). They conducted an experiment to examine the cross-game stability of the *k*-levels. They used four matrix undercutting games and six two-person guessing games and compared them at the individual level. They found no correlation between the levels of reasoning across games. However, they found some evidence of cross-game stability within the class of undercutting game. I studied a similar question to the cross-game stability of the level-*k* model. Instead of introducing a second family of games, I used cognitive load to mimic different strategic environments, and restricted the subjects to fixed pairs while playing the games. The belief space was therefore carefully controlled, and the uncertainty from playing against a new random player for each round was completely eliminated. The data suggested that systematic level changes can be predicted by the EDR model under certain conditions. In Section 1.2, I provide a brief introduction to the EDR model to cover some necessary background and theoretical predictions. In Section 1.3, the experimental design is introduced in detail. Sections 1.4 and 1.5 cover the data analysis procedure and the discussion of the results, respectively. Section 1.6 provides the concluding remarks.

1.2 Theoretical Consideration

1.2.1 Model

I adopted Alaoui and Penta's EDR model for theoretical predictions (Alaoui and Penta, 2016). In this model, players follow an endogenous reasoning process that determines the strategic bound in a particular context. With added structure on beliefs, the model is able to predict a player's actual level of play in any game that could use a *k*-level iterative best response reasoning process. The main benefit of using this model is that the structure of the model allowed me to conduct a comparative statics exercise on a player's reasoning process. One of the main goals of this study was to conduct a comparative static exercise on the cost side. Below, I provide more detailed descriptions of some key features of this model. These features are relevant to the experimental design and predictions for this paper.

A player's cognitive bound is a mapping from the incremental cost of reasoning ($c(k)$) and the incremental value of reasoning ($v(k)$) at each level to the intersection of the two terms.

$$\kappa(v,c) = min\{k \in \mathbb{N} \mid v(k) \geq c(k) \text{ } and \text{ } v(k+1) < c(k+1)\} \tag{1.1}$$

A player reaches their cognitive bound at the *k*th level by having a value of reasoning for that level exceeds cost of reasoning, but their cost–benefit analysis no longer supports the one-higher level (i.e., $k+1$) of reasoning. Further denote the cognitive bound of player i as $\bar{k}_i$, where:

$$\bar{k}_i = \kappa(v_i, c_i). \tag{1.2}$$

According to Alaoui and Penta, the value of reasoning is affected by the payoff of the game (Alaoui and Penta, 2016). The cost of reasoning is an endogenous characteristic of an individual, which is largely related to their cognitive or reasoning ability. In this paper, I take their assumption on the value of reasoning and continue to assume that the payoff is the only incentive for players to apply logical reasoning in the games. I provide a further discussion on the cost of reasoning. Beyond an individual's endogenous ability, the strategic environment (such as cognitive load) provides many challenges for a person in applying strategic reasoning, which alters the cost of reasoning.

A player's belief is represented as a tuple. Since the game in my design is symmetric in payoffs, a player's belief can be restricted to the beliefs about the cost of reasoning. Therefore, the first element of the tuple, c_i, represents player i's own cost of reasoning. The second element is player i's beliefs of his opponent's (player j) cost of reasoning, denoted as c^i_j. The last element c^{ij}_i is player i's second-order belief, which is their belief about player j's belief of themselves. Any higher-order beliefs could be nested to the first- and second-order beliefs; therefore, a player's belief is represented as:

$$t_i = (c_i, c^i_j, c^{ij}_i). \tag{1.3}$$

1.2.2 Theoretical Predictions

I formulated the testable predictions following the EDR construction discussed in Section 1.2.1. For any game $G = \{X_i, u_i\}_{i=1,2}$, let $k_i(x_i)$ be the reflected behavioral level of player i by choosing action x_i, where X_i is the set of actions available for player i and u_i is the payoff function for player i.

1. **Changing the cost of reasoning**: For any c^i_j and c^{ij}_i, $k_i(x_i)$ (weakly) decreases with c_i. Fixing player i's first- and second-order beliefs, their cognitive bound weakly decreases with the cost of reasoning. The observed level of player i from the game will also weakly decrease. In my design, for the first 16 games holding cognitive load and information structure constant for the opponent, players will display lower strategic levels when the memorization task is a string of seven letters.

2. **Changing the opponent's cost of reasoning**: For any c_i and c^{ij}_i, $k_i(x_i)$ (weakly) decreases with c^i_j. If $c^{ij}_i = c_i$, then player i's cognitive bound is binding if they regard their opponent as more sophisticated.

   Player i reacts to the change in the cost of reasoning of their opponents. More specifically, if he observes his opponent's cost of reasoning increasing, he will adjust their strategy in the game to best respond to his opponent. That means they may choose to take an action that corresponds to a lower level of strategic sophistication. However, such adjustments of strategies are binding by the cognitive bound when the player believes their opponent has a lower cost of reasoning compared to their own cost. In the context of my experiment, a player should choose a weakly lower level of strategy if he observes his opponent's memorization task becoming more difficult (i.e., from a string of three letters to a string of seven letters).

3. **Changing the second-order belief**: For any c_i and c_j^i, $k_i(x_i)$ (weakly) decreases with c_i^{ij}. If $c_i \geq c_i^{ij}$, then player i's cognitive bound is binding. By fixing player i's own cost of reasoning and his opponent's cost, through only changing player i's second-order belief, player i should adjust their strategic actions. For example, when a player has a low cost of reasoning in the game, if they believe that their opponent has a wrong belief about themselves, namely, they believe that their opponent thinks the cost of reasoning for them is very high, then they can switch to an action that is associated with a lower level of reasoning. However, this adjustment of strategic actions according to the second-order belief is restricted by player i's own cognitive bound, meaning that they cannot make any adjustments that requires a higher level of reasoning than their cognitive bound. In the context of my experimental design, players should adjust their actions when the information structure shifts from full revelation of cognitive load to partial revelation.

4. **Cognitive bound**: Given c_i, for any c_j^i and c_i^{ij}, $k_i(x_i)$ never exceeds $\bar{k}_i$. When fixing player i's own cost of reasoning, their behavioral level should never exceed their cognitive bound. In the context of this experiment, on an individual level, actions observed in games 17 and 18 should correspond to the highest level of reasoning that one player can achieve under the respective cognitive load.

1.3 Experimental Design

In this section, I present the details of the experimental design. The experiment captured the process of level-k thinking through the two-person guessing game (Costa-Gomes and Crawford, 2006). I provide a brief introduction to the game first, followed by the treatment design and the experimental timeline.

1.3.1 The Game

The two-person guessing game is an asymmetric, two-player game. Each player has a lower limit, $a_i > 0$, an upper limit, $b_i > 0$, and a target $p_i \in (0,2)$. Players are required to input a guess that is within their lower and upper limit. However, their actual choice is not restricted by the limit. Denote player i's input by x_i. If a player guesses a number x_i that falls outside the limit interval, then their guess will be adjusted to the closest bound. For example, if $x_i < a_i$, then the adjusted guess y_i will be $y_i = a_i$. If $x_i > b_i$, then the adjusted guess y_i is $y_i = b_i$. However, any guess falling within the limit interval will not be adjusted; i.e., $y_i = x_i$.

The goal of the game is to make a guess that minimizes the difference between the player's own guess and the product of their target and his opponent's guess. Denote the difference by $e_i = \mid y_i - p_i \cdot y_j \mid$. The payoff is a quasi-concave function minimized at zero. Player i receives $u_i = \max\{0, 200 - e_i\} + \max\{0, 100 - \frac{e_i}{100}\}$. Since a player's guesses that have the same adjusted inputs will yield the same outcome for the subject, I use the adjusted guess y_i as a proxy of how players perform in the game.

In this game, the level-0 player is assumed to play randomly according to a uniform distribution over the action space. Denote the theoretical predicted guess made by a k-level player as x_i^k. Given the assumption imposed on the level-0 player's strategy, level-1 players will best respond to the expected value of level-0 player's guess, i.e., $x_i^1 = p_i \cdot \mathbb{E}\{y_j \mid y_j \in [a_j, b_j]\}$. The level-2 player's strategy will then be $x_i^2 = p_i \cdot \{1(x_j^1 \in [a_j, b_j]) \cdot x_j^1 + 1(x_j^1 < a_j) \cdot a_j + 1(x_j^1 > b_j) \cdot b_j\}$. The reasoning process follows iterative best responses. It converges to the Nash equilibrium after finite rounds of iterations.

In this paper, I adopt 14 two-person guessing games used by CGC06 and 4 two-person guessing games used by Georganas et al. (Costa-Gomes and Crawford, 2006; Georganas et al., 2015). The parameters of each game are given in Table 1.1. All the players survive at least two rounds of iterative best responses before reaching the equilibrium (as stated in Table 1.1 "steps to eqm" column). Since in CGC06, only a few number of subjects reached level 3 in the reasoning process, the choice of parameters in this paper should be sufficient to identify a player's strategic levels in the game.

1.3.2 Cognitive Load

Before directed to the guessing game, subjects were required to memorize a string of letters and were told that they need to recall the given string after the guessing game. The string was composed of either three or seven random letters, for example, UMH or WIEZOFH. The subjects were given 15 s to memorize the string; then they were automatically directed to the guessing game.

I did not pay the subjects specifically for correct recalls. However, their payments on the guessing game were partially dependent on this memorization task. If the recall for the selected payment round was wrong, they were not paid for that round, and left the experiment with only the participation fee. Said payment scheme incentivized the subjects to memorize the cognitive load correctly, and therefore guaranteed the effects of different cognitive load treatments.

Table 1.1: The eighteen two-person guessing games.

Game #	P1's Limits & Target	P2's Limits & Target	Treatment	P1's role	P2's role	Steps to Eqm	Eqm at Boundary
1	[(100,900); 1.5]	[(300,500); 0.7]	[LL+]	role B	role A	5+	-
2	[(300,900); 1.3]	[(100,500); 0.7]	[HL-]	role B	role A	5+	lower
3	[(300,900); 1.3]	[(300,900); 1.3]	[HH+]	role B	role A	3	upper
4	[(300,900); 0.7]	[(100,900); 1.3]	[LH+]	role A	role B	5+	lower
5	[(100,500); 1.5]	[(100,500); 0.7]	[LH-]	role B	role A	5+	upper
6	[(100,500); 0.7]	[(100,900); 0.5]	[HL+]	role A	role B	5	lower
7	[(100,500); 0.7]	[(100,500); 1.5]	[LH-]	role A	role B	5+	-
8	[(300,500); 0.7]	[(100,900); 1.5]	[LL+]	role A	role B	5+	upper
9	[(100,500); 0.7]	[(300,900); 1.3]	[HL-]	role A	role B	5+	-
10	[(300,500); 0.7]	[(100,900); 0.5]	[HH-]	role B	role A	3	lower
11	[(100,500); 1.5]	[(100,900); 0.5]	[LL-]	role B	role A	5+	-
12	[(300,900); 1.3]	[(300,900); 1.3]	[HH+]	role A	role B	3	upper
13	[(100,900); 1.3]	[(300,900); 0.7]	[LH+]	role B	role A	5+	-
14	[(100,900); 0.5]	[(300,500); 0.7]	[HH-]	role A	role B	4	-
15	[(100,900); 0.5]	[(100,500); 0.7]	[HL+]	role B	role A	4	lower
16	[(100,500); 0.5]	[(100,500); 1.5]	[LL-]	role A	role B	5+	lower
17	[(100,900); 1.3]	[(100,500); 0.5]	L	-	-	5+	-
18	[(100,900); 1.5]	[(100,500); 0.7]	H	-	-	5+	-

1.3.3 Treatments

The experiment consisted of two blocks. In the first block, subjects were assigned into pairs. They played 16 two-person guessing games against each other within the fixed pairs. In the second block, they played two guessing games against the computer. There were a total of 18 two-person guessing games for them to complete for this experiment, and no feedback was given throughout the process.

Against Human

I adopted a $2 \times 2 \times 2$ design. For ease of explanation, I specify the two players in the guessing game as having role A and role B in this section. However, subjects were not aware of their role during the experiment. Each subject was given the role of A or B for each treatment exactly once. I used a within-subject design.

To examine the effects of changing the cost of thinking on a subject's level of reasoning, I varied the cognitive load for role A, holding role B's cognitive load constant. As mentioned in the previous section, role A needed to memorize a string of either three or seven random letters when playing the guessing game. To test the effects of changing the opponent's cost of thinking on a player's level of strategic sophistication revealed in the game, I also varied role B's cognitive load by two levels. Changing the cost of thinking of role B essentially tests the effects of changing the first-order belief for role A. Denote the cognitive load of three letters as low load (L) and seven letters as high load (H).

Lastly, I varied the disclosure of information on the cognitive load for role B. The exact cognitive load implemented on role A was either fully revealed to role B or partially revealed as a probability distribution. Denote full revelation as [+] and the counterpart as [-]. In the partial revelation treatment, role B was told that role A has a 0.5 probability of memorizing a string of three letters and a 0.5 probability of memorizing a string of seven letters. The full and partial revelations of the cognitive load information on role B were a method of measuring the effects of changing the second-order belief for role A. In the full revelation treatment, both roles A and B were aware that role A's memorization task is common knowledge. However, in the partial revelation treatment, role A knew their exact memorization task was hidden to role B; therefore, their second-order belief (i.e., their belief about role B's belief of their own cost of thinking) may not coincide with their actual cost of reasoning. A summary of treatments is provided in Table 1.2. In later sections, I used role A's label to identify the treatments, as I was essentially examining the treatment effects for role A only. The first letter in the label indicates role A's cognitive load (either L or H). The second letter indicates role B's cognitive load (opponent's cognitive load, either L or H), and the last element of the label indicates full or partial revelation (role A's second order belief, either [+] or [-]). Role B served as a supporting role to complete the information required for each treatment. The information presented to role B for each treatment is also presented in Table 1.2. However, when later discussing the experimental results, I only refer to each treatment using role A's label. Table 1.1 provides a summary of treatments and assignments of roles for each game. Each subject played as either role A or role B exactly once for each treatment. There are in total 16 games. For each treatment, the pair of games are symmetric in game parameters and cognitive load realizations. The games were played in two random orders (the first order was as game numbers listed in Table 1.1; the second order was: 2, 13, 14, 4, 3, 1, 16, 6, 11, 8, 12, 5, 10, 15, 7, 9, 18, 17. Since for each game, there were two players assigned with different cognitive loads, considering player 2's order of play, there were essentially four sequences. The number of subjects in each order was roughly balanced. After dropping subjects with missing data, there were 28 subjects playing the first order as player 1, 29 subjects playing the first order as player 2, and 27 subjects playing the second order as player 1 and player 2 respectively.). Before the start of each session, one of the two was randomly selected.

Table 1.2: The eight treatments.

	Role	Label	Cost	1st Order Belief	2nd Order Belief
1	Role A	[LL+]	Low	Low	Low (full revelation)
	Role B		Low	Low	Low (full revelation)
2	Role A	[HL+]	High	Low	High (full revelation)
	Role B		Low	High	Low (full revelation)
3	Role A	[LH+]	Low	High	Low (full revelation)
	Role B		High	Low	High (full revelation)
4	Role A	[HH+]	High	High	High (full revelation)
	Role B		High	High	High (full revelation)
5	Role A	[LL-]	Low	Low	50% Low, 50% High (partial revelation)
	Role B		Low	50% Low, 50% High	Low (full revelation)
6	Role A	[HL-]	High	Low	50% Low, 50% High (partial revelation)
	Role B		Low	50% Low, 50% High	Low (full revelation)
7	Role A	[LH-]	Low	High	50% Low, 50% High (partial revelation)
	Role B		High	50% Low, 50% High	High (full revelation)
8	Role A	[HH-]	High	High	50% Low, 50% High (partial revelation)
	Role B		High	50% Low, 50% High	High (full revelation)

Against Computer

Subjects played against the computer for the second block of the experiment. The computer always chooses a Nash equilibrium action. The concept of equilibrium was explained to the subjects. For example, subjects were told that ”a combination of guesses, one for each person, such that each person’s guess earns them as many points as possible, given the other person’s guess, is called an equilibrium guess.” A similar description of equilibrium guess is found in CGC06. Subjects were also given an example of an equilibrium guess following this description. However, they were not specifically taught how to derive an equilibrium guess. The reason for introducing the equilibrium concept was to encourage the subjects to perform as many rounds of iterative best responses as possible. The two guessing games in this part are labeled 17 and 18 in Table 1.1.

1.3.4 Experimental Timeline

A total of 111 subjects were recruited for this experiment. Sessions were conducted at the Incentive Lab at Rady School of Management, University of California—San Diego (San Diego, CA, USA). The experiment was programmed and conducted using z-tree (Fischbacher, 2007). The complete session lasted for 90 min. Subjects were given a 5 USD show-up fee for attending the experiment and an additional $5 if they passed the understanding test and completed the experiment. They earned an additional $8 on average depending on their decisions for the guessing games. For those who did not pass the understanding test, they spent about 30 min in this experiment and left with the $5 show-up fee.

Subjects were given instructions on the two-person guessing game first. After explaining the rules, I introduced four unincentivized practice rounds. During the practice rounds, subjects played against the computer and were told that the computer will always choose the mean of the target interval. After the subjects made a guess, feedback was provided for the subjects to reflect on the game rule and the payoff rule. An understanding test was then administered. The test was composed of six questions, similar to the understanding test in CGC06. Standard questions included calculatiosn of best responses and payoffs. Although subjects in the experiment were not restricted to following a level-*k* reasoning process, for the purpose of the experiment, I wanted to make sure the subjects were capable of calculating the best responses. A screenshot of the understanding test is shown in Figure 1.1. Subjects needed to answer four out of six questions correctly to proceed to the main part of the experiment.

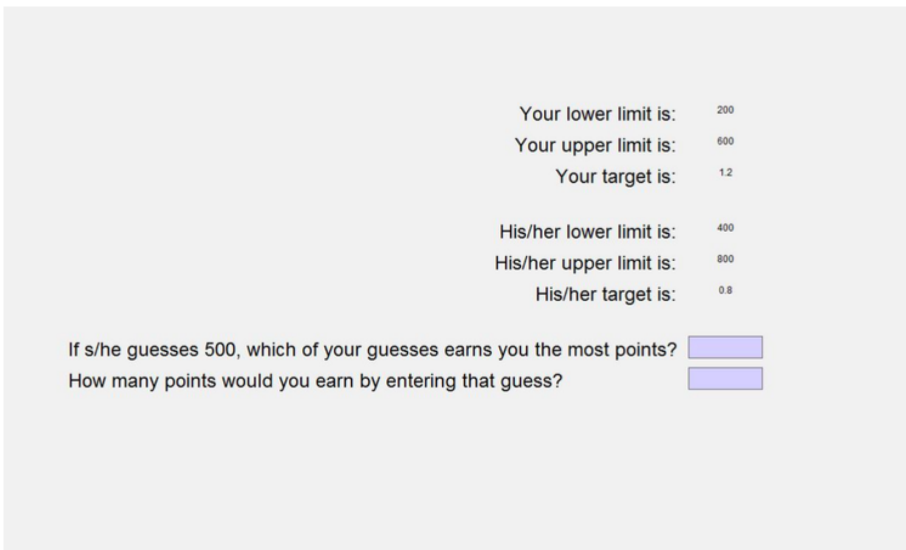

Figure 1.1: Screenshot of the understanding test.

Before playing the incentivized guessing games, subjects were introduced to the memorization task. They were given two unincentivized practice rounds for the low load and high load treatments. During the practice round, they had the standard 15 s to memorize the string of letters and were asked for immediate recall when the time was up. They, however, did no get to practice the guessing game with the cognitive load implemented.

The main experiment consisted of two parts, as discussed in Section 1.3.3. There were 18 two-person guessing games in total. For the first 16 games, subjects were randomly assigned into pairs and stayed within the same pair for all 16 decisions (one as player 1 and the other as player 2). For each game, subjects were given the same information set that consisted of the types of memorization task (either string of three or seven letters, or a probability distribution) for themselves and their opponents, whether their opponents knew about their exact memorization task, and the targets and limits for both players. An example of the actual decision screen is provided in Figure 1.2. Subjects were also asked to elicit their opponents' types of memorization task after they made their guesses and recalled the letters. This practice allowed me to check whether the subjects received and processed the correct information about their strategic environment. There was no feedback given in between the 18 guessing games. This prevented the subjects from learning anything about their opponents' past actions. Such practice also limited the subject's learning of the guessing game, as no payoff information was provided. (There was limited learning of the game. Upon checking a subject's levels with respect to the orders of the games they played, playing a later game was not associated with higher k-levels. The coefficient from the OLS regression was 0.005 and it was not statistically significant.)

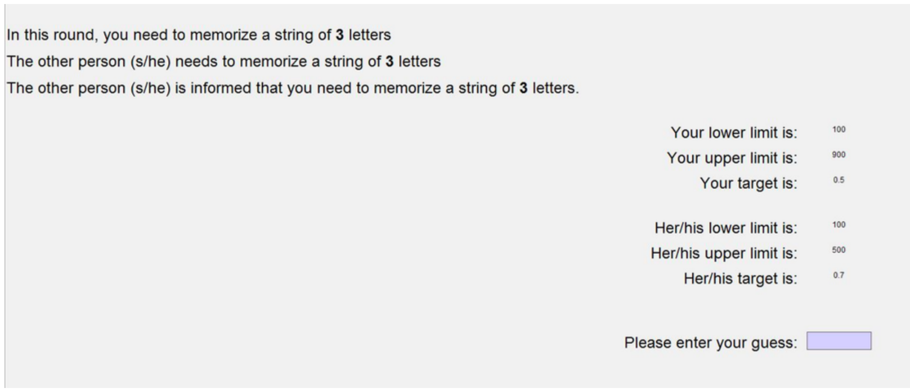

Figure 1.2: Screenshot of the incentivised two-person guessing game.

Subjects took a 10-question Mensa practice test at the end of the experiment. The test is used to measure the subject's analytical ability. Some questions ask the subject to identify the missing element that completes a sequence of patterns or numbers. Some questions are verbal math questions. A couple of studies in economics literature have used a similar test as a measure of cognitive ability (Georganas et al., 2015). I used this test in the experiment to measure whether there were any heterogeneous treatment effects on subjects with different exogenous cognitive abilities.

