Laboratory Experiments on Belief Formation and Cognitive Constraints

Manuel Puente

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy
in the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY

2020

ABSTRACT

Laboratory Experiments on Belief Formation and Cognitive Constraints

Manuel Puente

In this dissertation I study how different cognitive constraints affect individuals' belief formation process, and the consequences of these constraints on behavior. In the first chapter I present laboratory experiments designed to test whether subjects' inability to perform more rounds of iterated deletion of dominated strategies is due to cognitive limitations, or to higher order beliefs about the rationality of others. I propose three alternative explanations for why subjects might not be doing more iterations of dominance reasoning. First, they might have problems computing iterated best responses, even when doing so does not require higher order beliefs. Second, subjects might face limitations in their ability to generate higher order beliefs. Finally, subjects' behavior might not be limited by cognitive limitations, but rather justified by their beliefs about what others will play. I design two experiments in order to test these hypothesis. Findings from the first experiment suggest that most subjects' strategies (about 66%) are not the result of their inability to compute iterated best responses. I then run a second experiment, finding that about 70% of the subjects' behavior come from limitations in their ability to iterate best responses and generate higher order beliefs at the same time, while for the other 30% their strategies are a best response to higher order beliefs that others are not rational. In the second chapter I study whether a Sender in a Bayesian Persuasion setting (Kamenica and Gentzkow,

2011) can benefit from behavioral biases in the way Receivers update their beliefs, by choosing how to communicate information. I present three experiments in order to test this hypothesis, finding that Receivers tend to overestimate the probability of a state of the world after receiving signals that are more likely in that state. Because of this bias, Senders' gains from persuasion can be increased by "muddling the water" and making it hard for Receivers to find the correct posteriors. This contradicts the theoretical result that states that communicating using signal structures is equivalent to communicating which posteriors these structures induce. Through analysis of the data and robustness experiments, I am able to discard social preferences or low incentives as driving my results, leaving base-rate neglect as a more likely explanation. The final chapter studies whether sensory bottlenecks, as oppose to purely computational cognitive constraints, are important factors affecting subjects' inference in an experiment that mimics financial markets. We show that providing redundant visual and auditory cues about the liquidity of a stock significantly improves performance, corroborating previous findings in neuroscience of multi-sensory integration, which could have policy implications in economically relevant situation.

Contents

List of Figures

Acknowledgements

First I would like to thank my supervisor, Alessandra Casella, for her guidance and infinite patience. I also want to thank the other members of my Dissertation committee: Mark Dean, Navin Kartik, Qingmin Liu, and Jacopo Perego. The members of the Experimental Lab and the Microeconomics Colloquium at Columbia University provided invaluable feedback, without which this work would have been of considerably lower quality. The third chapter of this dissertation was possible thanks to the my coauthors Zeyang Chen, Mark Dean, Harrison Hong, and Juanjuan Meng. Finally, I would like to thank Oskar Zorilla, both for his emotional support throughout this process, and for sharing his thoughts in countless economic discussions in the past six years.

Dedication

Para mi madre, María Isabel;

y para mi padre, Juan Carlos.

Chapter 1

Disentangling Higher Order Rationality

1.1 Introduction

A fundamental assumption in Game Theory is that players' actions survive the process of iterated deletion of dominated strategies. That is, players do not choose dominated strategies, do not believe that others will choose dominated strategies, do not believe that others believe that others will choose dominated strategies, and so on. Although this process can theoretically go on forever, experimental studies typically find behavior to be consistent with only a few rounds of iterated deletion of dominated strategies.[1]

Formally, we will use the concepts of rationality and higher order rationality. We say that an agent is rational, or first order rational, if she best responds to her beliefs about what others will play. An agent is second order rational if she is rational and she believes other agents are rational. That is, if she believes others will best respond to their beliefs about what others will play. Following this reasoning, we say an agent is kth order rational if she is rational and she believes others are

[1] Crawford et al. (2013), and Gill et al. (2001) provide surveys of the literature. Arad and Rubinstein (2012) and Kneeland (2015) are more recent examples.

$(k-1)$th order rational. The process of iterated deletion of dominated strategies is the empirical equivalent of higher order rationality: a kth order rational individual must play strategies that survive k rounds of iterated deletion of dominated strategies. Conversely, a choice profile that survives k rounds of iterated deletion is kth order rationalizable, meaning that it is consistent with the subject being kth order rational.

Most empirical studies focus on estimating rationality orders of individuals, and usually find that subjects perform no more than three or four rounds of iterated deletion. However, there is some evidence that these estimates are not rigid, and they respond to beliefs about the sophistication of their opponents (e.g. Alaoui and Penta, 2016, Palacios-Huerta and Volij, 2009). In light of these findings, an open question is whether subjects' estimated higher order rationality is indicative of cognitive limitations, or of their beliefs and higher order beliefs about what others will play.

In this paper we propose and experimentally test three explanations for why subjects are not doing more rounds of iterated deletion. First, subjects can be bounded by their ability to compute iterated best responses. Second, they may have problems generating higher order beliefs about rationality (e.g. "A believes that B believes that C is rational"). Finally, subjects might be stopping the process of iterated deletion of dominated strategies before hitting cognitive bounds because of what they believe others will play.