1.3.5 Discussion of the Experimental Design

First, I used letters to compose strings for the cognitive load treatment, unlike the conventional use of binary numbers (Allred et al., 2016). This design restricts the subjects from using the cognitive load numbers as their inputs for the guessing game. It allows a clear separation of the two tasks, the memorization task and guessing game, and therefore increases the reliability of the treatment effects of cognitive load. I recognized subjects may be able to use other methods to memorize the string of letters, for example, using hand gestures. However, any such method also requires cognitive effort and therefore should not significantly lessen the effects of cognitive load for treatment purposes.

Subjects remained within the fixed pair for the first 16 incentivized guessing games. Since no feedback was given in between games, this design ensures the manipulation of cognitive load being the only source of changing beliefs for any subject. Subjects were different exogenously in terms of cognitive ability, so by staying in the same pair, they carried the same beliefs about their opponents' cognitive abilities throughout the whole session.

Lastly, for each of the 18 incentivized tasks, subjects were given 90 s to make a decision for the guessing game. According to Agranov et al., 90 s is enough for strategic players to make a decision in this type of guessing game (Agranov et al., 2015). To keep the effect of cognitive load constant across players, I only allowed the subjects to submit their guesses after the 90 s was up. Said practice avoids some subjects naïvely picking a guesses without strategic contemplation for the purpose of achieving correct recalls for the memorization task.

1.4 Data Analysis Procedure

All the subjects played 18 games in total, each against a fixed opponent during the experiment. There were 1998 observations of guesses. Grouping the guesses by games, I looked for level shifts observed with raw guesses. This exercise provided a general view of the effects of cognitive load on the games. I also used density plots of the guesses to visualize the treatment effects.

After the exploration of raw guesses, I estimated the level for each guess using the maximum likelihood method. Instead of assuming the subject's behaviors are determined by a single type across all the games, I assumed the subject's behavior in each game was determined by a single type and the types across games were allowed to be different. This was achievable with the design of my experiment with the variations on cognitive load.

Out of 1998 observations, 831 guesses correspond to a type's exact guesses. As about 40% of the observed guesses were a type's exact guesses, I followed the CGC06 approach in my estimation. Specifically, for each player i, game g, and level k, if player i was not making a type's exact guesses in game g, then I defined a likelihood function $L(y_{ig} \mid k, \lambda)$ for each level k for the player in that game, with beliefs $f_g^k(y)$ and sensitivity parameter λ, based on the assumption that they were trying to maximize their expected utility.

Formally, let x_{ig} be the raw guess observed for player i in game g. With the specification of lower limits a_{ig} and upper limits b_{ig}, the adjusted guess is then $y_{ig} = \min\{\max\{a_{ig}, x_{ig}), b_{ig}\}\}$. The density $f_g^k(z)$ represents a subject's belief about his opponent's action given their behavioral level being k. Although in the literature a subject's belief of the other player's level could follow a certain type of distribution, for example, Poisson distribution as in Camerer et al. (2004), in this study, I followed the standard approach that level-k player has point belief about his opponent, that his opponent is level-$(k-1)$ with probability 1. y_g^0 is defined as uniformly spread across the action space. The expected payoff of playing x_{ig} with behavioral level k's belief is then:

$$U_{ig}^k(y_{ig}) = \int_1^{1000} U_{ig}(y_{ig}, z) f_g(z) dz. \tag{1.4}$$

Let $U_{ig}^{k} = [\max(y_g^k - 0.5, a_{ig}), \min(y_g^k + 0.5, b_{ig})]$ be the interval of a type-k subject's exact adjusted guesses, allowing an error of 0.5. Any guess for game g, subject i, who is placed within U_{ig}^{k}, is then identified as an exact match for k-level. Conversely, define ${U_{ig}^{k}}^{\complement} = [a_{ig}, b_{ig}]/U_{ig}^{k}$ as the complement of U_{ig}^{k} within the limit interval for subject i's game g. The likelihood function is then the following:

$$L(y_{ig}|k,\lambda) = \frac{exp[\lambda U_{ig}^{k}(y_{ig})]}{\int_{{U_{ig}^{k}}^{\complement}} exp[\lambda U_{ig}^{k}(w)]dw}. \tag{1.5}$$

Since only one observation was used for the estimation, I took the sensitivity parameter (λ) as 1.33, which is the averaged estimated value of λ in CGC06 with only the subject's guesses. The maximum likelihood estimate of a subject's behavioral level in each game maximizes (1.5) over k, which is:

$$k_{ig} = \underset{k\in\{1,2,3,4,5,6\}}{\arg\max} \; L^{*}(y_{ig}|k). \tag{1.6}$$

To examine the treatment effects on behavioral levels, I pooled guesses into pairs for comparison. For example, to test the prediction on the changing cost of reasoning, I first identified games with the same first-order belief (either low or high cost of reasoning for opponent) and the same second-order belief (partial revelation), and then they were separated into comparison pairs by the subject's cognitive load tasks. The same selection was performed following the conditions listed in each testable prediction.

For each pair of games, I first conducted a binary comparison on their behavioral levels and I report the summary statistics. Since this is essentially a repeated measure of behavioral level from the same sample, I then conducted the Wilcoxon signed-rank test to check the distribution of behavioral levels. Lastly, I ran a GLS random effect regression to examine the treatment effects on behavioral levels. The regression was run by regressing the estimated level on the treatment variable. A subject's cognitive load was coded as 0 when it was in the low load treatment, and 1 when it was a high-load treatment. The same binary coding was also applied to the opponent's cognitive load treatment. The full revelation of information treatment was coded 0, whereas partial revelation was coded 1.

1.5 Results

1.5.1 General Examination of Raw Guesses

There were a total 1998 observations and 831 guesses corresponded to a specific level (levels 1 to 5, and equilibrium). When identifying levels, I assigned the lowest possible level to a guess that matched multiple types. For example, in game 3, equilibrium was reached after three rounds of iterative best responses, and the equilibrium was at the boundary of the target interval. In this case, although levels 3, 4, and 5, and the equilibrium all have corresponding guesses at 900, a subject's guess of 900 only assigned the subject to type level 3. This method of identification restricted over-assignments of the types.

Figure 1.3 shows the distribution of guesses that matched specific levels. Of the 831 guesses that matched a specific level, 43.92% were level 1 guesses, 31.41% were level 2 guesses, 14.20% were equilibrium guesses, and level 3 and higher corresponded to the remaining 10% of the guesses. To provide a clearer picture of the treatment effect, I used a Markov matrix for some treatments with these exactly matched guesses. Tables 1.3 and 1.4 present the level transitions between comparable games. For example, Table 1.3 consists of all the comparable pairs of changing a subject's own cost of reasoning, fixing the opponent with a high cognitive load (game 7 [LH-] and game 14 [HH-]). There were a total of 111 pairs of comparisons, 24 of which had both guesses that exactly matched a specific level. From games 7 to 14, 12 subjects reached level 1 in game 7 and 83.33% stayed at level 1 in game 14. Eight subjects reached level 2 in game 7, 87.5% of which stayed at level 2 and below in game 14. This result largely complies with the theory prediction that increasing cost of reasoning while fixing first- and second-order belief constant decreases the level of reasoning weakly. Likewise, Table 1.4 presents all the comparison pairs of changing the subject's first-order belief while fixing their own cost of reasoning and keeping their second-order belief constant (game 1 [LL+], game 8 [LL+], game 4 [LH+], and game 15 [LH+]). There were a total of 444 pairs of comparison, 99 of which had both guesses matched to a specific level. Forty pairs had level 1 guesses in the [LL+] treatment and 62.5% of them remained level 1 in the [LH+] treatment games. Similarly, 27 pairs had level two guesses in the [LL+] treatment. When changing the subject's first-order belief by increasing the cognitive load of their opponents, about 90% of these pairs had level 2 or lower guesses in the [LH+] treatment. These statistics largely coincided with the theoretical prediction—with increasing the cost of reasoning for the opponents, the subjects adjusted by weakly decreasing their behavioral levels of playing the game. Due to the limited number of exact matches, I was not able to conduct the same exercise for all the treatment pairs. However, complete discussion of the treatment effects is provided below with estimated behavioral levels.

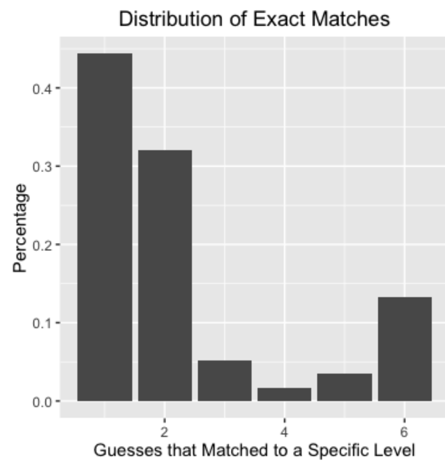

Figure 1.3: Distribution of exact matches.

Table 1.3: Markov matrix of level transitions for increasing cost of reasoning, opponent with high load.

↓ to →	Level 1	Level 2	Level 3	Level 4	Level 5	Eqm	Num
Level 1	**83.33%**	0	16.67%	0	0	0	12
Level 2	**25%**	**62.5%**	12.5%	0	0	0	8
Level 3	0	0	**100%**	0	0	0	1
Level 4	0	0	0	0	0	0	0
Level 5	0	0	0	0	0	0	0
Eqm	0	**66.67%**	**33.33%**	0	0	0	3

Table 1.4: Markov matrix of level transitions for changing first-order belief, subject with low load.

↓ to →	**Level 1**	**Level 2**	**Level 3**	**Level 4**	**Level 5**	**Eqm**	**Num**
Level 1	**62.5%**	7.5%	2.5%	0	15%	12.5%	40
Level 2	**3.7%**	**85.19%**	7.41%	0	0	3.7%	27
Level 3	0	0	0	0	0	0	0
Level 4	0	0	0	0	0	0	0
Level 5	0	0	0	0	0	0	0
Eqm	**31.25%**	**18.75%**	**28.13%**	0	**21.88%**	0	32

The pattern of subjects' adjustments to the changing strategic environment is also illustrated with density plots of each game. This time, all the raw guesses (after adjustments according to upper and lower limits) were used to plot the graphs. Figure 1.4 illustrates the treatment effects for the three theoretical predictions. To better compare across games, level 1 guesses were centered, and all the guesses were adjusted accordingly. The colored vertical lines illustrate the level-exact guesses. For example, in Figure 1.4a, the vertical red dashed line indicates level-1 guesses. Both density plots in the figure show peaks around the red vertical line, which indicate higher proportions of level-1 (or close to level 1) strategy used within the games across all the subjects. Notably, in the density plot for the [LH-] treatment (G7), there is another peak centered right at the level 2 guess for that game (indicated by blue dashed line). The density plot clearly shows that in the game where subjects have a lower cost of reasoning ([LH-]), guesses are congregated at both levels 1 and 2, whereas in the game where subjects have a higher cost of reasoning ([HH-]), only a peak at the level-1 guess is observed. Likewise, in Figure 1.4b, four games are plotted to illustrate the treatment effects of increasing cost of reasoning for the opponent. In Figure 1.4c, three games are used to demonstrate changing second-order beliefs. Note that both games 1 and 8 are relevant in both graphs, as the [LL+] treatment is relevant for both comparisons. As illustrated in the figure, in one of the games, the three peaks correspond to level 1, level 2, and equilibrium. When increasing the cost of reasoning for the opponent, the level 1 peak is still observable; however, only one game has a level-2 peak. Similarly, when changing the second-order belief from low load with probability 1 to (0.5, 0.5; L, H), only the level 1 peak remains, as then the subjects thought that their opponents thought there was a 50% probability that the subject was experiencing a high cognitive load. I omitted other vertical lines that indicated different levels due to the absence of peaks in the density plots.

1.5.2 Distribution of Levels

From the preview of results from raw guesses in the previous subsection, changing the strategic environment appeared to lead to some structured changes in the depth of reasoning. However, only about half of the guesses were type-exact guesses. To better understand the treatment effects of the other half, I used maximum likelihood estimation to assign types, and then conducted analyses based on the estimated levels.

There were a total 1998 observations of guesses. As discussed in the previous section, I assigned a behavioral level for each observation. Surprisingly, a few guesses corresponded to exact level 4 and level 5 guesses in my data. Therefore, I included levels 1 to 5 and the Nash equilibrium type in my estimation. Of all the observations, 1167 guesses were estimated. The distribution of estimated levels for these guesses is shown in Table 1.5.

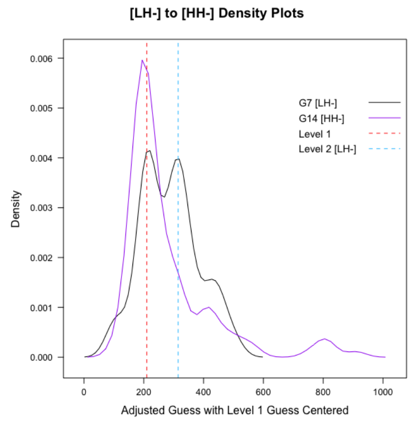

(a) Density plot of changing cost of reasoning (opponent high load)

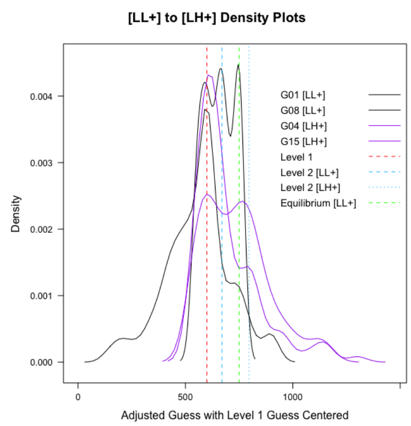

(b) Density plot of changing cost of reasoning for opponent (subject low load)

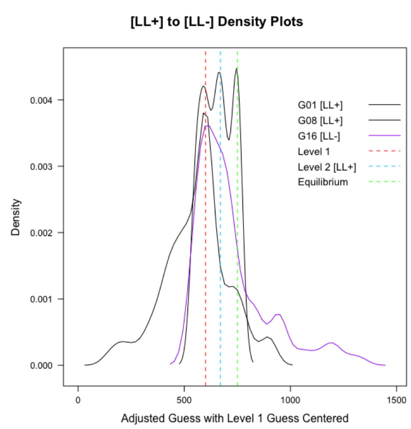

(c) Density plot of changing second-order belief (LL)

Figure 1.4: Density plot of raw guesses.

The majority of the guesses were assigned to level 1 guesses. The level distribution for all the guesses is shown in Table 1.6. The game number is referred to the game number list in Table 1.1. Since all the subjects played each game exactly once, for each game listed, there were 111 observations.

Table 1.5: Summary of estimation results.

	L1	**L2**	**L3**	**L4**	**L5**	**Nash**
Exact	43.92%	31.41%	5.54%	1.56%	3.37%	14.20%
Estimated	71.89%	5.57%	12.94%	5.57%	3.43%	0.60%
Total	60.26%	16.32%	9.86%	3.90%	3.40%	6.26%

Table 1.6: The frequency of levels by game.

Game #	**L1**	**L2**	**L3**	**L4**	**L5**	**Nash**
All Guesses	**60.26%**	**16.32%**	**9.86%**	**3.90%**	**3.40%**	**6.26%**
1	72.97%	13.51%	0	2.70%	4.50%	6.31%
2	35.14%	16.22%	19.82%	15.32%	3.60%	9.91%
3	44.14%	22.52%	33.33%	0	0	0
4	63.06%	14.41%	1.80%	7.21%	13.51%	0
5	72.07%	12.61%	0	1.80%	0	13.51%
6	83.78%	4.50%	2.70%	9.01%	0	0
7	53.15%	12.61%	16.22%	5.41%	6.31%	6.31%
8	62.16%	12.61%	1.00%	0	0	24.32%
9	59.46%	6.31%	18.02%	1.80%	7.21%	7.21%
10	54.05%	45.95%	0	0	0	0
11	69.37%	10.81%	7.21%	1.80%	1.80%	9.01%
12	44.14%	23.42%	32.43%	0	0	0
13	70.27%	16.22%	1.00%	1.00%	5.41%	6.31%
14	73.87%	11.71%	14.41%	0	0	0
15	65.77%	15.32%	11.71%	7.21%	0	0
16	64.86%	16.22%	5.41%	2.70%	10.81%	0
17	34.23%	28.83%	3.60%	10.81%	2.70%	19.92%
18	62.16%	9.91%	9.91%	3.60%	5.41%	9.91%

The distributions of the levels were fairly similar to the results in CGC06, except that levels 4 and 5 were then included. Level 1 was the most prominent behavioral level. Of 1998 observations, 60.26% were level 1 guesses. In some games, level 1 was even more frequently observed. For example, in game 1, about 70% of the guesses were classified as level 1. A number of observations were levels 2 and 3 and Nash guesses. In my data, the occurrence of level 3 was more frequent in a few games. For example, in game 2 and game 3, more than 20% of observations were assigned to level 3. Although some observations corresponded to exact level 4 or level 5 guesses, the overall frequency of these two higher levels was much lower. In about one-third of the games, no guesses were classified into these two levels.

As shown in Table 1.6, there are a pair of games that have almost identical level distribution, game 3 and game 12. These two games have identical parameters and treatments (as shown in Table 1.1). Besides these two games, the frequency of levels in other games differed considerably. In some games, behavioral levels congregated toward levels 1 or 2, for example, games 1 and 6. In some games, such as games 2 and 9, behavioral levels spread out across the six categories. The variations in the distribution of levels across games could be due to the differences in the cognitive load tasks. The exact impact of the memorization tasks is discussed in detail in the following subsections.

1.5.3 Result 1: Increasing Cost of Reasoning

As mentioned in Section 1.2.2, the first testable prediction involved fixing the subject's first- and second-order beliefs and examining the effect of the changing cost of reasoning on the subject's behavioral levels. There were essentially two comparisons in this case: a comparison between treatment [LL-] and treatment [HL-], and between treatment

[LH-] and [HH-]. Note that in both comparisons, the cost of reasoning for the subject varied from low to high; therefore, it was crucial to have partial revelation of the subject's (role A) memorization task. In the partial revelation treatment, role B (the opponent) only knew the probability distribution of the subject's memorization task (0.5, 0.5; L, H); therefore, even with the subject's own tasks varying between two treatments, the subject's second-order belief was controlled to be the same. There were 222 pairs of comparison in total. The summary statistics of the comparisons are presented in Table 1.7. The plotted distribution of behavioral levels is presented in Figure 1.5. To aid with the interpretation of the results, the behavioral levels in the figure are presented in a reverse order (i.e., higher level on the left and lower level on the right).

Table 1.7: The frequency of changing behavioral levels with increasing cost of reasoning.

Pair Name (From Game a to Game b)	# of Pairs	Treatment	Decreases	Constant	Increases
G16 to G9	111	LL- to HL-	20.72%	43.24%	36.03%
G7 to G14	111	LH- to HH-	39.64%	49.55%	10.81%
Combined	222	L?- to H?-	30.18%	46.40%	23.42%

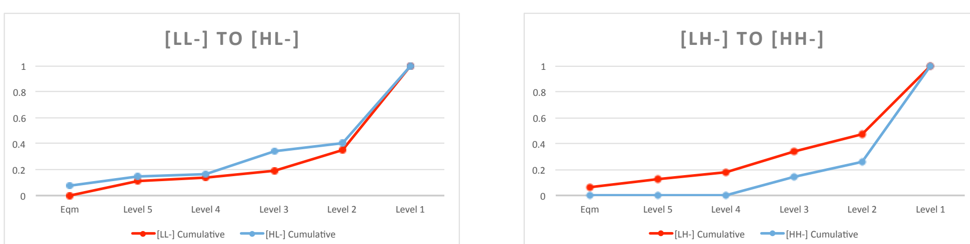

(a) **(b)**

Figure 1.5: Cumulative level distribution for increasing cost of reasoning.

When the opponent's cognitive load was controlled to be high and with partial revelation, subjects weakly decreased their behavioral levels 89.19% of the time (39.64% strict decrease). In Figure 1.5b, the [LH-] treatment is first-order stochastic dominant over the [HH-] treatment. The Wilcoxon test (Table 1.8) was significant at the 1% level for the comparison of the distributions of behavioral levels between these two strategic environments. When regressing the behavioral level on the treatment dummy, the result (Table 1.9) suggested that the coefficient for treatment dummy was 0.77, which was significant at the 1% level. This implied that the estimated behavioral level weakly decreased when the subject's own cognitive load increased when facing an opponent with high cognitive load. The finding is consistent with the EDR model. The relatively large proportion (49.55%) of constant levels may seem quite surprising at first look. One possible explanation is that these subjects may have had different cognitive bounds in the two treatments. In the [LH-] treatment, subjects may have adjusted their behavioral levels downward from their cognitive bound in that treatment due to some belief they formed when facing opponents with high cognitive loads. In the [HH-] treatment, subjects who had a lower cognitive bound (as they had a high cognitive load) may have displayed a lower behavioral level. When the two behavioral levels from two treatments coincided, I observed no changes in the behavioral levels in the treatment comparison.

The result for the comparison between [LL-] and [HL-] is less clear. As shown in Table 1.7, 63.97% of the comparisons had weakly decreasing behavioral levels (20.72% strict decrease), and a noticeable percentage (36.03%) of the comparisons had increasing levels. The Wilcoxon test statistic rejected the null hypothesis that the two strategic environments have the same distribution of behavioral levels at the 5% level. However, upon further checking using a one-tail Wilcoxon test, the distribution of behavioral levels shifted rightward when cognitive load changed from low to high when facing an opponent with a low cognitive load. When conducting the standard GLS random effect regression, the coefficient on the treatment dummy was positive and significant at the 10% level.

Table 1.8: Test results for equality of distribution.

Comparison Group (in Treatment)	**Wilcoxon *p*-Values (Two-Tailed)**	**Wilcoxon *p*-Values (One-Tailed)**
Changing Cost of Reasoning		
LL- to HL-	0.05 **	0.98
LH- to HH-	0.00 ***	0.00 ***
Changing Opponent's Cost of Reasoning		
LL+ to LH+	0.01 ***	0.00 ***
HL+ to HH+	0.00 ***	1
Changing Second Order Belief		
LL+ to LL-	0.11	0.05 **
LH+ to LH-	0.00 ***	0.99
HL- to HL+	0.00 ***	0.00 ***
HH- to HH+	0.00 ***	1
Against Computer (Nash)		
L to H	0.00 ***	0.00 ***

Notes: * indicates $<$ 10% significance, ** indicates $<$ 5% significance, and *** indicates $<$ 1% significance.

Table 1.9: Regression results for treatment effects.

Comparison Group (in Treatment)	Relevant Dummy	Constant	Number of Obs.
Changing Cost of Reasoning			
LL- to HL-	0.34 * (0.18)	1.78 *** (0.14)	222
LH- to HH-	-0.77 *** (0.16)	2.18 *** (0.12)	222
Changing Opponent's Cost of Reasoning			
LL+ to LH+	-0.27 * (0.14)	2.03 *** (0.11)	444
HL+ to HH+	0.33 *** (0.10)	1.55 *** (0.07)	444
Changing Second Order Belief			
LL+ to LL-	-0.25 (0.18)	2.04 *** (0.12)	333
LH+ to LH-	0.41 *** (0.15)	1.77 *** (0.09)	333
HL+ to HL-	0.57 *** (0.15)	1.55 *** (0.10)	333
HH+ to HH-	-0.48 *** (0.09)	1.89 *** (0.06)	333

Notes: * indicates $< 10\%$ significance, ** indicates $< 5\%$ significance, and *** indicates $< 1\%$ significance. Standard errors in parenthesis.

In this analysis, I treated equilibrium level as the highest level, since it requires the subjects to perform multiple steps of iterative best responses. However, since many games have equilibrium at the boundary of the limit interval (games 2, 4, 6, 10, 15, and 16 have the equilibrium at the lower limit; games 3, 5, 8, and 12 have the equilibrium at the upper limit), if the subject chooses an equilibrium action by naïvely playing at the boundary, then this behavioral level should not be considered as a higher level than any of the *k* levels. This was not the case for this comparison pair. Although game 9 ([HL-]) had 7.21% equilibrium guesses, those guesses were not at the boundary. However, upon further checking of games 16 ([LL-]) and 9 ([HL-]), I found that the level-5 type in game 16 had the same strategy as the equilibrium strategy of that game. Therefore, some of the equilibrium strategies in game 16 were pooled into level-5 type, which may be one possible explanation for the significant positive coefficient on the treatment dummy. Another explanation may be that the subjects felt more motivated to reason at higher strategic levels when they saw the opponents had easier strategic environments (memorizing three letters) as opposed to their own difficult strategic environments (memorizing seven letters). As a result, they displayed higher behavioral levels. This explanation suggests that other factors, such as motivation factor, may also play a role in determining a subject's behavioral levels.

1.5.4 Result 2: Increasing Cost of Reasoning for Opponent

To examine the effect of changing the first-order belief on a subject's behavioral level in games, I selected pairs of games with changing cognitive loads for the opponents. For example, a comparison of behavioral levels for games 1 and 4 served the purpose. In game 1 ([LL+]), player 1 has a low cognitive load when facing an opponent with a low cognitive load, and there is full revelation of each other's strategic environment. In game 4 ([LH+]), player 1 has a low cognitive load when facing an opponent with high cognitive load, and again, there is full revelation of the treatments. I found 444 pairs of comparison for the cases wherein the subjects had low cognitive loads, and another 444 pairs of comparison for the cases when they had high cognitive loads. The detailed comparison groups and summary statistics are shown in Table 1.10. The plotted distribution of behavioral levels is presented in Figure 1.6.

Table 1.10: The frequency of changing behavioral levels with increasing cost of reasoning for opponent.

Pair Name (From Game a to Game b)	# of Pairs	Treatment	Decreases	Constant	Increases
(G1, G8) to (G4, G15)	444	LL+ to LH+	23.87%	55.86%	20.27%
(G6, G13) to (G3, G12)	444	HL+ to HH+	15.54%	41.44%	43.02%
Combined	888	?L+ to ?H+	19.71%	48.65%	31.64%

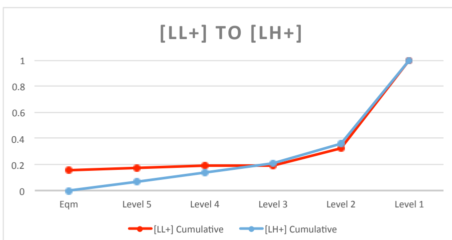

(a)

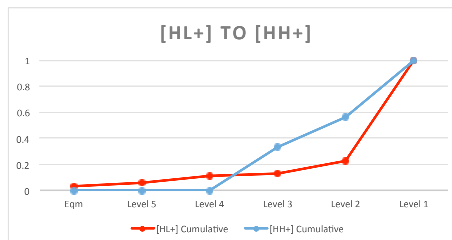

(b)

Figure 1.6: Level Distribution for increasing cost of reasoning for opponent.

The combined results were the opposite of the theory prediction, with a significant 31.64% of cases of increasing behavioral levels. However, upon further checking, the majority of the increasing cases occurred when subjects are having high cognitive load. When subjects had low cognitive load, 79.73% of the time, they weakly decreased their behavioral levels when their opponents' cognitive loads changed from low to high (23.87% strict decrease). Figure 1.6a illustrates that [LL+] games had more guesses at higher levels. This result is consistent with the EDR model. When a subject's cost and second-order belief was controlled across the two strategic environments, he was responsive to the changes in his opponent's cost of reasoning. However, some of these adjustments in behavioral levels were not strictly decreasing. If the subject believed that the increased opponent's cost of reasoning was not large enough to decrease the opponent's behavioral level by one, the subject's behavioral level remained the same across the two strategic environments. This partially explains the high percentage (55.86% and 41.44%) of constant behavioral levels in Table 1.10. When the subject had a high cognitive load and his opponent's cognitive load changed, the result did not comply with the EDR model. A total of 43.02% of the pairs showed increasing behavioral levels across the two strategic environments. The frequency of levels in Table 1.6 reveals that most subjects had level 1 guesses in games 6 (83.78%) and 13 (70.27%). This gave subjects much less room to adjust their behavioral levels downward compared to another strategic situation. Any behavioral level that was beyond level 1 in games 3 and 12 was considered as moving the behavioral level upward. This was one major limitation in observing the effects of changing the first-order belief when the subject had a high cost of reasoning (i.e., high cognitive load).

The Wilcoxon signed-rank test rejected the null hypothesis that the level distribution was the same for both treatment comparisons ([LL+] to [LH+] and [HL+] to [HH+]). However, the one-tail test suggested that when the subject had low cognitive load, increasing his opponent's cost of reasoning shifted the former's level to the left (to lower levels, significant at the 1% level). However, when the subject had high cognitive load, the level distribution shifted to the right. The regression coefficients suggested that increasing the opponent's cost of reasoning decreased the behavioral level when the subject had a low cognitive load (significant at the 10% level).

1.5.5 Result 3: Changing the Second-Order Belief

In the experiment, I used a (0.5, 0.5) probability distribution on the revelation of cognitive load treatments to control for the subject's second-order belief. In the full revelation treatment, role B knew the exact memorization task that was received by role A (the subject), either three (low load) or seven letters (high load) with a probability of one. Therefore, role A's (the subject) second-order belief was either ((1, 0); (L, H)) or ((0, 1); (L, H)). In the partial revelation treatment, role B knew that the probability of three or seven letters for role A was (0.5, 0.5), which made role A (the subject) have a second-order belief of ((0.5, 0.5); (L, H)). If comparing two games with different second-order beliefs for the subject, with everything else controlled as constant, then a second-order belief of low load with probability of one should be considered as more cognitively capable perceived by role B than a second-order belief of((0.5, 0.5); (L, H). The experiment, as shown in Table 1.11, supported that most subjects had a clear understanding of their opponent's cognitive load when the load was explicitly elicited, and they almost had uniform beliefs about their opponents' cognitive loads when they were in the partial revelation treatment as role B.

Table 1.11: Subject's belief about his opponent's cognitive load.

		Belief Elicitation			
		3 Letters	**7 Letters**	**Not Sure**	**Sum**
	3 Letters	591	51	24	666
Treatments	7 Letters	97	546	23	666
	(0.5 L, 0.5 H)	143	111	190	444

In the dataset, I found 888 pairs for comparison that allowed me to examine the effect of changing the second-order belief. I separated them into two groups: a comparison between the full revelation of low load to partial revelation, and a comparison between a partial revelation and a full revelation of high load. Both comparisons were performed in the direction of increasing second-order belief (i.e., c_i^{ij} increases). The detailed comparison pairs and summary statistics are listed in Table 1.12. The distribution of behavioral levels is plotted in Figure 1.7.

Table 1.12: The frequency of changing levels with changing second-order belief.

Pair Name (From Game a to Game b)	**# of Pairs**	**Treatment**	**Decreases**	**Constant**	**Increases**
Second order belief: Low to (0.5 Low, 0.5 High)					
(G1, G8) to G16	222	LL+ to LL-	25.68%	51.35%	22.97%
(G4, G15) to G7	222	LH+ to LH-	20.27%	43.69%	36.04%
Combined	444	L?+ to L?-	22.97%	47.52%	29.50%
Second order belief: (0.5 Low, 0.5 High) to High					
G9 to (G6, G13)	222	HL- to HL+	35.14%	52.70%	12.16%
G14 to (G3, G12)	222	HH- to HH+	14.86%	45.50%	39.64%
Combined	444	H?- to H?+	25.00%	49.10%	25.90%

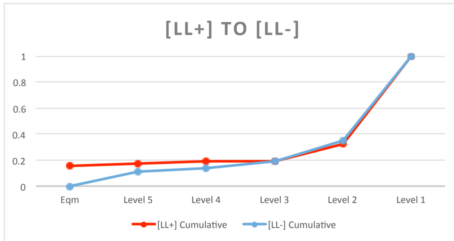

(a)

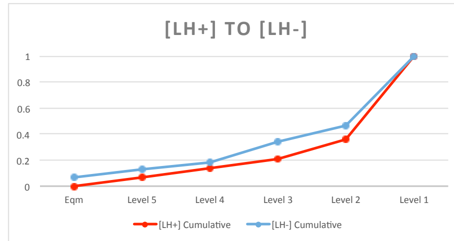

(b)

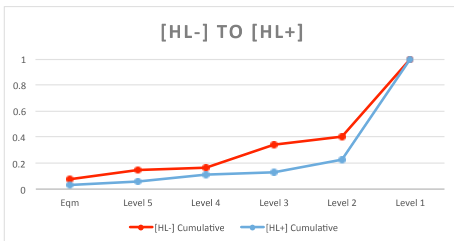

(c)

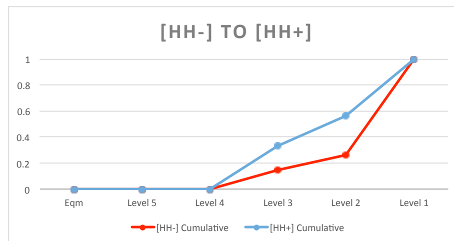

(d)

Figure 1.7: Level Distribution for changing second-order belief.

The effect of changing the second-order belief was generally weak, except for the cases where the subjects had high cognitive loads when facing opponents with low cognitive loads. For the treatment where both players had low cognitive loads, about 77.03% of the pairs had weakly decreasing behavioral levels when second-order belief changed from full to partial revelation. Among these comparisons, only 25.68% had strictly decreasing levels. This finding suggested that the changes in second-order belief may not have been strong enough for the subjects to adjust their behavioral level downward, even though both subjects had a low cognitive load and were relatively competent at contemplating over the strategic environment. To examine the effect of the second-order belief, it was first necessary to determine the effect of changing the first-order belief for the same group of subjects. In Table 1.10, the subject's behavioral responses to the changing opponent's cognitive load were limited when the subject had a high cognitive load. Now, consider the finding in the [LH+] to [LH-] comparison to the [HH-] to [HH+] comparison (Table 1.12); changing the second-order belief of the subject effectively changed his opponent's first-order belief. If the subject holds the belief about his opponent (who has a high cognitive load treatment) that the changes in his opponent's behavioral level are limited, then the subject should not decrease his behavioral level at all. This partially explains the low frequency of strictly decreasing behavioral levels for subjects who faced opponents with high cognitive loads.

The comparison between [HL-] and [HL+] is consistent with the EDR model. In Figure 1.7c, the [HL-] treatment is first-order stochastic dominant over the [HL+] treatment. Of the guesses, 87.84% had weakly decreasing behavioral levels, with 35.14% having strict decreases. Changing from partial revelation to full revelation of high cognitive load, the second-order belief decreased the subject's cognitive capability perceived by their opponent. Subjects were responsive to this change in the belief system, and adjusted their behavioral levels downward to best respond to their opponents. Testable prediction 3 suggests that if the subject's behavioral level is binding by their cognitive bound, then they are not able to make further adjustments according to their changing beliefs. The large percentage of constant levels for these comparisons supported this statement.

The Wilcoxon test results showed that the level distribution changed for changing second-order belief. When conducting a one-tailed test, the test result suggested that for [LH+ to LH-] and [HH- to HH+] treatments, the distribution of levels significantly (at the 1% level) shifted rightward (increasing behavioral levels). This may have occurred due to the subject's belief that their opponent with high cognitive load will engage in higher behavioral level. This result seems to comply with the results in Section 1.5.4, but the underlying reasons need further investigation.

The regression coefficient on the treatment dummy further supported the results. Since the treatment dummy was coded as zero with full revelation and one with partial revelation, the coefficient of 0.57 for [HL-, HL+] comparison suggested that the behavioral level decreased from partial to full revelation. It was significant at the 1% level. Again, the [LH+, LH-] and [HH-, HH+] comparisons were the opposite direction of model predictions, and they were also highly significant. In general, when the subjects faced opponents with high cognitive loads, they were responsive to changing second-order beliefs, but not in the direction that is predicted by the EDR model. However, when they faced more cognitively capable opponents, then they were mostly responsive to this change in the belief system because they thought their opponents were responsive to this information in their strategic environment. This finding is consistent with the EDR model when the opponent has a low cognitive load, which supports the opposite direction when the opponent is in a less cognitively capable situation.

1.5.6 Result 4: Cognitive Bound

In block 2 of the experiment, the subjects played against the computer. They were told that the computer was playing a Nash equilibrium strategy, and the equilibrium concept was explained. However, they were not taught the method to derive the equilibrium. The behavioral levels from the guesses in these two games should be considered as the highest levels they could achieve under each cognitive load treatment. I selected all the games with the same cognitive load treatment, either low cognitive load or high cognitive load, and pooled the results. A pairwise comparison between the pooled data and behavioral level obtained from games 17 and 18 allowed me to examine the existence of cognitive bounds. There were 888 pairs of comparison for each type of cognitive load, and the summary statistics are shown in Table 1.13.

Table 1.13: The frequency of changing behavioral levels comparing to cognitive bound.

Cognitive Bound	**# of Pairs**	**Decreases**	**Constant**	**Increases**
Low cognitive bound				
(G1, G2, G4, G7, G8, G11, G15, G16) to G17	888	18.36%	33.45%	48.20%
High cognitive bound				
(G3, G5, G6, G9, G10, G12, G13, G14) to G18	888	21.96%	48.20%	29.84%
Combined	1776	20.16%	40.82%	39.02%

The result for low cognitive load treatment was interesting: 48.20% of the guesses from block 1 games had behavioral levels lower than the subject's respective cognitive bound (level in game 17). Less than 20% of guesses had higher behavioral levels. This suggested that in many block 1 games, subjects purposely adjusted their behavioral levels downward due to different strategic situations, even though they had reached higher levels. For high cognitive load treatment, about 30% of behavioral levels increased from block 1 games to game 18. However, about 50% of the guesses had the same behavioral level across the two situations. Since high cognitive load inherently restricts the subject's cognitive ability, there may have been less room for downward adjustments for block 1 games. Due to the large percentage of weakly increasing levels from block 1 to block 2 games, I concluded that cognitive bound existed in most cases. In some situations, cognitive bound was strictly higher than the subject's behavioral levels in games. In some situations, cognitive bound was the same as the behavioral levels. Such cases were largely observed in the high cognitive load treatment.

To examine whether high cognitive load had a lower level distribution, I conducted a Wilcoxon signed-rank test on the estimated behavioral levels of games 17 (low load) and 18 (high load). Table 1.8 shows that the distributions of levels for the two treatments were significantly different at the 1% level. The one-tailed test indicated that the distribution of low load cognitive bound levels was to the right of the distribution of high load cognitive bound levels. This finding indicated that subjects had a higher cognitive bound when receiving low cognitive load treatment (memorizing a string of three letters) compared to receiving a high cognitive load treatment (memorizing a string of seven letters).

1.5.7 Robustness Check

During the guessing games, subjects needed to memorize a string of three or seven letters and recall the letters after they finished the guessing game. In this subsection, I present the results of this memorization task. Although the subjects were fully aware that if they failed to recall all the letters correctly, they would earn zero points for that round of the game, there were still some cases of wrong recalls due to reasons such as lack of attention or being too focused on the guessing game. I wanted to control the experimental results for such cases, as the subjects may have engaged in reasoning at higher levels when cognitive load did not fully apply. Table 1.14 shows the results of the memorization tasks. Most of the memorization tasks were perfectly performed. Not surprisingly, low cognitive load (three-letter memorization task) had more correct recalls, about 7% more than the high cognitive load task. The difference was significant at the 1% level.

Table 1.14: Results of the memorization task.

	3 Letters *** (Low Load)	**7 Letters (High Load)**	**Total**
Correct	97.30%	90.89%	94.09%
Wrong	2.70%	9.11%	5.91%
# of tasks	999	999	1998

Notes: * indicates $<$ 10% significance, ** indicates $<$ 5% significance, and *** indicates $<$ 1% significance.

To check whether poor performance of the memorization task affected the treatment results, I excluded the data with wrong recalls and performed the analysis again. The comparison pair was dropped from the sample if either game of the pair had incorrect recalls. This was performed to ensure that the cognitive load was fully in effect, so that high cognitive load added difficulties to thinking through the guessing games at higher levels, and the cost of reasoning was higher.

Table 1.15 presents the treatment results after the robustness check. Treatments that involved high cognitive load had more data points dropped. For example, the [HH-] to [HH+] comparison had 444 pairs of comparison in the original sample. After robustness check, about 100 pairs were dropped. However, the results did not change much compared to the results presented in results 1 to 3 (Sections 1.5.3–1.5.5). The changes were mostly within 1%. I can therefore safely conclude that the original results were robust. The quality of the memorization task (i.e., whether the letters were correctly recalled) was almost independent of the treatment effects. Even in the cases of wrong recalls, the effect of cognitive load still applied to the subjects.

Table 1.15: Summary of the robust results with incorrect recalls dropped.

Pair Name (From Game a to Game b)	# of Pairs	Treatment	Decreases	Constant	Increases
Increasing Cost of Reasoning					
G16 to G9	107	LL- to HL-	20.56%	43.93%	35.51%
G7 to G14	77	LH- to HH-	40.26%	48.05%	11.69%
Combined	184	L?- to H?-	28.80%	45.65%	25.54%
Increasing Cost of Reasoning for Opponent					
(G1, G8) to (G4, G15)	426	LL+ to LH+	23.71%	56.10%	20.19%
(G6, G13) to (G3, G12)	392	HL+ to HH+	16.33%	42.09%	41.58%
Combined	818	?L+ to ?H+	20.17%	49.39%	30.44%
Second order belief: Low to (0.5 Low, 0.5 High)					
(G1, G8) to G16	214	LL+ to LL-	26.17%	50.93%	22.90%
(G4, G15) to G7	211	LH+ to LH-	21.33%	44.08%	34.60%
Combined	425	L?+ to L?-	23.76%	47.53%	28.71%
Second order belief: (0.5 Low, 0.5 High) to High					
G9 to (G6, G13)	202	HL- to HL+	35.64%	52.48%	11.88%
G14 to (G3, G12)	152	HH- to HH+	15.13%	45.39%	39.47%
Combined	354	H?- to H?+	26.84%	49.44%	23.73%

1.5.8 Cognitive Tests

In this subsection, I examine the results of the Mensa practice test. The test is composed of 10 questions and has a time limit of 10 min. Some subjects finished earlier, but they could never run overtime. Each correct answer is worth 1 point and all the unattempted questions are marked as 0 points. The score distribution of 104 subjects (seven missing) is presented in Table 1.16. There are a few very low points (2 or 3), and six subjects had scores of 10. Most subjects earned seven or eight points in this test.

Table 1.16: Summary statistics of cognitive test score and the counts of level changes following theory predictions.

	(1)	(2)	(3)	(4)	(5)	(6)
	Test Score	**Sum.Strict**	**Sum.Weak**	**Cost.Weak**	**1st.Weak**	**2nd.Weak**
Points possible	10	18	18	2	8	8
Max	10	13	18	2	8	8
Min	2	0	7	0	2	1
Median	7	4	13	2	6	6
Mean	7	4	13	1.5	5.5	6

To examine whether there are heterogeneous treatment effects in this experiment due to exogenous cognitive ability, I first determined a measure of the treatment effect. Out of all the results discussed in results 1 to 3 (Sections 1.5.3–1.5.5), there are in total 18 pairs of comparison. For each subject, I recorded one for the pair if the level change followed the theory prediction, and zero otherwise. As listed in Table 1.16, column Sum.Strict includes all the 18 comparisons, and only strict changes of levels are recognized. For example, if the pair game 16–game 9 had level 2 in both games, it is coded zero under Sum.Strict. However, column Sum.Weak allows weak changes; therefore, the above-mentioned scenario is coded as one under this column. The EDR model mostly discusses weak behavioral level changes because, in some cases, the changes in belief system or costs are not big enough to shift a behavioral level downwards by one level (evidenced by a large percentage of constant levels). Due to this reason, I considered ”weak” changes, and decomposed them into columns (4) to (6), which cover the three main results. When limited to strict changes, a number of subjects had zero pairs following theory prediction (10 out of 111 subjects), and most subjects had only three or four pairs that had changes that could be predicted by the EDR model. However, when allowing weak changes, seven subjects had all the comparison pairs that were theory-predicted directional level changes, and most subjects had about 13 to 14 comparisons that could be predicted by the EDR model. The last three columns in Table 1.16 present results for each treatment separately.

To test whether cognitive ability had any correlation with the treatment effects, I ran a regression after dropping the subjects with missing test scores. The result is presented in Table 1.17. I used gender, class standing, and major as control variables. This information was collected at the end of the experiment. It appears that the cognitive test score and the female dummy variable were positively correlated with weak changes (at a 5% significance level), and the treatments changing the opponent's cost of reasoning (changing first-order belief) and changing second-order belief. The results showed some heterogeneous treatment effects in which the more cognitively capable subjects were more responsive to the treatments as predicted by the EDR model, especially in those requiring adjustments in response to the changing strategic environment of their opponents. When the strategic environment changed, these subjects were more likely to actively adjust their actions to gain possible strategic advantages.

Table 1.17: Regression results for cognitive test scores on correct directional changes of behavioral levels in block 1 games.

	Sum.Strict	**Sum.Weak**	**Cost.Weak**	**First Order.Weak**	**Second Order.Weak**
Test Score	0.03	0.36 **	-0.04	0.23 **	0.17*
	(0.16)	(0.18)	(0.04)	(0.11)	(0.10)
Gender (F)	-0.73	1.43 **	0.03	0.83 **	0.57 *
	(0.57)	(0.60)	(0.13)	(0.37)	(0.34)
Class Standing	0.33	-0.18	-0.01	-0.13	-0.04
	(0.24)	(0.25)	(0.05)	(0.15)	(0.14)
Major	-0.02	0.00	-0.03	-0.03	0.01
	(0.12)	(0.13)	(0.03)	(0.08)	(0.07)
Constant	3.09 *	10.04 ***	1.76 ***	3.94 ***	4.34 ***
	(1.67)	(1.75)	(0.37)	(1.08)	(1.00)
# of Obs.	104	104	104	104	104

Notes: ”Weak” includes constant levels and decreasing levels, while ”strict” only includes strictly decreasing levels.* indicates $< 10\%$ significance, ** indicates $< 5\%$ significance, and *** indicates $< 1\%$ significance. Standard errors in parenthesis.

Since the result above suggested that more cognitively capable subjects' responses to changing strategic environment were more coherent with the EDR model, I separated the subjects into two groups according to cognitive test scores. Subjects with scores of eight or above were labeled as high cognitive subjects (high), and the remainder were labeled as low cognitive subjects (low). Table 1.18 presents results 1–3 again, separated by the cognitive test scores. As discussed in results 1 to 3 (Sections 1.5.3–1.5.5), I found significant asymmetries arising from the different strategic environments. Separating the subjects into two groups according to cognitive test scores allowed a closer examination of the source of the asymmetry. In Table 1.18, result 2 and result 3.1 highlight the relatively stable performance for the high cognitive subjects. As discussed in Section 1.5.4, subjects' responses to their opponents' changing cost of reasoning depended on their own cost of reasoning. In general, their adjustments in behavioral levels only followed the EDR model when they had a low cost of reasoning. This observation is untrue for the high cognitive subjects, who showed relatively stable performance regardless of their own strategic environment, with about 20% of the comparisons strictly following the EDR model. I observed a slight increase of 10% for those that did not follow the model; however, in general, the performance did not vary considerably. For the low cognitive subjects, the difference was huge. The 27.54% for comparison pairs that strictly followed the model decreased to 12.29%, and, more strikingly, the percentage of pairs that did not follow the model increased from 19.07% to 51.27%. This huge difference showed that the asymmetry found in the previous results was mostly due to these low cognitive subjects. There was a similar observation for result 3.1, where the high cognitive subjects had relatively stable performance regardless of their opponents' cognitive loads, whereas the low cognitive test score subjects were very sensitive to their opponents' strategic environments. Therefore, I concluded that the majority of asymmetric results found in results 2 and 3.1 were primarily driven by the low cognitive subjects. They were responsive to the treatments under the condition that they were in a more cognitively advanced situation. For results 1 and 3.2, both high and low cognitive subjects responded asymmetrically toward the treatment. However, as evidenced in Table 1.18, the changes from the low cognitive group were much greater than those of their counterparts.

Table 1.18: Results 1 to 3 separated by cognitive test scores.

			Changes in Behavioral Levels		
Treatment	**Cog Test**	**# of Pairs**	**Decreases**	**Constant**	**Increases**
Result 1: Increasing Cost of Reasoning					
LL- to HL-	High	45	22.22%	37.78%	40.00%
LL- to HL-	Low	59	16.95%	47.46%	35.59%
LH- to HH-	High	45	28.89%	55.56%	15.56%
LH- to HH-	Low	59	35.59%	57.63%	6.78%
Result 2: Increasing Cost of Reasoning for Opponent					
LL+ to LH+	High	180	19.44%	59.44%	21.11%
LL+ to LH+	Low	236	27.54%	53.39%	19.07%
HL+ to HH+	High	180	20.00%	47.22%	32.78%
HL+ to HH+	Low	236	12.29%	36.44%	51.27%
Result 3.1: Second order belief: Low to (0.5 Low, 0.5 High)					
LL+ to LL-	High	90	23.33%	53.33%	23.33%
LL+ to LL-	Low	118	28.81%	50.85%	20.34%
LH+ to LH-	High	90	20.00%	44.44%	35.56%
LH+ to LH-	Low	118	19.49%	42.37%	38.14%
Result 3.2: Second order belief: (0.5 Low, 0.5 High) to High					
HL- to HL+	High	90	35.56%	48.89%	15.56%
HL- to HL+	High	118	34.75%	55.08%	10.17%
HH- to HH+	High	90	21.11%	50.00%	28.89%
HH- to HH+	High	118	11.02%	40.68%	48.31%

The impact of cognitive ability on treatment effects was further evidenced by the regression results. In Table 1.19, the interaction term is significant for the comparison pairs that did not follow the EDR model ([HL+ to HH+], [LH+ to LH-], and [HH+ to HH-]). This implied that higher cognitive test scores skewed the effects of the treatment in the direction pointed by the EDR model. It seems that cognitive ability plays an important role for the subjects to display behavioral changes that can be predicted by the EDR model. The cognitive ability was captured endogenously by the treatment design in this experiment with two kinds of cognitive load. As discussed previously, the results differed systematically according to the amount of cognitive resources. Cognitive ability was also captured exogenously by the Mensa practice test, as discussed in this section. Within the asymmetric findings, subjects with higher cognitive test scores had more stable performance regardless of their own cognitive load, and were generally more predictable by the EDR model.

Table 1.19: Regression results for treatment effects and cognitive test scores on behavioral levels.

Comparison	Dummy	Score	Dummy $*$ Score	# Obs.
Changing Cost of Reasoning				
LL- to HL-	0.81 (0.79)	0.07 (0.07)	-0.06 (0.11)	208
LH- to HH-	-2.13^{***} (0.73)	-0.11 (0.10)	0.19^{*} (0.10)	208
Changing Opponent's Cost of Reasoning				
LL+ to LH+	-1.15 (0.77)	-0.04 (0.09)	0.12 (0.10)	416
HL+ to HH+	1.67^{***} (0.44)	0.08 (0.05)	-0.19^{***} (0.06)	416
Changing Second Order Belief				
LL+ to LL-	-1.06 (0.75)	-0.04 (0.09)	0.11 (0.10)	312
LH+ to LH-	1.77^{**} (0.80)	0.08 (0.05)	-0.19^{*} (0.11)	312
HL+ to HL-	1.08 (0.75)	0.08 (0.05)	-0.07 (0.10)	312
HH+ to HH-	-1.85^{***} (0.37)	-0.11^{***} (0.04)	0.19^{***} (0.05)	312

Notes: * indicates $< 10\%$ significance, ** indicates $< 5\%$ significance, and *** indicates $< 1\%$ significance. Clustered individual standard errors in parenthesis.

1.6 Concluding Remarks

In this study, I designed a laboratory experiment to examine the consistency of players' strategic sophistication formulated by the level-k model. Following the endogenous depth of reasoning framework, I controlled the strategic environment by varying the cost of reasoning for the subjects, and their first- and second-order beliefs about their opponents.

My findings were consistent with the EDR model under some conditions. When the strategic environment was carefully controlled, subjects were very responsive towards the changes in the environment. Subjects who have more cognitive resources (in a low cognitive load treatment) or subjects who are facing opponents with less cognitive resources (in a high cognitive load treatment) change strategies systematically. This behavior can be predicted by the EDR model. Subjects in a strategically disadvantaged situation (high cognitive load treatment) have less room for strategic adjustments. In some of my findings, subjects appeared to try to achieve higher behavioral levels when they were under the high cognitive load treatment. The reason for this is still unclear. It may due to the awareness of the strategic disadvantage and the extra effort of the subjects under such situations, or some other behavioral factors existed that were not captured by the EDR model. The underlying reason needs further investigation. The effect of cognitive ability on the treatments was also captured by the cognitive test. Subjects with higher test scores were more predictable by the EDR model, regardless of the strategic environment. This finding is in line with the asymmetry observed in my results. As the source of asymmetry was mainly the amount of cognitive resources, it is not surprising that subjects with higher cognitive test scores adjusted better in these tasks.

A level of cognitive bound existed for subjects in different strategic situations. When playing games under the same amount of cognitive resources, subjects rarely had behavioral levels that exceeded their respective cognitive bounds for that strategic situation. Significant downward adjustments occurred from the cognitive bound in response to different strategic environments. Overall, when there is a strict control over the strategic environment, changes in k-levels across games are systematic. They can be explained by the EDR model to some extent, especially for subjects in a more cognitively advantaged situation. This study only discusses the directional changes in the levels. Further studies could examine the criteria and accuracy of such predictions.

1.7 Acknowledgement

Chapter 1, in full, is a reprint of the material as it appears in Games, vol. 11(3), 2020. The dissertation author was the primary investigator and author of this paper.

Chapter 2

Predictive Accuracy for Measures of Higher Order Rationality - An Experiment Using Ring Games

2.1 Introduction

Understanding the extent of strategic sophistication is critical to making accurate predictions for outcomes in strategic interactions. Traditionally, analysis of strategic sophistication begins with a representation of a coherent system of beliefs and best responses. Rationalizable strategies (Bernheim, 1984; Pearce, 1984) are ones that can be arrived at through such a system. Orders of rationalizable actions, termed orders of rationality, have been used as the basis for analyzing the extent of strategic sophistication. Kneeland (2015) provides one prominent empirical example, recovering the distribution of strategic sophistication based on orders of rationalizable actions.

To cleanly identify orders of rationality, Kneeland (2015) introduces a novel experimental design called a "ring game." Within a ring game, subjects play in one of four roles with a specified interaction. Each (P)layer's payoff depends on the action of the player indexed above them: P1's payoff on P2's action, P2's payoff on P3's action, P3's payoff on P4's action; and P4's payoff on P1's action. The last interaction is largely inconsequential as P4 is given a dominant strategy, which 93% of Kneeland's subjects follow. Given virtually all P4 subjects are rational to a first order, obeying dominance, a P3 subject who is rational to a second order can only rationalize playing the corresponding best response to P4's dominant strategy. Correspondingly a P2 subject who is rational to a third order can only rationalize playing the best response to P3's best response to P4's dominant strategy; and so on. Between two game variants, G_1 and G_2, the action associated with P4's dominant strategy is altered, but all other aspects remain the same. Orders of rationality are identified by examining actions in all roles, and whether they respond to the experimental manipulation of the dominant strategy labeling for P4. Figure 2.1, Panel A reproduces the two Kneeland ring games, and Table 2.1 reproduces the k^{th}-order rationalizable strategies corresponding to the order of rationality in each role and game variant.

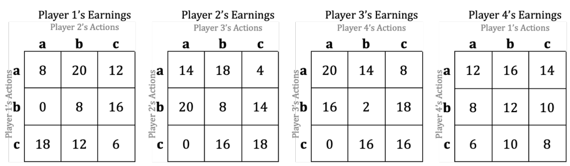

(**a**) G1

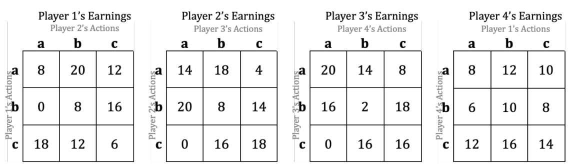

(**b**) G2

Panel A: Kneeland ring games

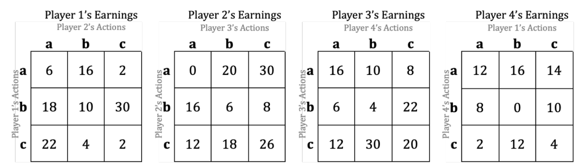

(**c**) G3

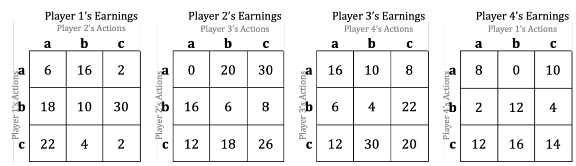

(**d**) G4

Panel B: Modified ring games

Figure 2.1: Two sets of ring games

Table 2.1: Actions and orders of rationality in Kneeland ring games

	Player and Game										
	P1			P2			P3			P4	
Order	G1		G2	G1		G2	G1		G2	G1	G2
R1	(a,a)	(b,b)	(c,c)	(a,a)	(b,b)	(c,c)	(a,a)	(b,b)	(c,c)	(a,c)	
R2	(a,a)	(b,b)	(c,c)	(a,a)	(b,b)	(c,c)		(a,b)		(a,c)	
R3	(a,a)	(b,b)	(c,c)		(b,a)			(a,b)		(a,c)	
R4		(a,c)			(b,a)			(a,b)		(a,c)	

This manuscript seeks to understand the predictive accuracy of measures of strategic sophistication derived from experimental ring game methods.[1] Given the importance of

[1]Using other methods for identifying the extent of strategic sophistication, research has reached mixed conclusions on predictive accuracy. For example, in a set of matrix games Stahl and Wilson (1995) documents consistent predictions for 73% of subjects. Using the set of two-person guessing games developed by Costa-Gomes and Crawford (2006), Zhao (2020) provides evidence of changing strategic behavior according

making accurate behavioral predictions in a range of strategic interactions, this is a critical research question. To date little is known on this question. In one initial test, Kneeland (2015) documents robustness of the distribution of rationality to the order of presentation of G_1 and G_2 between samples. Missing from the literature to date are broader explorations of predictive accuracy of ring game rationality measures at both the aggregate and individual levels.

Our test begins with a simple within-subject manipulation of the original ring game, presented in Figure 2.1, Panel B. In our modified ring games, P4 continues to have a dominant strategy with identical payoffs to the Kneeland configuration. Our modification is implemented for players P3, P2, and P1. Though not a best response to P4's dominant strategy, for P3 action *c* yields a high arithmetic mean payoff and no possibility of earning zero. This 'focal' strategy may be attractive for P3. Further, the existence of this focal strategy for P3 may present a dilemma for P2: can they be assured the P3 won't play it? Of course, if they reason at order 2 or higher they must believe that P3 will play best response to P4's dominant strategy and not this focal strategy. Note that P2 also has a focal strategy, action *a*, which is neither a best response to P3's focal strategy nor a best response to P3's best response to P4's dominant strategy. An identical strategy configuration exists for P1: an iterated best response strategy, a focal strategy, and a best response to P2 playing their focal strategy.

Our modification generates a compelling environment in which to test for the predictive accuracy of measures of strategic sophistication revealed in the ring game. The

to subject's cognitive resources and their understanding of the strategic environment. Georganas et al. (2015) report some cross-game consistency of strategic levels in undercutting games, but little consistency in two-person guessing games.

sets of k^{th}-order rationalizable strategies at each role are unchanged. Hence, the natural prediction is that revealed orders will be identical in the original and modified ring games at both the aggregate and individual level. Nonetheless, our modification is not merely a replication with alternate payoffs or a different ordering of tasks. We design a 'focal' strategy for each player position. The focal strategies were designed to carry the highest arithmetic mean payoff, and have no possibility of earning a zero payoff.[2] Cooper et al. (1990) show that such focality in payoffs for off-equilibrium strategies affects players' choices of Pareto-dominant Nash equilibrium strategies. In our work, we similarly induce focality in payoffs associated with non-rationalizable strategies. Players might be attracted to such focal strategies; and additionally, they might forecast others being attracted to focal strategies, and best respond thereto.[3]

The focal strategies in our modified ring games provide a plausible non-rational strategy that a given player might use. Players iteratedlly best responding to this focal strategy would, themselves, appear non-rational, despite the fact that they are iteratedly best-responding to their beliefs. These best-responding subjects would appear less strategically sophisticated than they truly are if the researcher ignores the presence of focal strategies. Hence, if focal strategies are used and believed to be used by subjects, one may expect quite limited consistency when measuring rationality orders in both the original and modified ring games. **?** theoretically and empirically analyze a related separation between rationality and strategic sophistication, driven by non-degenerate beliefs about rationality. Players without point beliefs on the rationality of others may appear non-rational even if they are

[2]An alternative source of focal point is through labeling. See Mehta et al. (1994), Bardsley et al. (2010).

[3]Crawford et al. (2008) shows iterative reasoning to focal strategies in a set of coordination games. Their focal strategy is based on focality in the labelings. They show that the effects of focality in the labeling are offset by payoff differences in coordination games.

iteratedly best-responding to said beliefs. Such players' strategic sophistication would thus be mis-measured. In our design, focal strategies and best responses thereto create a unique prediction for a type of strategically sophisticated but non-rational play in the **?** framework. In developing our experiments we pre-specified the potential attraction to our focal strategies, and developed our predictions and identification of strategic sophistication based on this focality.[4]

In a sample of 200 subjects, we conducted the standard Kneeland ring games and the modified ring games presented in Figure 1. In the standard ring games, we closely reproduce the findings of Kneeland (2015). Only 5% of subjects fail to choose the dominant strategy as P4. Using Kneeland's identification approach, 24.5% of subjects are classified as having level 1 rationality, 37% as level 2, 13% as level 3, and 18.5% as level 4 (the remaining 7% are unclassified). This classification is both similar to and similarly precise as that of Kneeland.[5]

In modified ring games, again only 5% of subjects fail to choose the dominant strategy as P4. However, the presence of focal strategies has a dramatic effect on player behavior. Subjects appear to have qualitatively lower levels of rationality in our modified games. Using Kneeland's identification approach, 55% of subjects are classified as having level 1 rationality, 9.5% as level 2, 6.5% as level 3, and 0% as level 4 (the remaining 29% are unclassified). Within classified subjects, there is limited correlation in the levels of rationality revealed in the original and modified ring games (Spearman correlation: 0.11, $p = 0.11$). At both the aggregate and individual level there is limited predictive power

[4]Interested readers are referred to our pre-analysis plan, which lays out the details of focal strategies in the identification of strategic sophistication: https://doi.org/10.1257/rct.5937-1.0.

[5]Similar to Kneeland (2015), we find that 93% of subjects are within one choice of their identified level.

from the extent of strategic sophistication measured in original ring games.

We attempt to understand the lack of predictive validity by examining in more detail the strategies enacted in our modified ring game. Perhaps the modification we conducted generated substantial confusion or added too much complexity to the environment, muddying responses and hampering prediction. Such ex-post arguments are contradicted by important regularities within the modified ring game data. First, as noted above, 95% of P4 players choose their dominant strategy. Importantly, 80% of P1 subjects play their focal strategy, making clear the attraction to focal actions. Echoes of this attraction are also seen for other players. For example, 39% of P2 subjects play their focal strategy and 15% play the best response to P3 playing their focal strategy. Such regularity in play contradicts the claim that our modified ring game merely generates noisy response.

In developing our study we pre-specified the strategies that we believed would be focal for each player. We incorporate this focality in the identification of levels of strategic sophsitication (see subsection 3.3 for details). Incorporating focal strategies into the identification generates a distribution of strategic levels that appears more similar to that derived in the original ring game: 52% are classified as level 1, 30% as level 2, 11.5% as level 3, and 6.5% as level 4. At the individual level, strategic sophistication also aligns much more closely between games once focality is incorporated into the identification, (Spearman correlation: 0.35, $p < 0.01$). This restoration of predictive accuracy demonstrates the value of ring game measures of strategic sophistication, but requires a detailed understanding of the strategic environment and the influence of focal strategies on all players.

Our results highlight the challenge and the promise of predictive validity for measures of strategic sophistication. When viewed at face value, out-of-sample predictive accuracy is quite limited. Nonetheless, once the researcher is endowed with a reasonable model of the strategic environment, accuracy is restored. Measures of strategic sophistication are consistent when appropriately grounded, but how is a researcher to know what that grounding should be? Interestingly, the existence of a dominant strategy in our modified ring game, chosen by an overwhelming majority of subjects, does not provide sufficient grounding to organize the rest of play.

One may consider our results and the sensitivity of rationality measures as a basis for preferring more abstract measures of strategic sophistication, such as cognitive ability or psychological assessments. Our data indicate that this is unlikely to be a valuable path forward. In addition to our modified ring games we also collected data on cognitive ability measures and psychological tests such as the "reading the mind in the eyes" task (Baron-Cohen et al., 2001) that have been used in prior work. These more abstract measures have limited correlation with strategic sophistication in either the original or our modified ring games.

The lesson from our work is that researchers must be extremely careful when porting measures of strategic sophistication out-of-sample. Prediction is possible from the measures of Kneeland (2015), but requires a thorough advanced assessment of the strategic environment in which the researcher is attempting to make predictions.

This manuscript proceeds as follows: we introduce our experimental design in section 2. Section 3 presents our main results. In section 4, we discuss more abstract measures of strategic sophistication. Section 5 concludes.

2.2 Experimental Design

We implemented two sets of ring games to elicit subjects' orders of rationality — Kneeland's original ring games and our modified ring games. We also issued four sets of more general ability and cognitive tests — best response task, an "eye gaze" test for Theory of Mind ability (Baron-Cohen et al. (2001)), Mensa IQ practice problems, and the three-item cognitive reflection test (Frederick (2005)).

2.2.1 Original Ring Games and Modified Ring Games

We adopted Kneeland's design of the ring game as the original ring game in our experiment (see Figure 2.1, Panel A). To examine the predictive power of orders of rationality derived from the original ring game, we modified the original ring game by introducing 'focal' actions for player positions P3, P2, and P1. These focal strategies have the highest arithmetic mean payoff for each respective player position, while avoiding zero payoffs. For example, in Figure 2.1, Panel B, action *b* is the focal strategy for P1, action *c* is the focal strategy for P2 and action *c* is the focal strategy for P3. These focal strategies are never dominant strategies, nor are they iterated best responses to P4's dominant strategy. In addition to having an iterated best response to P4's dominant strategy and a focal strategy, P1(P2) also has a best response to the P2(P3) playing their own focal strategy. Our modified ring games do not make any changes to P4's dominant strategy payoffs, but do alter the payoffs associated with the other P4 strategies. Appendix B.1 provides screenshots of the experimental instructions.

2.2.2 General Strategic Ability and Cognitive Tests

To provide more abstract measures of strategic sophistication, we also implemented a set of strategic ability and cognitive tasks. First, we issued a set of Best Response (BR) tasks. The BR tasks test best response ability by directly defining the actions of all the counterparts in the game. This portion of the study was composed of four additional ring games. Each BR ring game was matched to a modified ring game, by reducing each payoff entry of the modified ring game by 1 (except the zeros). In the BR tasks, the three 'other players' are replaced with computer players. Subjects were informed that they were playing against three computers that choose the actions that give them the highest earnings. The subjects were also informed that the computer P4 had chosen the dominant action when they were not playing at the P4 position. Subjects played 4 BR games, at each player position exactly once. Appendix B.1 provides screenshots of the BR tasks.

Subjects also completed three sets of cognitive tests. First, to measure Theory of Mind (ToM) ability, we use the "Reading the Mind in the Eyes" task (also known as the "eye-gaze test") (Baron-Cohen et al., 2001). In the eye gaze test, subjects identify the emotions being expressed by a pair of eyes in 36 photographs. ToM ability is assessed from the number of correct identifications. A few economics papers have incorporated this test in their design. The findings suggest that higher test scores in the eye-gaze test seem to be an indication of stronger ability in strategic thinking (Bruguier et al., 2010; Georganas et al., 2015). Second, we selected 10 questions from the Mensa practice problems as a measure cognitive ability. Subjects had 10 minutes to complete the 10 problems. Third, we implemented the standard 3-item Cognitive Reflection Test (CRT) of Frederick (2005). Appendix B.1 provides screenshots of these additional tasks.

2.2.3 Experimental Timeline

Two-hundred subjects were recruited from the subject pool of the Incentive Lab (University of California - San Diego, USA). We conducted the experiment using the Z-Tree program (Fischbacher, 2007). Subjects received a $5 show-up fee and earned an additional $17 on average (The minimum payment was $5, and the maximum payment was $35).

The experiment was conducted in three main sections. First, subjects were given instructions for the eye-gaze test. After practicing for 1 round, they completed the 36 questions of the task without feedback between questions. This part of the experiment was unincentivized.

Second, subjects were introduced to the ring games using the same instructions as Kneeland's (2015) experiment. Experimental instructions were read aloud by the experimenters at the beginning of this portion of the experiment. Subjects were then given an understanding test consisting of 4 questions. For each understanding question, subjects were able to see feedback with correct answers and explanations immediately after they submitted their answers to ensure that the subjects understood the experiment clearly. Subjects were then randomly matched into groups of four. They played the original ring games (G1 and G2) and the modified ring games (G3 and G4) within these fixed groups. Each subject played in each player position for each ring game exactly once, making 16 games in total. The 16 ring games were issued in alternate order (one modified game, one original game, one modified game, etc.). There was no feedback given between games. For each ring game, subjects were given at least 90 seconds to make a decision. They were allowed to use more time if needed, but subjects were not allowed to proceed prior to 90 seconds. After the 16 ring games, subjects completed the 4 BR tasks. Subjects were informed at the beginning of the BR tasks that they were playing against the computer players. One of the 20 total games (16 ring games and 4 BR tasks) were randomly selected for payment at the end of section 2 of the experiment.

Third, subjects completed the ten Mensa problems and the three CRT questions without incentives. We also collected some demographic information at the end of the experiment. Sessions lasted for about 75 minutes.

2.3 Results

The results are presented in three subsections. First, we examine behavior in our implementation of the Kneeland ring games, reproducing the general patterns and distribution of types observed in Kneeland (2015). Having demonstrated this replication, in a second subsection we investigate modified ring game behavior. We document limited consistency in rationality measures across the original ring game and our modified ring game at both the aggregate and individual level. In a third subsection, we attempt to understand the lack of predictive accuracy for rationality measures identified in the Kneeland ring game. We investigate whether subjects appear drawn to the focal strategies of our design, and whether incorporating focal strategies into the identification of strategic sophistication helps to organize behavior. We show the plausibility of this identification approach, delivering both more similar distributions of orders and a three-fold increase in the individual-level correlation in strategic sophistication across games.

2.3.1 Replication of Kneeland (2015)

Figure 2.2 presents histograms of the strategies played in each player role and game for the 200 subjects in our study. The strategy labeling "(p,q)" corresponds to strategies in both versions of the game noted in Table 2.1. White bars correspond to our implementation of the standard Kneeland ring game. Beginning with P4, 95% of subjects play the dominant strategy, (a,c). As P3, 68% of subjects play (a,b), the best response to P4's dominant strategy, indicating rationality of at least level 2. As P2, 32% of subjects play (b,a), the iterated best response to P4's dominant strategy, and indicating rationality of at least level 3. As P1, 27.5% of subjects play (a,c), the iterated best response to P4's dominant strategy, and indicating rationality of at least level 4.

The raw data of Figure 2.2 provides a clear indication of the distribution of levels of strategic sophistication, but cannot be used to infer the distribution if subjects switch between levels of play across roles. Kneeland (2015) provides an identification technique based on strategies at each player-role. For example, to be classified as R2, a subject must choose the dominant strategy as P4, choose the best response to this dominant strategy as P3, and choose constant strategies in all other player roles. Kneeland's identification approach permits one deviation from this rule. Table 2.2 implements this approach on our original ring game and contrasts the resulting distribution with Kneeland's results. For our 200 subjects in the standard ring game 7% of subjects are not within one deviation of a given level and so classified as R0, 24.5% are R1, 37% are R2, 13% are R3, and 18.5% are R4. The same values for Kneeland (2015) are 6% R0, 23% R1, 27% R2, 22% R3, and 22% R4.

Kneeland's seminal ring game results were notable partially because they indicated

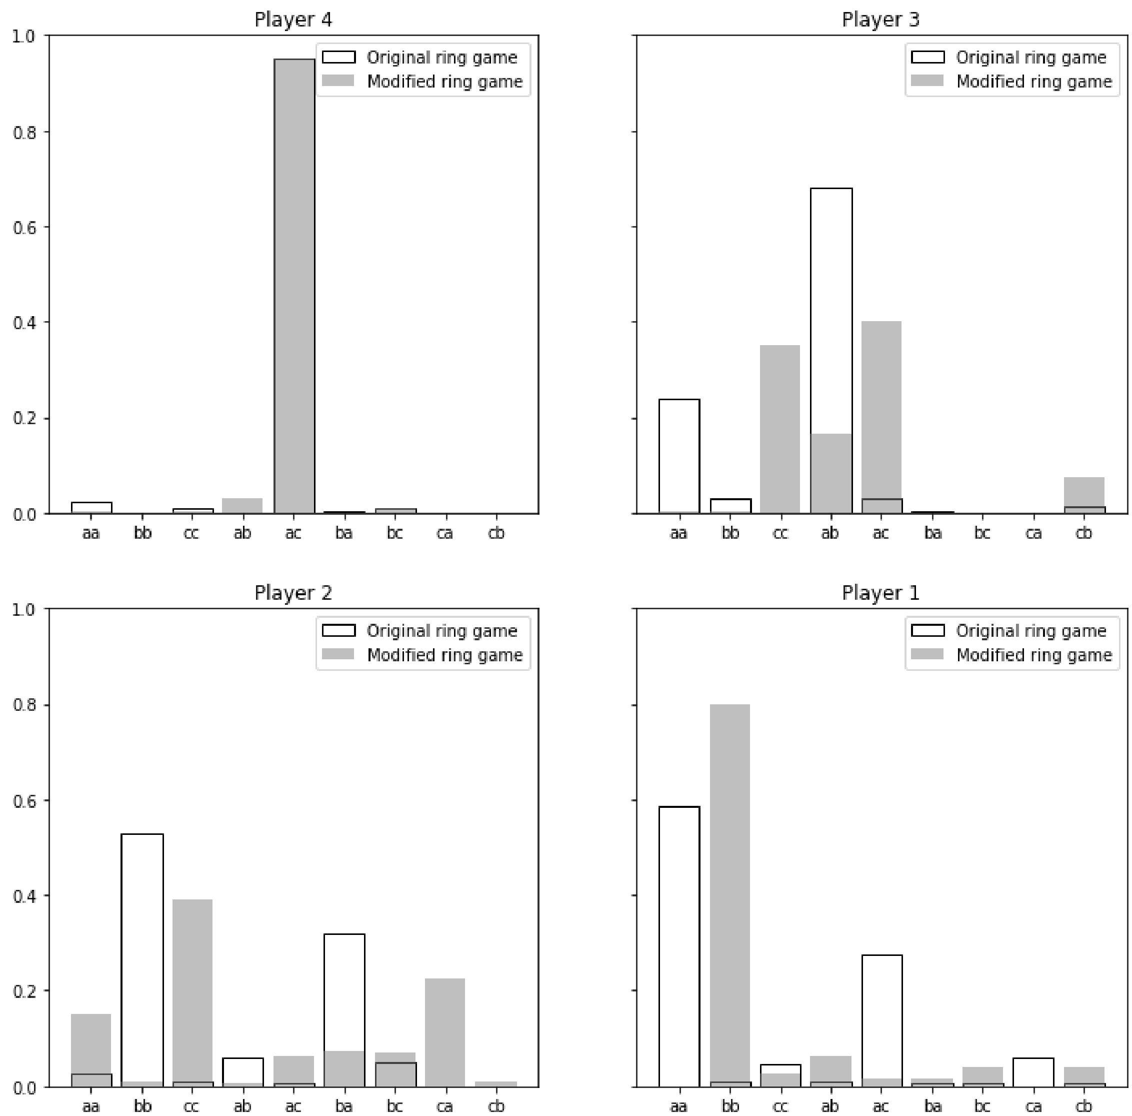

Notes: The first letter in strategy labeling refers to subject's choices in game G1 or G3, and the second letter refers to subject's choices in game G2 or G4.

Figure 2.2: Action profiles for the 16 ring games

a greater proportion of higher rationality levels than generally documented in other designs. Kneeland's results showed that of classified (non-R0) types, 76% were above R1 and around 46% were R3 or higher. We reproduce this core finding: of classified (non-R0) types, 74% are above R1 and around 34% are R3 or higher. Indeed, our distribution of levels of strategic sophistication is statistically indistinguishable from that of Kneeland (2015) in a Mann-Whitney test ($p = 0.18$).

Table 2.2: Subjects classified by order of rationality, by game

	R0	R1	R2	R3	R4	#
Original ring game	7%	24.5%	37%	13%	18.5%	200
Modified ring game	29%	55%	9.5%	6.5%	0	200
Kneeland (2015) results	6%	23%	27%	22%	22%	116

Notes: Mann-Whitney test p-values: Original ring game and Kneeland's results, 0.18; Modified ring game and Kneeland's results, < 0.01; Original ring game and modified ring game, < 0.01.

2.3.2 Modified Ring Game

Figure 2.2 contrasts strategies in our modified ring games with those played in the original versions. Beginning with P4, 95% of subjects play the dominant strategy, (a,c), identical to the standard ring game. At the other player roles, behavior diverges considerably. Only 16.5% of P3 subjects play (a,b), the best response to P4's dominant strategy, indicating rationality of at least level 2. More prevalent strategies are (c,c), a level 1 response, and (a,c), which is not rationalizable. As P2, only 7.5% of subjects play (b,a), the level 3 strategy, and as P1 only 1.5% of subjects play (a,c), the level 4 strategy. The most prevalent strategies at these player roles are (c,c) and (b,b), respectively, both level 1 responses. In player roles P3, P2, and P1, we reject the null hypothesis of identical distributions of strategies across the original and modified ring game.[6]

Table 2.2 conducts the Kneeland identification on our modified ring game behavior and provides distributional contrasts with the original ring game. For our 200 subjects in the modified ring game 29% of subjects are not within one deviation of a given level and so classified as R0, 55% are R1, 9.5% are R2, 6.5% are R3, and 0% are R4. In addition to having more R0 types than the standard ring game, the modified ring game provides a substantially different image of the proportion of higher types. Of classified (non-R0) types, only 23% are above R1 and around 9% are R3 or higher. We reject the null hypothesis that the distributions of levels of strategic sophistication are equal across our original and modified ring games in a Mann-Whitney test ($p < 0.01$).

Table 2.3: Subject's types by modified ring game and original ring game

		Original RG					
		R0	R1	R2	R3	R4	Total
Modified RG	R0	3%	4%	11%	4.5%	6.5%	29%
	R1	4%	20%	21%	5%	5%	55%
	R2	-	0.5%	3.5%	2.5%	3%	9.5%
	R3	-	-	1.5%	1%	4%	6.5%
	R4	-	-	-	-	-	
	Total	7%	24.5%	37%	13%	18.5%	

Notes: Spearman correlation: 0.11

Despite the reliable replication of dominant strategy play for P4, the aggregate distributions of strategic sophistication appear dramatically different across our original and modified ring games. In Table 2.3, we examine individual level behavior by correlating the classified levels across the two games. The Spearman correlation in strategic sophistication

[6]Mann-Whitney tests: P1, $p < 0.01$; P2, $p < 0.01$; P3, $p < 0.01$; P4, $p = 0.39$.

across the two games is 0.11. In addition to yielding a qualitatively different message on the distribution of strategic sophistication, individual classifications in the modified ring game are largely independent from those in the original ring game. This independence highlights a marked inability to predict behavior from ring game rationality classifications. Relying on the original ring game to make out-of-sample predictions to the modified ring game, one would accurately predict for only 27.5% of subjects (the diagonal sum of Table 2.3), which is only marginally better than accurate prediction with uniform random choice, i.e., 20%.[7]

2.3.3 Focal Strategies: Understanding Predictive Inaccuracy

Though both the original and modified ring games reliably generate dominant strategy play for P4, the remaining players in the modified ring game do not reliably iteratedly best respond to this play. A natural question is whether the modified ring game generates confusion, leading to unreliable, noisy data. Under this interpretation, the lack of predictive validity from strategic sophistication identified in the original ring game is merely a consequence of complicated experimental design. An alternative is that subject behavior is organized, but around salient features of the modified ring game rather than P4's dominant strategy play. We discern between these two interpretations by assessing an alternative taxonomy of strategic sophistication, which incorporates focal strategies into the identification strategy.

Importantly, one cannot provide the identification approach *after* having examined the data, it must be pre-registered. If a researcher does not pre-state their proposed method for identification, then they may make use of additional degrees of freedom afforded by

[7] Test of proportions, $p = 0.11$.

having peeked at the data. Our modified ring game design had a specific objective of providing players P3, P2, and P1 a focal strategy with high arithmetic mean payoffs, a best response to the focal strategy of the player above them, and an iterated best response to P4's dominant strategy. We presumed that P4's dominant strategy is also focal. This understanding of our design led us to construct an alternate taxonomy, which we label F to denote its basis in focality. Table 2.4 provides this taxonomy, which was pre-registered with our pre-analysis plan and can be found at https://doi.org/10.1257/rct.5937-1.0.

Table 2.4: Predicted actions under rationality and assumptions about focality in modified ring games

	Modified Ring Games							
	P1		P2		P3		P4	
Type	G3	G4	G3	G4	G3	G4	G3	G4
F1	(b,b)		(c,c)		(c,c)		(a,c)	
F2	(b,b)		(a,a)		(a,b)		(a,c)	
F3	(c,c)		(b,a)		(a,b)		(a,c)	
F4	(a,c)		(b,a)		(a,b)		(a,c)	

Under the F measure of strategic sophistication, an $F1$ player plays their focal strategy, an $F2$ player best responds to the player above them playing their focal strategy, an $F3$ player iteratedly best responds to the player two above them playing their focal strategy, etc. Table 2.4 presents the strategies associated with each level in the modified ring game. The admissible strategies under the F-measure are a strict subset of those for the R-measure in the modified ring game noted in Table 2.1. Note, however, that a single behavioral pattern would be interpreted differently by the two measures. For example, consider the P1 strategy (c,c), which would be interpreted under the R-measure as either an $R1, R2$, or $R3$ strategy, but is interpreted under the F-measure as only an $F3$ strategy.

Figure 2.2 demonstrates the value of incorporating focality into the identification of strategic sophistication. Consider P3: where (a,a), (b,b), and (c,c) strategies are admissible under $R1$, but only (c,c) is admissible under $F1$. Among these three, only the focal strategy, (c,c), is chosen in the modified ring game. Additionally, (c,c) is the second most frequent choice for P3, and more frequently chosen than the best response to Player 4's dominant strategy, (a,b).[8] For P2 and P1, the most frequently chosen strategy are focal with (c,c) chosen 39% of the time by P2 , and (b,b) chosen 80% of the time by P1. Among the constant strategies, (a,a), (b,b), (c,c), subjects much more frequently choose the one that is F-admissible at each player role relative to the others (51.3% vs 6.5% of strategies for players P3, P2, and P1). In addition to focal strategy play, echoes of the effect of focal strategies for other players are observed: 15% of P2 play (a,a), the best response to P3 playing their focal strategy.

Because the admissible strategies under the F-measure at each level are subsets of those admissible for the R-measure in the modified ring game, the identification method of Kneeland would necessarily classify fewer subjects under levels of $F1-F4$ than $R1-R4$.[9] One benefit of the increased restrictiveness of the F-taxonomy is that for each player role at each level, an exact prediction for behavior is made. This allows us to implement spike-logit estimation techniques frequently used in the estimation of strategic sophistication in other

[8] The modal choice for P3 is (a,c), which is not rationalizable, but is within one deviation of both the focal strategy and the best response.

[9] Indeed, under the identification technique of Kneeland 39.5% of subjects are not within one deviation of a given level and so classified as F0, 43% are F1, 14% are F2, 2.5% are F3, and 1% are F4. The increase of 10%-age points in unclassified subjects relative to the R-taxonomy is due to the increased restrictiveness of admissible strategies under the F-taxonomy. Interestingly, despite this increased restrictiveness, the Spearman correlation in levels of strategic sophistication between the standard ring game R-level and the modified ring game F-level is 0.06, a similar value to the correlation in R-levels across the two games.

games[10] to provide a best fitting F-level for each subject.[11] Appendix A.1 provides the details of this estimation strategy and its in-sample properties. Notably, this estimation method accurately predicts 82% of choices in-sample for the modified ring game.

Table 2.5: Subject's types by modified ring game(F) and standard ring game(R)

		Original RG					
		R0	R1	R2	R3	R4	Total
Modified RG	F1	3.5%	20.5%	18%	5%	5%	52%
	F2	2%	3%	14%	5.5%	5.5%	30%
	F3	-	1%	4%	2.5%	4%	11.5%
	F4	1.5%	-	1%	-	4%	6.5%
	Total	7%	24.5%	37%	13%	18.5%	

Notes: Spearman correlation: 0.35

Table 2.5 provides the distribution of F-levels recovered from our maximum likelihood exercise along with the individual correlation between R and F levels across games. Because the classification technique provides a best fitting F-type for each subject, no one is unclassified.[12] For our 200 subjects, 52% are classified as F1, 30% as F2, 11.5% as F3, and 6.5% as F4. This classification leads to a different representation of the levels of sophistication than previously obtained in the modified ring game: 48% of subjects are classified above F1 and 18% are F3 or higher. These values are around twice as large as those from Kneeland's R-level classification in this game (23% above R1, and 9% R3 or

[10]See, for example, Costa-Gomes and Crawford (2006), García-Pola et al. (2020).

[11]We cannot conduct the same estimation for the R-measure. Lower R types are allowed to choose any actions at higher player positions, as long as the actions are consistent across the two games. As a result, one cannot pin down a subject's action and his belief about other player's actions if he is not choosing a rationalizable strategy. However, we are able to conduct the same estimation for the F-measure in the original ring games. See appendix A.1 for details.

[12]There are a few subjects for whom the fit is very poor. We include these subjects because the corresponding risk of misclassification should work against finding predictive accuracy across games.

higher), and more closely corresponds with the message from original ring game on the distribution of strategic sophistication.

Table 2.5 also reports the Spearman correlation from Kneeland's *R*-classification for the standard ring game and the *F*-classification for the modified ring game. Here, we identify much more stability in identified types than previously documented. The Spearman correlation in levels of strategic sophistication is 0.35. Indeed, now relying on the standard ring game to make out-of-sample predictions to the modified ring game, one would accurately predict for 41% of subjects, substantially more than the uniform chance rate of 25%.[13]

We have shown that the *R*-classification has limited predictive accuracy across games, while the *F*-classification in the modified ring game and the *R*-classification in the standard ring game are more highly correlated. One may naturally wonder about how well the *F*-classification performs across games. Kneeland's original ring game also features strategies with high arithmetic mean payoffs and avoidance of zeros. However, these strategies and best responses thereto coincide with iterated best response to P4's dominant strategy. Hence, it is challenging to disentangle rationality from focality in the original Kneeland game. Nonetheless, in Appendix A.1 we provide the *F*-classification applied to the original ring game. As expected, the distributions of *F*- and *R*-orders are quite similar for the original game. Thus the correlation between *F*-classifications across the original and modified games is quite similar to that reported in Table 2.5, (Spearman correlation = 0.38).

[13]The uniform chance rate is 25% as there are no F0 types in our classification technique. Test of proportions, $p < 0.01$.

2.4 Discussion: More General Measures of Strategic Sophistication

The results to here show clear challenges for the predictive validity of measures of strategic sophistication derived from experimental ring games. Using measures based on orders of rationalizability, the extent of strategic sophistication is largely uncorrelated across the original ring game of Kneeland (2015) and our modified ring game. This lack of predictive power is not a product of noisy response or decision errors. Our modified ring games incorporate focal, high average-payoff strategies into the original ring game. Pre-specified measures of strategic sophistication incorporating focal strategies in the identification strategy in the modified ring game correlate much more highly with original ring game play than standard rationality measures do.

Insights into the organizing principles of play between the original and modified ring games permit us to make more reliable predictions across games. Without such insights, how can researchers make predictions across different contexts? If different contexts generally yield different strategic approaches, as ours apparently do, any given measure may lack predictive validity.

One potentially promising approach to restoring validity is to use more general measures of strategic sophistication, untied to a specific game or belief structure on the play of others. We consider four such measures. First, to remove dependencies of strategic play on the beliefs of others' strategies, we implemented a Best Response (BR) task. Our BR task is identical to the ring game, except subjects are told the other players are computers maximizing their payment.[14] We calculate the level of best response BR1-BR4 following the logic of our prior taxonomies.[15] Second, we conduct a general psychological test for general understanding the emotions and intentions of others, the ‘eye-gaze test’ of Baron-Cohen et al., (2001). Third, to provide a general measure of cognitive ability, we consider an IQ test composed of 10 Mensa practice problems. Fourth, we consider the standard 3-item cognitive reflection test (CRT) of Frederick (2005) as an alternate cognitive ability measure. Our eye-gaze and cognitive ability results are comparable to those of previous studies.[16]

If more general measures yield more portable lessons on strategic sophistication, they should explain behavior to a similar degree across different strategic environments. In Table 2.6, we provide corresponding analysis, regressing the degree of strategic sophistication in our standard and modified ring games on our four more general measures, controlling for observable characteristics. Beginning with BR tasks in columns (1), (3), and (5) of Table 2.6, we find that either alone or conjunction with the other measures, best response ability correlates significantly with standard ring game and modified ring game behavior.

[14]See appendix B.1 for detailed experimental instructions.

[15]With the elimination of strategic uncertainty, the game becomes plausibly ‘easier’ and more subjects will be classified as higher order rational. This is indeed what we observe with our subjects. Only 2% of the subjects are unclassified. There are 21.5% of the subjects choosing focal strategies at each player position, therefore classified as BR1 — best responding to their own payoff. There are 16.5% of subjects choosing actions that are best responses to the focal strategies of one player above them (1st order payoff) and classified as BR2. We also observe 14.5% of subjects being BR3, playing actions that are iterated best responses to the focal strategies of two players above them (2nd order payoff). Lastly, there are 45.5% of subjects that are classified as BR4, playing actions of best responses to P4’s dominant strategy at each player positions. This significant proportion of higher types is very different from our results of the modified ring games, where we only observe 17.5% F3 or higher types.

[16]Baron-Cohen et al. (2001)Baron-Cohen et al. (2001) reports an average score of 81%. Our average score is 78%. Zhao (2020)Zhao (2020) uses the same set of Mensa practice problems and reports an average score of 7. We report an average score of 6.7. For the CRT test, Frederick (2005)Frederick (2005) reports an overall average percentages of players scoring (0,1,2,3) as (33%, 28%, 23%, 17%). Our scores are distributed as (33.5%, 25.5%, 23%, 18%). This is in fact very close to the results from the meta-study.

Importantly, the degree of predictive power across the two games differs qualitatively. The partial R^2 values for BR measures are three times higher for the R-taxonomy in the standard ring game than for R-taxonomy in the modified ring game. In the modified ring game, BR measures explain only 4 percent of the variation in R-values. Following our experimental findings, one may imagine that BR measures would be predictive of modified ring game behavior if one considered the F-taxonomy. That is not the case as BR measures have even less predictive power for the F-taxonomy in the modified ring game, explaining only around 3 percent of the variation in F-values. In addition to BR measures, we consider the the alternate general measures for strategic sophistication in columns (2), (4), and (6) of Table 2.6. The alternate measures add very little explanatory power regardless of game or measure of strategic sophistication under consideration.

Table 2.6: Regression results

	DV: R-type (standard)		DV: R-type (modified)		DV: F-type (modified)	
	(1)	(2)	(3)	(4)	(5)	(6)
BR test	0.36^{***} (0.06)	0.35^{***} (0.06)	0.13^{***} (0.04)	0.12^{***} (0.05)	0.12^{**} (0.05)	0.12^{**} (0.05)
Eye gaze test	-	0.00 (0.02)	-	0.01 (0.02)	-	-0.03 (0.02)
IQ test	-	0.11^{**} (0.04)	-	0.04 (0.03)	-	-0.00 (0.04)
CRT	-	-0.07 (0.09)	-	-0.01 (0.06)	-	0.01 (0.07)
age	-0.02 (0.02)	-0.02 (0.02)	-0.00 (0.01)	-0.00 (0.01)	-0.01 (0.01)	-0.01 (0.01)
gender	-0.09 (0.14)	-0.07 (0.15)	0.01 (0.10)	0.02 (0.11)	-0.07 (0.11)	-0.04 (0.12)
year in school	0.13^{*} (0.07)	0.12^{*} (0.06)	0.06 (0.05)	0.05 (0.05)	-0.04 (0.05)	-0.04 (0.05)
major	-0.04 (0.04)	-0.03 (0.04)	0.01 (0.03)	0.02 (0.03)	0.01 (0.03)	0.01 (0.03)
race	0.01 (0.04)	0.02 (0.04)	0.01 (0.03)	0.02 (0.03)	0.03 (0.03)	0.02 (0.03)
Multiple R^2	0.18	0.21	0.05	0.06	0.05	0.06
Adjusted R^2	0.16	0.17	0.02	0.02	0.02	0.01
Partial R^2	0.15	0.18	0.04	0.06	0.03	0.04

Notes: Column (1) (3) and (5) preserves the results without additional cognitive test. Column (2) (4) and (6) includes the cognitive test results as regressors. The reduced model includes age, gender, year in school, major and race. Partial R^2 is for everything above the demographics. $*$ indicates $< 10\%$ significance. $**$ indicates $< 5\%$ significance. $***$ indicates $< 1\%$ significance. Standard error in parenthesis.

The findings of Table 2.6 suggest that more general measures of strategic sophistication, while potentially correlated with behavior in different environments, cannot generally be used to make accurate predictions in different strategic contexts. Even our most credible general test, the BR task, where subjects experience the same game form as our ring games, consistent predictions are not obtained across contexts.

2.5 Concluding Remarks

Knowing a player's level of strategic sophistication from one environment, can we make reliable predictions about their behavior in another? This is the central question addressed in this paper. In a laboratory experiment, we provide two sets of ring games (Kneeland, 2015), and examine the consistencies of resulting rationality measures. The games we use retain the key organization of the original ring game of Kneeland (2015), and our modification introduces focal, high average-payoff strategies for each player position. When using the standard rationality measures to identify subjects' degree of strategic sophistication, we observe substantial differences between the two sets of ring games at both the aggregate and individual level. At the behavioral level, this inconsistency is driven by subjects playing focal strategies and best responses to focal strategies in our modified games.

When we incorporate focal strategies in identification of strategic sophistication, we are able to restore predictive power. We further investigate if more general ability and psychological tests may be easier routes to provide out-of-sample predictions in new strategic environments. There is little success with these more general measures.

What can researchers do if they wish to make predictions for strategic play across contexts. Our findings point to one possible path: investment into understanding the organizing principles of play. If each context's additional considerations are understood in detail, accurate predictions may be obtained. Predictive validity is restored in our setting by specifying organizing principles of play around focal strategies. In other contexts, the organizing principles will clearly be different than ours, but they must be understood in detail. Without that understanding, external validity for any given measure will be compromised and the measurement of strategic sophistication may lack purpose.

2.6 Acknowledgement

Chapter 2, in part is currently being prepared for submission for publication of the material. It is coauthored with Sprenger, Charles. The dissertation author was the primary author of this chapter.

Chapter 3

Learning of Strategic Sophistication

3.1 Introduction

Nash equilibrium is the benchmark solution concept to predict players' behavior in games. Equilibrium requires that players are perfectly rational, with common identical beliefs about equilibrium play. Many experimental studies to date have shown that people's initial responses in strategic games may not strictly follow equilibrium predictions. Researchers have identified a number of non-equilibrium behavioral rules that better describe initial play (for example: beauty contest game, Nagel (1995); normal-form games, Stahl and Wilson (1994, 1995), Costa-Gomes et al. (2001); two-person guessing game, Costa-Gomes and Crawford (2006); 11-20 money request game, Arad and Rubinstein (2012b), among many others).

Two challenges to the standard equilibrium concept have been identified in these settings. First, the assumption of common knowledge of rationality may face challenges simply because players may not be strategic enough to understand the games fully and solve for the equilibrium immediately. Second, even if the players are aware of the equilibrium of the game, they may not have common identical beliefs about others also knowing it. Absent the assurance of common knowledge of rationality, players' beliefs can be potentially influenced by many factors including past experience and past opponent's play.

Strategic players may update their beliefs about other players when historical information about actions or payoff is given in repeated games. Researchers have utilized statistical learning models to predict subjects' future choices, after observing his responses to such historical information (e.g., Fudenberg and Kreps (1993), Fudenberg and Levine (1998a,b), and Camerer and Ho (1999)). They do so by analyzing the subject's initial choices, estimating an updating rule, and assigning the subjects with a probability distribution of actions and corresponding behavioral types according to the actions with the highest probability.

In this manuscript, we analyze how learning from past strategic interactions is deployed in potentially new strategic setting. Can subjects learn their opponent's strategic types through historical information of actions even if those actions occurred in a different context? Given that there is relatively limited literature on learning across strategic context as compared to the broader literature on learning, this question is of crucial importance. If players can learn about the strategic types of other players through their historical actions in different environments and respond to the learned types, researchers will be able to make broader predictions about behavior.

We conduct a set of online 'laboratory' experiments to investigate this question. Subjects in our experiment proceed through 15 strategic games. In the first part of the experiment, they play five rounds of 11-20 money request games. This game provides a natural trigger for level-*k* reasoning model (Stahl and Wilson, 1994, 1995; Nagel, 1995), which is an off-equilibrium behavioral model used to characterize subject's initial responses in games. We collect the subject's initial play with the first part of the experiment. In the second part of the experiment, subjects play a subsequent ten additional strategic games. In treatment 1, they play the 11-20 games again, while in treatment 2, they play ten rounds of 6-15 money request games. In the second part of the experiment, subjects are revealed their opponent's historical information of actions in part-one games fifty percent of the time. We examine if learning happens by comparing the strategic levels of the subjects for games with and without historical information, both at the aggregate distributional level and at the individual level.

For the first treatment, where the games for the two parts are the same, we document a significant shift in the distribution of choices. Subject choose strategies that correspond to higher-level types much more often when knowing about their opponent's historical actions. At the individual level, knowing the opponent's most recent strategy has a significant effect in shaping the subject's behavior in part-two games. In contrast, for the second treatment with different part-two games, we do not find a significant difference in the aggregate distributions of actions. However, at the individual level, knowing the historical information of the opponent's past choices in the 11-20 game still has an effect on the subject's decisions in the 6-15 game. More importantly, when the strategic games are different, subjects rely more on the historical information of their opponent's initial choices in the other game.

We interpret the portability of historical information through the lens of levels of strategic sophistication. Although strategic contexts are different, as the case in our second treatment, the common element of people's behaviors in such strategic games is the levels of strategic sophistication. The strategic levels, reflected from their choices in the games, carry some information about the players, that could be utilized in following interactions. Stahl (1996) constructs a rule learning model that is similar to our approach. He discussed the learning model using Nagel's beauty contest game data (Nagel, 1995). Instead of showing subjects the actions from the previous rounds, subjects are shown the winning number (mean guess) from the previous round. The rule learning model assumes that subjects have an initial propensity towards all the behavioral rules, and update the propensity according to the hypothetical relative performance of each behavioral rule in the previous rounds, given the summary statistics of the games. Although this approach also considers behavioral rules, it is conceptually different from our approach. The rule learning model is game-specific as the updating on the propensities is payoff driven. Rather than learning about other player's behavioral types, this model focuses on learning about which behavioral type performs better. Therefore it will be challenging to utilize the rule learning dynamics across strategic settings.

There are a number of manuscripts that study learning across games. Cooper and Kagel (2008) study learning across signaling games with changing payoffs. They document a significant fraction of sophisticated learners in the population, and report an increasing proportion of sophisticated learning with experience. Cooper and Kagel (2009) examine the effect of using meaningful context, as compared to abstract context, and find that meaningful context aids positive transfers across games. Rick and Weber (2010) study learning across dominance-solvable games. Instead of providing their subjects with feedback, they provide 20 seconds of reflection time after each game. They show that learning that happened in earlier games is transferred to later games. Garcia-Pola and Iriberri (2019) study learning with feedback in a set of normal-form games. They examine subject's naivete and sophistication in initial and repeated play, and they find no correlation.

There are also papers on learning across games that are more different in the setup. Dufwenberg et al. (2010) compare subject's behaviors in the game of 6 and the game of 21, and provide evidence for learning from simpler games to more complex games. Gneezy et al. (2010) study the sequential race game, and show that learning happens within the same game by playing repeatedly. There are fewer errors made by the subjects towards the end of the game. They also show that the concept of solving the game can be transferred to a similar game. Our study adds to the literature by providing a unique angle of learning across games. Studies in the literature have an emphasis on learning through the contextual setup of the games or the underlying solution concepts of the games. Our study points to the direction of learning about the opponent's strategic types. Subjects may not necessarily appear to be more strategic with this type of learning, but they may play more optimally in the new games, as they are able to update their beliefs about their opponents[1] and subsequently choose strategies that are best responses to the beliefs.

Our results highlight the plausibility of learning the levels of strategic sophistication across strategic contexts. We show that subjects are responsive towards historical information of their opponent's past actions. More importantly, our results reveal the relative importance of historical information in different situations. In the learning literature, most studies focus on learning within a single strategic context; and even for those who study learning across games, learning about other player's actions is seldom discussed. However, in real-world settings, repeated interactions with changing strategic contexts should be considered the norm. Without an understanding of learning from other players' historical actions in a different context, how do researchers make reliable predictions about people's behaviors in such situations? Our study points to a plausible way of interpreting historical information, through the lens of levels of strategic sophistication. We argue that strategic levels are mutual components in many strategic setups, and should be considered when collecting information about historical interactions.

This paper proceeds as follows: we introduce the experimental design in section 2. Section 3 presents our main results, and section 4 concludes.

[1]Other works on updating beliefs about opponents include learning through exogenous labeling, see Dufwenberg et al. (2005), Levitt et al. (2011), **?**; through cognitive tests, see Alaoui and Penta (2016), Gill and Prowse (2016); through understanding the strategic environment, Zhao (2020).

3.2 Experimental Design

Our experiment consists of three parts. Part one is identical for all experimental sessions. Part two has two different treatments, explained below. Part three consists of a 5-question Raven test, and a few demographic questions. They are the same across all sessions as well. Subjects are aware that the experiment has three parts at the beginning of the session. However, the instruction for each part is given only when the experiment progresses to that part. The details for each part and the experimental timeline will be explained below.

3.2.1 Part one of the experiments

In part one of the experiments, we ask the subjects to play five rounds of 11-20 money request game (Arad and Rubinstein, 2012b). In each game, subjects play in fixed pairs. They each pick a number between 11 and 20, and receive the amount they pick, in tokens. In addition, if they pick a number that is exactly one less than their opponent, they receive extra 20 tokens. There is no feedback given in part one of the experiment, and the subjects are not informed that they are playing in fixed pairs for this part of the experiment.

3.2.2 Part two of the experiment

Part two of the experiment has two treatments. In the first treatment, subjects play ten rounds of 11-20 money request games. The instruction for these games is the same as the instruction for the part-one games. In addition, the subjects are informed that they may receive some additional information (which later during the game, revealed to them as their opponent's historical actions) for each game. In the second treatment, subjects play ten rounds of 6-15 money request games. In each game, subjects play in pairs. They are asked to pick a number between 6 and 15. They receive the amount they pick, in tokens. In addition, if they pick a number that is exactly one less than their opponent, they receive extra 20 tokens. They also receive additional 5 tokens when calculating payoffs for each 6-15 game (detailed experimental instruction in appendix B.1). This is to ensure that the incentive structure is the same across the two treatments.

During part two of the experiment, subjects switch opponents every two rounds. There are in total ten rounds, therefore each subject plays against five different opponents during part two. In odd number rounds, they do not receive any historical information about their opponent's past choices. In even number rounds, we reveal to the subjects their opponent's choices during part one of the experiment (i.e. subjects see a list of five numbers, for example, "18, 18, 19, 20, 19"). In all ten rounds, we tell the subjects that their opponent may or may not receive some historical information about themselves. This is to ensure that the subject's second-order belief stays constant across games.

To control for the repeated game effect, in each round, we reveal to the subjects that if they have played with their opponent in part one of the experiment. Since the subjects play in fixed pairs in part one, they each only have one other player they have interacted with before. In the first two rounds, we make the subjects play with their part-one opponent, and tell them they have played with him/her in the previous games during part one of the experiment. In the following eight rounds, we explicitly tell the subjects that they are playing with another person in the session that they have never interacted with in part-one games. Beyond the information about the identity of their opponents, and their opponent's past choices, there is no additional feedback given during part two of the experiment.

3.2.3 Discussion of the Experimental Design

In our design, during part two of the experiment, we manipulate the revelation of historical information in alternating order, so that subjects are revealed about their opponent's past actions every other round. Although the subjects are playing with the same opponent for two rounds, we do not explicitly tell them that their opponents do not change from the game without historical information, to the game with historical information. Such setup of the design will not allow us to examine the isolating effects of historical information, fixing the subject's belief about their opponent's other strategic or cognitive traits. Rather, this design provides us an opportunity to compare the strategic behaviors when historical information is available to the baseline behavior where subjects play with a random opponent, without any information. Without informing the subjects of the matching protocols for the second part of the experiment, we recognize the limitations of allowing the subjects to have free beliefs about the person they are playing with. In this sense, the treatment effects examined by comparing the two games with and without historical information are entangled with the potential changing beliefs about the identity of the opponents. We want to point out this feature of the design; and acknowledge that the treatment effect discussed in this paper represents the effect of knowing a random opponent's historical information, rather than the effect of updating about a known opponent based on the availability of their past action information.

3.2.4 Experimental timeline

One hundred and fifty-six subjects were recruited from the subject pool of the Incentive Lab (University of California - San Diego, USA). Subjects were randomly selected to the two treatments. Seventy-two subjects participated in treatment one, and eighty-four subjects participated in the second treatment. We conducted the experiment online using the O-Tree program (Chen et al., 2016) and AWS server. Subjects received a $2 show-up fee and earned an additional $6.5 on average (The minimum payment was $5.6, and the maximum payment was $13.7).

At the beginning of the session, subjects reported to the experimenter using ZOOM video conference. The experimenter sent each subject a private chat message with a personal link to the study. All the information about the study, including the consent form, the instruction, and the experiments are deployed to the online server. Subjects interacted with the web page during the entire experiment. However, they were required to stay in the ZOOM session and keep the video camera on. This was to ensure that the subjects were engaged with the study during the entire session.

The study consists of three parts. At the beginning of each part, the instructions were read aloud by the experiment in the ZOOM conference. Subjects were also able to read the instructions on their screens. We tried to mimic the laboratory setup as much as possible to ensure the subjects understood the instructions fully and concentrated on the study. In the last part of the study, subjects were required to complete a 5-question raven test. For each question, they saw a pattern with a piece missing on the screen. There were also eight pieces below the pattern. Subjects were asked to choose the piece that is the right one to complete the pattern. They had five minutes to complete the test. They were allowed to submit their answers early. They were informed about their score at the end of the experiment. However, this part was not incentivized.

At the end of the session, we collected some demographic information from the subjects. One game of the fifteen money request games from the first two parts of the experiment was randomly selected for payments. Subjects collected tokens during the experiment, and each token was worth 30 cents. The payments were made through Venmo. To ensure that the subjects were properly incentivized during the experiment, the experimenter sent each subject one cent of validation payment at the beginning of the sessions. The final payments were made to the subjects within 24 hours after the sessions. Sessions lasted for about 30 minutes.

3.3 Results

The results are presented in three subsections. First, we examine behavior in part one of the experiment, reproducing the distribution of choices observed in Arad and Rubinstein (2012b) (henceforth AR2012). Having established the baseline for subjects' behaviors in the 11-20 games, in the second subsection we investigate the 11-20 game behaviors under treatment one. When historical information about the opponent's actions in part-one games is given, we document a shift in the distribution of choices. Moreover, we show that knowing opponent's last action from part-one games has a significant effect on shaping behaviors in the 11-20 games in part two. In the third subsection, we attempt to understand if the learning behavior observed in treatment one is a mere result of replicating the opponent's past choices. We investigate whether subjects behave differently in the set of 6-15 games, with and without information about their opponent's actions in part-one games. We show the plausibility of learning on the strategic levels through a set of different games. Subjects not only adjust their strategic levels according to the information given, but also show a tendency to rely on the information about their opponent's initial actions in the part-one games.

3.3.1 Baseline results of 11-20 Games

Figure 3.1 shows the density plots of the choices played in the part-one 11-20 games for the 156 subjects in our study. The black line represents AR2012 results, the red line corresponds to treatment one's part-one results and the green line corresponds to treatment two's part-one results. Although the density plots for our results do not completely overlap with AR2012, there is no significant difference between the choice distributions (p=0.33, χ^2test, with our two treatments pooled together). The choices made by our subjects across the two treatments have essentially identical distributions, as observed in the figure (p=0.24, χ^2 test). Table 3.1 shows the distribution of actions for part-one games.

Table 3.1: Relative frequencies of actions in part1 11-20 game

	11(%)	12(%)	13(%)	14(%)	15(%)	16(%)	17(%)	18(%)	19(%)	20(%)	N
AR(2012)	4	0	3	6	1	6	32	30	12	6	108
T1 (Part1)	3	1	4	4	8	7	14	28	21	11	360
T2 (Part1)	2	2	3	5	6	9	14	24	23	13	420

Due to the natural setup of the 11-20 game, every strategy could be considered as a level-*k* strategy for some *k*. Table 3.2 shows the results for strategic levels. Similar to AR2012 results, we document the vast majority of subjects (62 percent) chose the actions

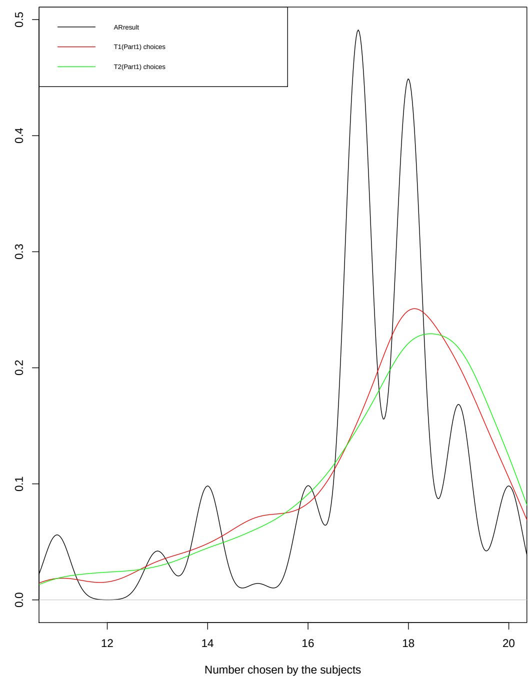

Figure 3.1: Choice densities for Part 1 11-20 Games

17-18-19, which correspond to 1-2-3 levels of reasoning respectively. For choices of smaller numbers, however, as Arad and Rubinstein (2012b) mentioned, choices of 11-16 do not necessarily mean that they are results of 4 to 9 rounds of iterated reasoning. We document about 8% of choices for both 15 and 16 respectively. We still consider them as level-4 (L4) and level-5 (L5) strategies as they represent a considerable amount of choices made by our subjects. For choices of 11-14, we pool them together and assign them to level-5+. Notably these choices were mostly made in earlier rounds (14% in round-1 and 7% in round 5). As our part-one games consist of five rounds of 11-20 games, we document a small order effect of the games. Although no feedback was given in-between rounds, subjects chose slightly larger numbers at later rounds [2].

Table 3.2: Relative frequencies of k-levels for part1 11-20 games

	L0(%)	L1(%)	L2(%)	L3(%)	L4(%)	L5(%)	L5+(%)	N
Part1(T1)	11	21	28	14	7	8	11	360
Part1(T2)	13	23	24	14	9	6	12	420

AR2012's 11-20 game serves as an excellent baseline for this study because the game design provides a very intuitive level-0 specification, and straightforward iterated reasoning pathway. If following the iterated best-responding strategy, the choices made for this game are relatively easy to be interpreted as steps of reasoning. This will become a crucial point for our subjects in part two of the experiment, where they observe their opponent's part-one actions. We document stable distributions of actions in part-one games across our two treatments, which make the two treatments comparable when we discuss the results later.

3.3.2 Results for learning within identical games

Under treatment one, subjects played ten more rounds of 11-20 games in part 2 of the experiment. They each played against five different opponents[3]. As subjects interacted with the same opponents twice, once with information about their opponent's part-one

[2] coefficient is 0.11 when regressing subject's choices on round number, p=0.03.

[3] Subjects played against their part-one opponents in the first two rounds. There is no repeated game effect. We check by regressing their choices in part 2 on a binary variable which indicates whether the opponent is the part-one opponent. The coefficient is -0.08, p=0.64.

actions and once without, we are able to do a side-by-side comparison of their decisions in this part of the experiment. Figure 3.2 presents the density plots of the subject's actions. The black line indicates choices without information (INFO0), and the red line indicates choices with information (INFO1). As shown in the figure, there is a significant shift in the distribution of choices. With additional information, subjects chose smaller numbers considerably more often[4].

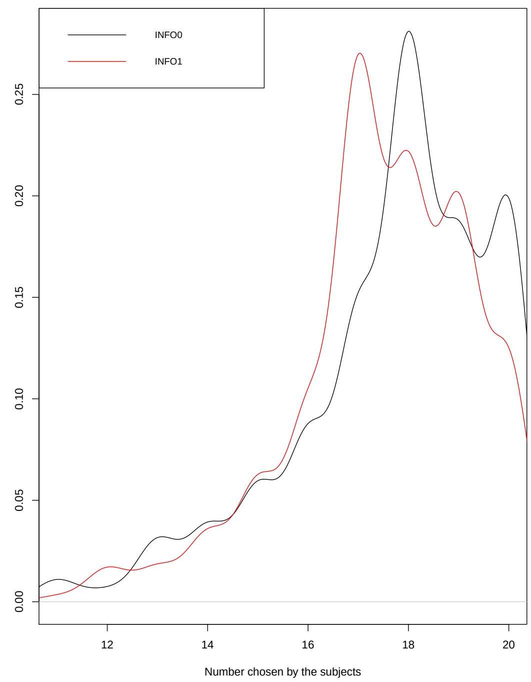

Figure 3.2: Choice densities for Part 2 (Treatment1) 11-20 Games

Table 3.3 conducts the level-*k* identification on the part-two 11-20 games and provides distributional contrasts for games with and without information. Notably, the L3 strategy was played 26% of the time when the opponent's information was revealed, as

[4]p= 0.02, Wilcoxon signed-rank test with paired samples.

compared to 14% when no information was revealed and 14% in the baseline games. This significant increased frequency of playing higher-level strategies provides some evidence on the effect of knowing the opponent's historical actions. In the contrast, lower-level strategies, especially the L0 strategy was favorable when no information was revealed, as it provides a guaranteed six dollars in payment in this experiment. It was played 20% under INFO0, while merely 12% under INFO1.

Table 3.3: Relative frequencies of k-levels for part2 11-20 games (T1)

	L0(%)	L1(%)	L2(%)	L3(%)	L4(%)	L5(%)	L5+(%)	N
INFO0	20	17	28	14	8	6	8	360
INFO1	12	19	21	26	9	6	7	360
Part2(T1)	16	18	24	20	9	6	8	720

Given that revealing the opponent's historical actions has an impact on the subject's strategies, we next investigate the underlying mechanism of this impact. Our subjects saw a list of five choices made by their opponents in the first part of the experiment, whenever this information was given to them. If their strategies in part-two games were dependent on the information, they would need to interpret the list of five numbers, and generate an assessment of their opponent's strategic levels based on those choices. Here we consider three metrics of the historical actions: the initial strategy, the most recent strategy (in this case, the strategy for the last round of part-one games), and the strategy with the highest frequency (mode). In Table 3.4, we provide corresponding analysis, regressing the strategic levels in our part-two 11-20 games on the interacting terms of the dummy treatment variable (info) and the strategic levels as reflected by the three metrics, controlling for observable characteristics. Column (1)-(3) examines the three metrics respectively. We find that the treatment effects for each of the metrics correlate significantly with part-two game behavior. However, when examining the relative explanatory power of these three metrics, in column (4) we find that knowing the opponent's strategic type based on their most recent actions has a significant effect on shaping the strategy for the part-two 11-20 games.

Table 3.4: Regression results for treatment1

	DV: Strategic levels for part2 11-20 games			
	(1)	(2)	(3)	(4)
Initial ∗ info	0.15** (0.06)	-	-	-0.01 (0.25)
Last ∗ info	-	0.27*** (0.08)	-	0.20** (0.10)
Mode ∗ info	-	-	0.23*** (0.07)	0.12 (0.11)
Initial	-0.05 (0.03)	-	-	0.01 (0.05)
Last	-	-0.08** (0.04)	-	-0.06 (0.05)
Mode	-	-	-0.07** (0.03)	-0.04 (0.06)
Info	-0.17 (0.22)	-0.34 (0.23)	-0.33 (0.22)	-0.45* (0.25)
Raven	0.08 (0.16)	0.08 (0.16)	0.08 (0.16)	0.08 (0.16)
Constant	2.65** (1.21)	2.68** (1.22)	2.69** (1.20)	2.70** (1.21)
Control	Yes	Yes	Yes	Yes
R^2	0.07	0.09	0.08	0.09
N	720	720	720	720

Notes: Control variables include age, gender, year in school, major and race. ∗ indicates $<$ 10% significance. ∗∗ indicates $<$ 5% significance. ∗∗∗ indicates $<$ 1% significance. Robust clustered individual standard error in parenthesis.

We next examine how subjects respond to their opponent's strategic types, based on their most recent actions. Table 3.5 provides the distribution of strategic levels for our subjects in the part-two 11-20 games in conjunction with their opponent's strategic types for the last game in part one of the experiment. Notably, for each opponent's type, subjects chose to respond with a strategy that was one level higher, the majority of the time (except when the opponent was L5). This could be viewed as choosing a strategy that was a direct best response to their opponent's types and was used 41.4% of the time. The second most used strategy was to replicate their opponent's actions, which was adopted 14.2% of the time. This part of the analysis is exploratory, but provides an interesting direction for future work to study the predictive accuracy of strategic levels in games with feedback about historical actions.

Table 3.5: Subject's type by part-two games with information about opponent's part-one's last game

		Opponent's type for part1's last game							
		L0	L1	L2	L3	L4	L5	L5+	Total
Part2 11-20 games	L0	1.9%	1.7%	1.9%	1.7%	0.6%	2.2%	1.9%	11.9%
	L1	**10.8%**	2.5%	3.3%	1.1%	-	0.8%	0.6%	19.1%
	L2	2.8%	**8.9%**	4.7%	1.9%	0.3%	1.7%	0.3%	20.6%
	L3	3.3%	3.3%	**15.8%**	1.9%	0.6%	0.8%	0.6%	26.3%
	L4	0.8%	0.8%	2.8%	**3.6%**	0.3%	0.3%	0.6%	9.2%
	L5	0.3%	-	1.7%	1.1%	**0.6%**	1.1%	1.1%	5.9%
	L5+	0.6%	-	1.4%	1.1%	0.6%	**1.7%**	1.7%	7.1%
	Total	20.5%	17.%	31.6%	12.4%	3%	8.6%	6.8%	

3.3.3 Results for learning across different games

Though information about the opponent's historical actions generates different responses in the set of 11-20 games, questions remain for whether the treatment effects were channeled through the revealed actions or the underlying strategic levels of those revealed actions. To answer this question, in the second treatment of the study, we implement a set of 6-15 money request games in the second part of the experiment. Again, similar to treatment 1, in part two of the experiment, subjects were given information about their opponent's actions in part-one 11-20 games fifty percent of the time. Since part-two games were different from 11-20 games in this treatment, subjects were not able to directly replicate or choose an action that is the best response to the historical actions that were revealed to them. If they were responsive towards the historical information, then they had to convert that information to something that is a mutual component to both games, which is the levels of strategic sophistication.

Figure 3.3 presents the density plots of the subject's actions in the set of 6-15 games. The black line indicates choices without information (INFO0), and the red line indicates choices with historical information (INFO1). As presented in the figure, there is a slight shift in the distribution of actions. With additional information, subjects chose smaller numbers slightly more often. However, there is no significant difference for the two choice distributions[5].

Similar to the 11-20 money request game, every strategy of the 6-15 game could be considered as a level-k strategy for some k. The L0-type is non-strategic, therefore choosing the salient number in the set of choices (in this case, 15), without considering other people's

[5] p=0.51, one tail Wilcoxon signed-rank test with paired samples

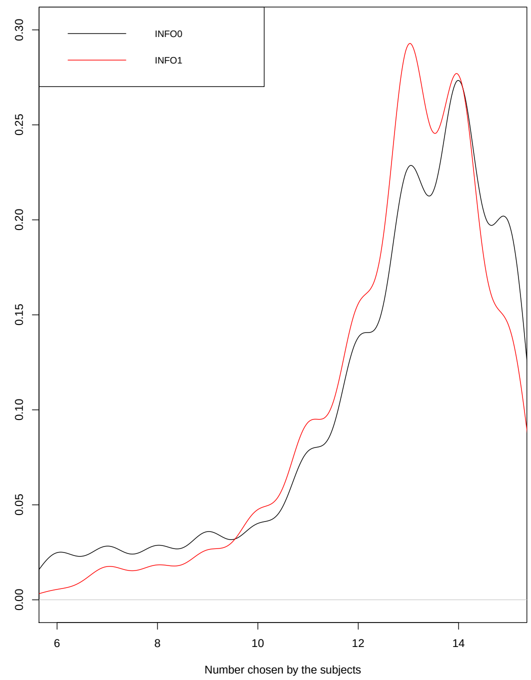

Figure 3.3: Choice densities for Part 2 (Treatment2) 6-15 Games

actions. An L1-type chooses an action that is the best response to the L0-type's expected action. In this game setup, the L1-type's strategy is 14. Following the level-*k* reasoning model, each L*k*-type chooses an action that is the best response to L$k-1$-type, which is one number smaller in this setup. Table 3.6 provides the level-*k* identification on the set of 6-15 games and provides distributional contrasts for games with and without the opponent's historical information. There was a slightly lower frequency of L0 strategy for the games with information (13% of the time, as compared to 19% under INFO0), and slightly higher frequency of L2 strategy (28% of the time, as compared to 21% under INFO0). However, in general, there is no significant difference in the distribution of *k*-levels.

Table 3.6: Relative frequencies of k-levels for part2 6-15 games (T2)

	L0(%)	L1(%)	L2(%)	L3(%)	L4(%)	L5(%)	L5+(%)	N
INFO0	19	26	21	13	7	4	11	420
INFO1	13	26	28	14	9	4	6	420
Part2(T2)	16	26	24	13	8	4	9	840

Although aggregate strategic levels do not differ much for the two comparison groups (INFO0 and INFO1), one may wonder if at the individual level, knowing the opponent's historical actions for 11-20 games has an effect on the subject's decisions in the 6-15 games. As discussed previously, if historical information of actions in a different game has an effect, it has to be related to the underlying strategic levels of those actions. Again, we consider three metrics, the revealed strategic levels of the initial strategy, that of the most recent strategy, and that of the mode strategy. Table 3.7 provides the corresponding analysis, regressing the strategic levels in 6-15 money request games on the interacting terms of the dummy treatment variable (info) and the strategic levels as reflected by the three metrics, controlling for observable characteristics. When examining the treatment effects of each metric respectively, column (1)-(3) shows that revealing each of the metrics has a significant effect on the 6-15 game behavior. This is similar to the results when the two sets of games are the same. This shows that at the individual level, knowing the opponent's historical actions of a different game helps shaping the behavior in a new game. Even though it is impossible to replicate the strategies directly, some beliefs about the opponent's strategic levels could be updated with the revealed information. Learning in fact happened at a more fundamental level, where the opponent's levels of strategic sophistication were studied, instead of strategies or actions at their face values.

Interestingly, when examining the relative explanatory power of the three metrics, column (4) shows that only information about the opponent's initial strategy significantly correlates with behaviors in the part-two games. This is very different from our previous results, where the opponent's most recent strategy plays a significant role in shaping behaviors. One explanation for the difference is that games in the second part of the experiments are different. In the first treatment, part one and part two have identical games. When assessing the opponent's levels of strategic sophistication, his most recent strategy has a higher weight than strategies from other rounds of the games, as it includes not only the information about the opponent's understanding of the game, but also some information about his experience level of the repeated games. However, when the game changes (as in the case of treatment two), initial responses to a game provide more information about a player's potential strategic levels in a new strategic context. Our results for the second treatment present the effects of revealing historical information of the opponent's past actions from another strategic context in shaping the strategic responses in a new environment. More importantly, the findings point to the importance of understanding which piece of historical information matters. When playing repeated games, perhaps the opponent's most recent strategy carries the most information. However, when interacting in a new strategic context, subjects care more about their opponent's initial responses to other games, and really utilize that information to understand their opponent's strategic levels when possible.

Table 3.7: Regression results for treatment2

	DV: Strategic levels for part2 6-15 games			
	(1)	(2)	(3)	(4)
Initial ∗ info	0.12^{***} (0.05)	-	-	0.14^{**} (0.19)
Last ∗ info	-	0.10^{*} (0.06)	-	0.07 (0.0)6
Mode ∗ info	-	-	0.10^{*} (0.06)	-0.07 (0.08)
Initial	-0.00 (0.03)	-	-	-0.00 (0.05)
Last	-	-0.02 (0.03)	-	-0.3 (0.04)
Mode	-	-	0.00 (0.03)	0.03 (0.07)
Info	-0.32^{*} (0.18)	-0.22 (0.16)	-0.22 (0.17)	-0.37^{*} (0.19)
Raven	-0.27^{**} (0.11)	-0.27^{**} (0.11)	-0.27^{**} (0.11))	-0.27^{**} (0.11)
Constant	2.41^{***} (0.76)	2.44^{***} (0.76)	2.39^{***} (0.75)	2.44^{***} (0.76)
Control	Yes	Yes	Yes	Yes
R^2	0.12	0.11	0.12	0.12
N	840	840	840	840

Notes: Control variables include age, gender, year in school, major and race. $*$ indicates $< 10\%$ significance. $**$ indicates $< 5\%$ significance. $***$ indicates $< 1\%$ significance. Robust clustered individual standard error in parenthesis.

3.4 Concluding Remarks

Knowing other people's historical information of past actions from a strategic game, can players learn about other people's levels of strategic sophistication? This is the central question addressed in this paper. In an online laboratory experiment, we provide two sets of strategic games, and examine if the subject's behavior in the latter games changes with information of their opponent's historical actions in the first set of games. When the two sets of games are identical, we find a significant shift in the distribution of choices when historical information is revealed. We also discover that information about the opponent's most recent strategy plays a significant role in determining strategies in the second set of games. When the two sets of games are different, we no longer find a significant difference in the distribution of choices. However, on the individual level, historical information still matters, but the relative importance of the information changes. When the games are different, subjects rely more on the information about their opponent's initial choices in the other game.

We show the plausibility of learning of strategic sophistication in this paper. Repeated interactions of individuals in different strategic contexts should no longer be considered as independent events. Even in the case when the contexts are drastically different, it is still possible for individuals to have some assessment about their counterpart's strategic levels through past interactions, and utilize that information in the new strategic interactions. In this paper, we provide an exploratory analysis of how individuals respond to the perceived strategic levels from historical actions. Future work could be done on improving the predictive accuracy under such circumstances. We utilize the level-k reasoning model in our design to study the question. However, we recognize that this behavioral model does not apply to all the strategic contexts. Learning of strategic sophistication with changing behavioral models remains an important open question for future studies.

3.5 Acknowledgement

Chapter 3, in part is currently being prepared for submission for publication of the material.

Appendix A

Appendix for Chapter 2

A.1 Estimation Strategy

The modified ring game offers unique action profiles for each player-type. With that benefit, we are able to implement the spike-logit maximum likelihood estimation to assign types for all of our subjects. Below we first describe the maximum likelihood function in detail, and then present the estimation results.

For each subject i, and set of eight modified ring games $g \in G$, we estimate a type k by defining a likelihood function $L_i(\varepsilon_i, \lambda_i, k \mid s_{ig})$. We assume subjects follow type-k strategy with probability $(1-\varepsilon)$ and with probability $\varepsilon \in [0,1]$, make mistakes that have a spike-logit error structure. When they make a mistake, they only play actions that are not type's predicted strategies with positive probability. With such error structure, the probability of selecting a type-inconsistent strategy is sensitive to the payoffs. This captures the attraction created by the focal strategies in our modified ring games. We assume type-1 subjects always choose the focal strategy, regardless of his player positions. Based on type-1's predicted strategy profiles, we are able to pin down a unique strategy profile for each type, as type-k only respond to $(k-1)th$-order payoffs. Following these unique strategy profiles, we further assume that type-k subjects have point belief about their opponent's actions at each player position.

Formally, let $S_{i,g}$ be the set of actions available to player i in game g, and s_g^k be type-k's predicted strategy in game g. Since player i's payoff is also determined by his immediate opponent's action, let $s_{j,g}$ be player i's immediate opponent's action in game g, and $f_g^k(s_{j,g})$ be type-k's belief in game g. λ_i is the sensitive parameter that will be estimated at individual level. Then $d_g^k(s_{i,g},\lambda_i)$ denotes type-k's error density when choosing $s_{i,g}$.

$$d_g^k(s_{i,g},\lambda_i) = \frac{\exp[\lambda_i \sum_{s_{j,g}\in S_{j,g}} u_{i,g}(s_{i,g},s_{j,g}) f_g^k(s_{j,g})]}{\sum_{s_{i,g}\in S_{i,g}\backslash s_g^k} \exp[\lambda_i \sum_{s_{j,g}\in S_{j,g}} u_{i,g}(s_{i,g},s_{j,g}) f_g^k(s_{j,g})]} \tag{A.1}$$

for $s_{i,g} \in S_{i,g}\backslash s_g^k$, and 0 elsewhere.

Let $I_i(s_i,g,k)$ be the indicator function that subject i follows type-k's strategy for game g. Then the likelihood of observing $s_{i,g}$ in game g, for type-k is given by

$$L_i(\varepsilon_i,\lambda_i,k \mid s_{i,g}) = (1-\varepsilon_i)I_i(s_i,g,k) + \varepsilon_i(1-I_i(s_i,g,k))d_g^k(s_{i,g},\lambda_i) \tag{A.2}$$

For the set of ring games, the likelihood function for observing $s_{i,G}$ for type-k is given by

$$L(\varepsilon_i,\lambda_i,k \mid s_{i,G}) = \prod_{g\in G} L_i(\varepsilon_i,\lambda_i,k \mid s_{i,g}) \tag{A.3}$$

In practice, we estimate ε_i in the sample through subject's actual error rate for each type. We jointly estimate k and λ_i by grid search over 405 possible values. Table A.1 shows the estimation results. The accuracy rate is obtained by counting the number of type-exact actions over the total number of type-predicted actions. For example, F1-type have 104 subjects, predicting a total of 832 actions with the eight modified ring games. There are 89% of the actions by these 104 subjects that comply with F1-type's prediction. Note that lower types have a generally higher accuracy rate. This is mainly due to the higher proportion of type-exact subjects that exist for these lower types. Figure A.1 shows the action profiles of the subjects for each F-types. There is a significant shift from large proportion of focal strategies observed in the lower types, to some non-focal strategies observed in the higher-order rational types.

Table A.1: Estimation results

	# of subjects	Accuracy rate
F1	104	89%
F2	60	81%
F3	23	64%
F4	13	61%

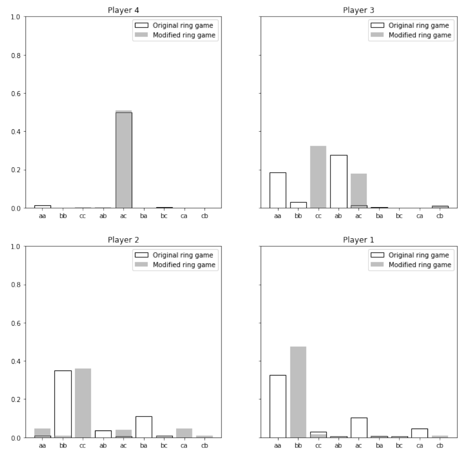

(a) type-F1

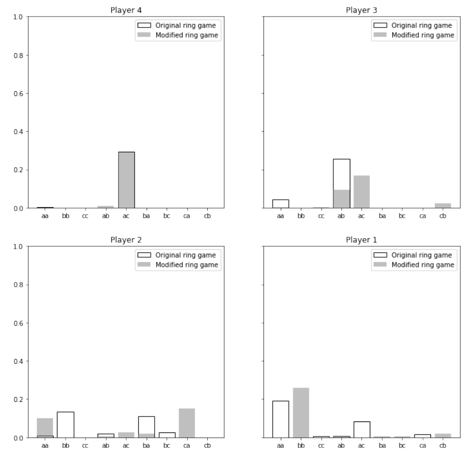

(b) type-F2

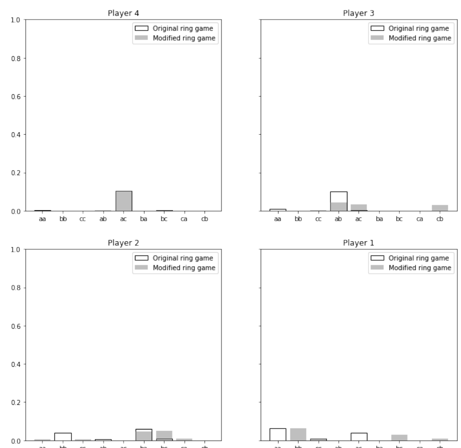

(c) type-F3

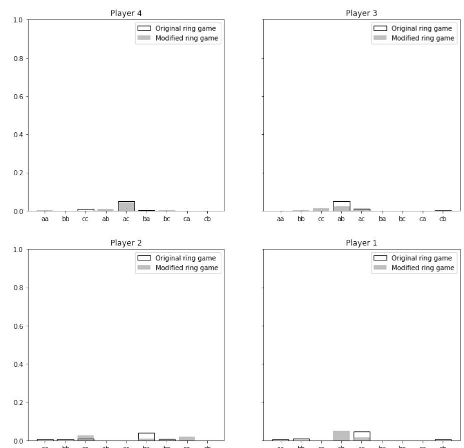

(d) type-F4

Figure A.1: Action profiles for each F-types

Table A.2 provides identification strategy under F-taxonomy for the original ring games, and table A.3 shows the results of estimation, as well as the individual correlation between F-levels across the original ring game and modified ring game. In the original ring games, the 'focal' strategies coincide with rationalizable strategies. As a result, there is not a clear separation of rationalizable action profiles and the action profiles centered around focal strategies or iterated best responses to focal strategies (table A.4 shows the correlation of F-levels and R-levels in the original ring game). The individual level correlation across the two games (Spearman correlation = 0.38) is thus similar to that using the R-taxonomy for the original ring game and the F-taxonomy in the modified game, (Spearman correlation = 0.35).

Table A.2: Predicted actions under rationality and assumptions about focality in original ring games

	Original Ring Games							
	P1		P2		P3		P4	
Type	G3	G4	G3	G4	G3	G4	G3	G4
F1	(a,a)		(b,b)		(a,a)		(a,c)	
F2	(a,a)		(b,b)		(a,b)		(a,c)	
F3	(a,a)		(b,a)		(a,b)		(a,c)	
F4	(a,c)		(b,a)		(a,b)		(a,c)	

Table A.3: Subject's types by modified ring game(F) and original ring game(F)

		Original RG				
		F1	F2	F3	F4	Total
Modified RG	F1	18.5%	21.5%	4%	8%	52%
	F2	5%	13%	6%	6%	30%
	F3	1%	4%	1.5%	5%	11.5%
	F4	-	0.5%	1.5%	4.5%	6.5%
	Total	24.5%	39%	13%	23.5%	

Notes: Spearman correlation: 0.38***

Table A.4: Subject's types by original ring game(R) and original ring game(F)

		Original RG					
		R0	R1	R2	R3	R4	Total
Original RG	F1	3%	21%	0.5%	-	-	24.5%
	F2	2%	3%	34%	-	-	39%
	F3	1.5%	-	1%	10.5%	-	13%
	F4	0.5%	0.5%	1.5%	2.5%	18.5%	23.5%
	Total	7%	24.5%	37%	13%	18.5%	

Notes: Spearman correlation: 0.84^{***}

A.2 Experimental Instructions

A.2.1 Experimental Instruction for Eye-gaze Test

Welcome to the experiment! You are about to participate in an experiment in the economics of decision-making. If you follow these instructions closely and consider your decisions carefully, you can earn a considerable amount of money, which will be paid to you in cash at the end of the experiment.

To ensure best results for yourself, please DO NOT COMMUNICATE with the other participants at any point during the experiment. If you have any questions, or need assistance of any kind, raise your hand and one of the experimenters will approach you.

There are two parts for today's experiment. The instruction for part 2 of the experiment will be given separately after the completion of part 1.

PART 1 of the Experiment

During part 1 of the experiment, you will answer 36 questions. For each question, you will see a picture of a set of eyes of a person on the screen. Below the picture, there will be four words. Choose and select the word that best describes what the person in the picture is thinking or feeling. You may feel that more than one word is applicable but please choose just one word, the word which you consider to be most suitable. Before making your choice, make sure that you have read all 4 words. You should try to do the task as quickly as possible but you will not be timed. If you really don't know what a word means, you can look it up in the definition handout in the following pages.

If you have understood the instruction for part 1 of the experiment, please click the button on the screen to proceed to the practice screen. You will have a chance to see what part 1 of the experiment looks like, before going to the actual questions.

A.2.2 Experimental Instruction for Ring Game

Welcome back for part 2 of the experiment!

The Basic Idea

You will 16 4-player games. In each of these games, you will be randomly matched with other participants currently in this room. We do our best to ensure that you and your counterparts remain anonymous. For each game, you will choose one of three actions. Each other participant in your game will also choose one of three actions.

Your Earnings

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

Player 2's Earnings

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

Player 3's Earnings

Player 3's Actions \ Player 4's Actions	j	k	l
g	20	12	8
h	6	8	18
i	0	16	4

Player 4's Earnings

Player 4's Actions \ Your Actions	a	b	c
j	10	12	8
k	6	20	18
l	16	4	0

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

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

Player 2, Player 3, and Player 4's earnings are listed in the other three tables. Player 2 may choose action d, e, or f, Player 3 may choose action g, h, or i, and Player 4 may choose action j, k, or l. Player 2's earnings depend upon the action he chooses and the action Player 3 chooses. Player 3's earnings depend upon the action he chooses and the action Player 4 chooses. Player 4's earnings depend upon the action he chooses and the action you choose.

For example, if you choose c, player 2 chooses e, player 3 chooses h, and player 4 chooses k, then you would earn 18 dollars, player 2 would earn 12 dollars, player 3 would earn 8 dollars, and player 4 would earn 18 dollars.

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

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

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

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

If you have understood the instruction for part 2 of the experiment, please click the button on the screen to proceed to the understanding quiz. There are 5 questions in the quiz. Answer those questions, and review the feedback given in the following screens. During the time, please raise your hand if you have any question regarding the quiz questions or the quiz answers. One of the experimenters will approach you and answer your questions.

A.2.3 Experimental Instruction for BR test

You have finished playing 16 games. You will now play additional 4 games with the computer. The computer will choose an action for player 2, player 3, and player 4 respectively. You will be given additional information about hte computer's (player 2, player 3, and player 4's) choices in the game. Please pay attention to those information.

One of these games may also be selected for payment at the end of the experiment. So you should treat each of these games like it will be the one determining your payment.

Please raise your hand if you have any question. Otherwise, click the 'continue' button below.

Your Earnings

Player 2's Actions

Your Actions	d	e	f
a	7	0	9
b	1	11	3
c	11	15	13

Player 2's Earnings

Player 3's Actions

Player 2's Actions	g	h	i
d	1	3	21
e	29	9	17
f	1	15	5

Player 3's Earnings

Player 4's Actions

Player 3's Actions	j	k	l
g	11	17	25
h	15	5	7
i	0	19	29

Player 4's Earnings

Your Actions

Player 4's Actions	a	b	c
j	15	9	7
k	5	3	21
l	11	29	19

Your earnings are given by the blue numbers. You may choose action a, b, or c. Your earnings will depend upon the action you choose and the action that Player 2 chooses.

In the above game, Player 2 (the computer), Player 3 (the computer) and Player 4 (the computer) have chosen the actions that give them the highest earnings.

Please choose your action.

Your Earnings

Player 2's Actions

Your Actions	d	e	f
a	15	9	7
b	5	3	21
c	11	29	19

Player 2's Earnings

Player 3's Actions

Player 2's Actions	g	h	i
d	11	15	13
e	7	0	9
f	1	11	3

Player 3's Earnings

Player 4's Actions

Player 3's Actions	j	k	l
g	1	3	21
h	29	9	17
i	1	15	5

Player 4's Earnings

Your Actions

Player 4's Actions	a	b	c
j	11	17	25
k	15	5	7
l	0	19	29

Your earnings are given by the blue numbers. You may choose action a, b, or c. Your earnings will depend upon the action you choose and the action that Player 2 chooses.

In the above game, Player 2 (the computer) has chosen the action ’d‘. Player 3 (the computer) and Player 4 (the computer) have chosen the actions that give them the highest earnings.

Please choose your action.

Your Earnings

Your Actions \ Player 2’s Actions	**d**	**e**	**f**
a	11	17	25
b	15	5	7
c	0	19	29

Player 2’s Earnings

Player 2’s Actions \ Player 3’s Actions	**g**	**h**	**i**
d	15	9	7
e	5	3	21
f	11	29	19

Player 3’s Earnings

Player 3’s Actions \ Player 4’s Actions	**j**	**k**	**l**
g	7	0	9
h	1	11	3
i	11	15	13

Player 4’s Earnings

Player 4’s Actions \ Your Actions	**a**	**b**	**c**
j	1	3	21
k	29	9	17
l	1	15	5

Your earnings are given by the blue numbers. You may choose action a, b, or c. Your earnings will depend upon the action you choose and the action that Player 2 chooses.

In the above game, Player 2 (the computer) has chosen the action that gives him the highest earnings. Player 3 (the computer) has chosen the action ’i‘. Player 4 (the computer) has chosen the action that gives him the highest earnings.

Please choose your action.

Your Earnings

Your Actions \ Player 2's Actions	d	e	f
a	1	3	21
b	29	9	17
c	1	15	5

Player 2's Earnings

Player 2's Actions \ Player 3's Actions	g	h	i
d	11	17	25
e	15	5	7
f	0	19	29

Player 3's Earnings

Player 3's Actions \ Player 4's Actions	j	k	l
g	15	9	7
h	5	3	21
i	11	29	19

Player 4's Earnings

Player 4's Actions \ Your Actions	a	b	c
j	11	15	13
k	7	0	9
l	1	11	3

Your earnings are given by the blue numbers. You may choose action a, b, or c. Your earnings will depend upon the action you choose and the action that Player 2 chooses.

In the above game, Player 2 (the computer) and Player 3 (the computer) have chosen the actions that give them the highest earnings. Player 4 (the computer) has chosen the action 'j'.

Please choose your action.

Appendix B

Appendix for Chapter 3

B.1 Experimental Instructions

B.1.1 Experimental instruction for Part 1 of the experiment

In part 1 of the experiment, you will be randomly matched with other players in this session. We do our best to ensure that you and your counterpart remain anonymous.

For each game, you will be asked to pick a number between 11 and 20. You will always receive the amount that you announce, in tokens. In addition:

-If you give a number that's exactly one less than your opponent, you receive extra 20 tokens.

Example:

-If you pick 17 and your opponent picks 19, then you receive 17 tokens and he receives 19 tokens

-If you pick 12 and your opponent picks 13, then you receive 32 tokens and he receives 13 tokens.

-If you picks 16 and you opponent picks 15, then you receive 16 tokens and he receives 35 tokens.

Each token is worth 30 cents.

Earnings:

You will earn a show-up payment of $2 for arriving to the experiment on time and participating. In addition to the show-up payment, one game may be randomly selected for payment at the end of the experiment. Every participant will be paid based on their actions and the actions of their randomly chosen counterparts in the selected game. Any of the games could be the one selected. So you should treat each game like it will be the one determining your payment.

Note:

To ensure the online experiment runs smoothly, each game in this experiment has a time limit of 2 minutes. If you do not make a choice before the time runs out, you will be logged out of this experiment. You will see a timer on the screen to show you how much time is left for each game. Thank you for your understanding.

If you have understood the instruction, please click the button on the screen to proceed to part 1 of the experiment.

Please only select the option below if the experimenter has finished reading the instructions for this part of the experiment.

B.1.2 Experimental instruction for Part 2 of the experiment (treatment2)

In part 2 of the experiment, you will be randomly matched with other players in this session. We do our best to ensure that you and your counterpart remain anonymous.

For each game, you will be asked to pick a number between 6 and 15. You will always receive the amount that you announce, in tokens. In addition:

-If you give a number that's exactly one less than your opponent, you receive extra 20 tokens.

Example:

-If you pick 12 and your opponent picks 14, then you receive 12 tokens and he receives 14 tokens.

-If you pick 7 and your opponent picks 8, then you receive 27 tokens and he receives

8 tokens.

-If you picks 11 and you opponent picks 10, then you receive 11 tokens and he receives 30 tokens.

You will also receive additional 5 tokens for each game. Each token is worth 30 cents.

In addition, in each game, you may or may not receive some additional information about your counterpart. Those information, whether given to you or not, will be clearly stated on your screen.

Earnings:

Together with games in part 1 of the experiment, one game may be randomly selected for payment at the end of the experiment. Every participant will be paid based on their actions and the actions of their randomly chosen counterparts in the selected game. Any of the games could be the one selected. So you should treat each game like it will be the one determining your payment.

Note:

To ensure the online experiment runs smoothly, each game in this experiment has a time limit of 2 minutes. If you do not make a choice before the time runs out, you will be logged out of this experiment. You will see a timer on the screen to show you how much time is left for each game. Thank you for your understanding.

If you have understood the instruction, please click the button on the screen to proceed to part 2 of the experiment.

Please only select the option below if the experimenter has finished reading the instructions for this part of the experiment.

B.1.3 Experimental instruction for Part 3 of the experiment

Welcome to Part 3 of the experiment!

For this part of the experiment, you will complete 5 questions. For every question, there is a pattern with a piece missing and a number of pieces below the pattern. You have to choose which of the pieces below is the right one to complete the pattern. In each case, one and only one of the pieces is the right one to complete the pattern.

For each question, please select your answers below the patterns. You will score 1 point for every right answer. You will not be penalized for wrong answers. You will have 5 minutes to complete the questions. You will see a timer on the top of your screen. You can submit your answers early. However, once the time is up, you answers will be automatically submitted. You will have the chance to know your score at the end of the experiment.

After completing the 5 questions, you will be directed to answer a few demographic questions. At the end of the experiment, we will show you the game number that has been randomly selected for payment for this session, as well as your choice and the choice of your counterpart for that game. Please allow up to 24 hours for the payments to arrive through Venmo.

If you have understood the instruction, please click the button on the screen to proceed to part 3 of the experiment.

Please only select the option below if the experimenter has finished reading the instructions for this part of the experiment.

Bibliography

Agranov, Marina, Andrew Caplin, and Chloe Tergiman, "Naive play and the process of choice in guessing games," *Journal of the Economic Science Association*, 12 2015, *1* (2), 146–157.

Alaoui, Larbi and Antonio Penta, "Endogenous Depth of Reasoning," *The Review of Economic Studies*, 10 2016, *83* (4), 1297–1333.

Allred, Sarah, Sean Duffy, and John Smith, "Cognitive load and strategic sophistication," *Journal of Economic Behavior & Organization*, 2016, *125*, 162–178.

Arad, Ayala and Ariel Rubinstein, "Multi-dimensional iterative reasoning in action: The case of the Colonel Blotto game," *Journal of Economic Behavior & Organization*, 11 2012, *84* (2), 571–585.

_ **and** _ , "The 11-20 Money Request Game: A Level- k Reasoning Study," *American Economic Review*, 12 2012, *102* (7), 3561–3573.

Bardsley, Nicholas, Judith Mehta, Chris Starmer, and Robert Sugden, "Explaining focal points: Cognitive hierarchy theory versus team reasoning," *Economic Journal*, 3 2010, *120* (543), 40–79.

Baron-Cohen, Simon, Sally Wheelwright, Jacqueline Hill, Yogini Raste, and Ian Plumb, "The " Reading the Mind in the Eyes " Test Revised Version : A Study with Normal Adults, and Adults with Asperger Syndrome or High-functioning Autism," Technical Report 2 2001.

Bernheim, B. Douglas, "Rationalizable Strategic Behavior," *Econometrica*, 7 1984, *52* (4), 1007.

Bosch-Domènech, Antoni, JOSÉ G. Montalvo, Rosemarie Nagel, and Albert Satorra, "One, two, (three), infinity, ... : Newspaper and lab beauty-contest experiments," *American Economic Review*, 2002, *92* (5), 1687–1701.

Bruguier, Antoine J., Steven R. Quartz, and Peter Bossaerts, "Exploring the Nature of "Trader Intuition"," *The Journal of Finance*, 10 2010, *65* (5), 1703–1723.

Burnham, Terence C., David Cesarini, Magnus Johannesson, Paul Lichtenstein, and Björn Wallace, "Higher cognitive ability is associated with lower entries in a p-beauty contest," *Journal of Economic Behavior and Organization*, 2009, *72* (1), 171–175.

Camerer, Cf, Th Ho, and Jk Chong, "A cognitive hierarchy model of games," *The Quarterly Journal of Economics*, 2004, *119* (3), 861–898.

Camerer, Colin and Teck Hua Ho, "Experience-weighted attraction learning in normal form games," *Econometrica*, 7 1999, *67* (4), 827–874.

Chen, Daniel L., Martin Schonger, and Chris Wickens, "oTree-An open-source platform for laboratory, online, and field experiments," *Journal of Behavioral and Experimental Finance*, 3 2016, *9*, 88–97.

Cooper, David J and John H Kagel, "Learning and transfer in signaling games," *Economic Theory*, 2008, *34*, 415–439.

Cooper, David J. and John H. Kagel, "The role of context and team play in cross-game learning," *Journal of the European Economic Association*, 2009.

Cooper, Russell, Douglas Dejong, Robert Forsythe, and Tom Ross, "Selection Criteria in Coordination Games: Some Experimental Results," *American Economic Review*, 1990, *80* (1), 218–33.

Costa-Gomes, Miguel A and Vincent P Crawford, "Cognition and Behavior in Two-Person Guessing Games: An Experimental Study," *American Economic Review*, 11 2006, *96* (5), 1737–1768.

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

Crawford, Vincent P. and Nagore Iriberri, "Level-K Auctions: Can a Nonequilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions?," 2007.

_ , Uri Gneezy, and Yuval Rottenstreich, "The power of focal points is limited: Even minute payoff asymmetry may yield large coordination failures," *American Economic Review*, 9 2008, *98* (4), 1443–1458.

Crawford, Vincent P Vp, Miguel a Costa-Gomes, and Nagore Iriberri, "Structural Models of Nonequilibrium Strategic Thinking: Theory, Evidence, and Applications," *Journal of Economic Literature*, 2013, *51* (1), 5–62.

Duffy, John and Rosemarie Nagel, "On the Robustness of Behaviour in Experimental 'Beauty Contest' Games," *The Economic Journal*, 1997, *107*, 1684–1700.

Dufwenberg, Martin, Ramya Sundaram, and David J. Butler, "Epiphany in the Game of 21," *Journal of Economic Behavior and Organization*, 2010.

— **, Tobias Lindqvist, and Evan Moore**, "Bubbles and experience: An experiment," 12 2005.

Fischbacher, Urs, "Z-Tree: Zurich toolbox for ready-made economic experiments," *Experimental Economics*, 6 2007, *10* (2), 171–178.

Fragiadakis, Daniel E, Daniel T Knoepfle, and Muriel Niederle, "Who is Strategic? *," Technical Report 2016.

Frederick, Shane, "Cognitive reflection and decision making," 9 2005.

Fudenberg, Drew and David Levine, "Learning in games," *European Economic Review*, 5 1998, *42* (3-5), 631–639.

— **and** — , *The Theory of Learning in Games*, Cambridge, MA: The MIT Press, 1998.

— **and David M. Kreps**, "Learning mixed equilibria," *Games and Economic Behavior*, 7 1993, *5* (3), 320–367.

Garcia-Pola, Bernardo and Nagore Iriberri, "Naivete and Sophistication in Initial and Repeated Play in Games," 2019.

García-Pola, Bernardo, Nagore Iriberri, and Jaromír Kovářík, "Non-equilibrium play in centipede games," *Games and Economic Behavior*, 3 2020, *120*, 391–433.

Georganas, Sotiris, Paul J. Healy, and Roberto A. Weber, "On the persistence of strategic sophistication," *Journal of Economic Theory*, 2015, *159*, 369–400.

Gill, David and Victoria Prowse, "Cognitive ability, character skills, and learning to play equilibrium: A level-k analysis," *Journal of Political Economy*, 12 2016, *124* (6), 1619–1676.

Gneezy, Uri, Aldo Rustichini, and Alexander Vostroknutov, "Experience and insight in the Race game," *Journal of Economic Behavior and Organization*, 8 2010, *75* (2), 144–155.

Grosskopf, Brit and Rosemarie Nagel, "The two-person beauty contest," *Games and Economic Behavior*, 2008, *62* (1), 93–99.

Ho, Teck-Hua and Xuanming Su, "A Dynamic Level- k Model in Sequential Games," *Management Science*, 2013, *59* (2), 452–469.

Ho, Teck Hua, Colin Cambrer, and Keith Weigelt, "Iterated Dominance and Iterated Best Response in Experimental "p-Beauty Contests"," *American Economic Review*, 9 1998, *88* (4), 947–969.

Ivanov, Asen, Dan Levin, Muriel Niederle, Source Econometrica, and No July, "Can Relaxation of Beliefs Rationalize the Winner's Curse?: An Experimental Study," *Econometrica*, 2010, *78* (4), 1435–1452.

Kneeland, Terri, “Identifying Higher-Order Rationality,” *Econometrica*, 2015, *83* (5), 2065–2079.

Levitt, Steven D., John A. List, and Sally E. Sadoff, “Checkmate: Exploring backward induction among chess players,” *American Economic Review*, 4 2011, *101* (2), 975–990.

Mehta, Judith, Chris Starmer, and Robert Sugden, “Focal points in pure coordination games: An experimental investigation,” *Theory and Decision*, 3 1994, *36* (2), 163–185.

Nagel, Rosemarie, “Unraveling in Guessing Games: An Experimental Study,” *The American Economic Review*, 1995, *85*, 1313–1326.

Pearce, David G., “Rationalizable Strategic Behavior and the Problem of Perfection,” *Econometrica*, 7 1984, *52* (4), 1029.

Rick, Scott and Roberto A. Weber, “Meaningful learning and transfer of learning in games played repeatedly without feedback,” *Games and Economic Behavior*, 3 2010, *68* (2), 716–730.

Stahl, Dale O., “Boundedly rational rule learning in a guessing game,” *Games and Economic Behavior*, 10 1996, *16* (2), 303–330.

__ **and Paul W. Wilson**, “Experimental evidence on players’ models of other players,” *Journal of Economic Behavior & Organization*, 12 1994, *25* (3), 309–327.

__ **and** __ , “On Players? Models of Other Players: Theory and Experimental Evidence,” *Games and Economic Behavior*, 7 1995, *10* (1), 218–254.

Zhao, Wanqun, “Cost of Reasoning and Strategic Sophistication,” *Games*, 9 2020, *11* (3), 40.