To illustrate these three alternative explanations, think of the following game. Ann, Ben, and Carlos have to simultaneously decide whether to go to the movies or to the stadium for the day. Their preferences are as follows. Ann only cares about spending the day with Ben. Ben only cares about spending the day with Carlos. Finally, Carlos wants to go to the movies, and doesn't care about who else goes with him. Where will Ann go? Notice that in this game, Ann's best action depends only on what Ben will do, which in turn may depend on what he believes Carlos will do. This game is solvable by the process of iterated deletion of dominated strategies. The fist step of the process is to notice that if Carlos is rational he will go to the movies, since it dominates going to the stadium. Next, if Ben is rational and believes Carlos is rational (i.e. second order rationality), he will also go to the movies, since that is his best response. Finally, if Ann is rational, believes Ben is rational, and believes that Ben believes that Carlos is rational (i.e. third order rationality), she will go to the movies too. This game has the same logical structure as the Ring Games used in this paper, with one player having a dominant strategy, a second player best responding to him, a third player best responding to the second player, etc.

Lets say Ann goes to the stadium, and we conclude that she is not third order rational. What can we say about her from her choice? Our first explanation for Ann's behavior is that she is not able to calculate what her best response to Ben's best response to Carlos' best response is. That is, she hit a bound on her ability to compute iterated best responses. In this case, we say that Ann is $BRI2$, meaning that she has a Best Response Iteration ability of two. Our second explanation is based on the idea that Ann's thinking process involves another cognitively demanding exercise: she must form higher order beliefs. That is, she must have beliefs about what Ben believes that Carlos will play. Following this argument, going to the stadium can be explained by Ann only being able to generate at most second order beliefs. In this case, we say that she has a Belief Formation ability of two, or $BF2$. The final explanation is that Ann's choice is not justified by her cognitive constrains, but because she believes that Ben believes that Carlos is irrational. If Ben believes that Carlos is going to the stadium, he will go to that place himself. Ann's best responses is then to go to the stadium herself.

Experiment 1 is designed to study whether Best Response Iteration ability can explain estimated higher order rationality. The identification strategy is based on a comparisons between two treatments. As a control, treatment E1-H estimates rationality orders by observing behavior in simultaneous Ring Games (see Kneeland, 2015). The second treatment, called E1-C, tests whether estimated rationality is caused by the number of iterations of best responses subjects can compute. This treatment has subjects play the same games as in E1-H, but with two modifications: we transform the games into sequential ones and we replace every opponent by a "rational" computer. In the example above, imagine Ann now chooses first, Carlos chooses second, and Ben observes Carlos decision and chooses third. Assume also that a computer decides what the best option for Carlos is and chooses for him, and another computer observes where Carlos went and then chooses the best option for Ben. By making the game sequential, Ann doesn't need to form second order beliefs since she knows that Ben's computer observes Carlos decision before playing. Our second modification generates the belief for Ann that others are rational. The key insight about this experiment is that the number of iterations of best responses required to find the iterated dominance solution remain the same across treatments for each Player. For Ann to find the correct decision, she must still compute what her best response, to Ben's best response, to Carlos optimal choice is. So, if we find that Ann goes to the stadium in the new setup, we know that her original decision was not due to her not being able to do more best response iterations. When comparing performance across the two treatments, we find that rationality orders can be justified by Best Responses Iteration ability for around 45% of our subjects.

Next, we move to using a broader range of cognitive limitations, by testing whether Best Response Iteration together with Belief Formation ability can better explain subjects rationality. Experiment 2 compares decisions between two new treatments. In both of them subjects play the same modified version of Ring Games where only one player -who has a dominant strategy- plays in a second stage after observing others' decisions, but the two treatments differ in the identity of that player. In the first treatment, E2-H, the second mover is another subject; while in the second treatment, E2-C, this player is a "rational" computer. Treatment E2-H plays a similar role as E1-H in Experiment 1 and provides baseline estimates of subjects' higher order rationality. Treatment E2-C tests whether these estimates are justified by cognitive limitations or higher order beliefs about rationality. To illustrate, imagine Ann plays the same game as before, except only Carlos plays second after observing everyone else's choice. Since for Carlos going to the movies is a dominant strategy, this modification should not affect Ann's choice when compared to the simultaneous game. Thus, if she goes to the stadium we can conclude that she is not third order rational. Suppose now that we have a computer play for Carlos in the same way as we did before. The introduction of a computer does not modify cognitive requirements: Ann still has to iterate best responses three times (i.e. $BRI3$) and form third order beliefs (i.e. $BF3$) to find the sequentially dominant action. But introducing the computer may affect beliefs about Carlos' rationality. Thus, if Ann changes her choice and goes to the movies, we can say that her rationality order is explained by her belief that Ben believes that Carlos is irrational. If, on the other hand, she doesn't change her decision and still goes to the stadium, we infer that she is either not able to iterate best responses three times or to form third order beliefs. Our results from this experiment are that for around 34% of our subjects, baseline estimates of higher order rationality can be explained by higher order beliefs of irrationality. This leaves 66% of our subjects having low rationality orders because of cognitive difficulties.

This paper is a contribution to the experimental literature on iterated deletion of dominated strategies and higher order rationality. From this literature, Kneeland (2015) is most closely related to this work. She utilizes Ring Games and a novel identification strategy to estimate rationality order. The games and estimation strategy used in treatments E1-H and E2-H in this paper are borrowed from that paper. Kneeland finds that most subjects have low orders of rationality, with 66% being second order rational or lower, and 88% being third order rational or lower. Our study is also related to the Level-k model (Nagel, 1995; Stahl and Wilson, 1994) and the Cognitive Hierarchy model (Camerer et al., 2004), which on top of assuming that players have finite rationality orders, impose restrictions on the distribution of beliefs and higher order beliefs about others' rationality. Crawford et al. (2013) provide a detailed review of empirical studies for this topic.

These initial papers had a follow-up literature that studies the relation between the sophistication of strategies chosen by individuals and their cognitive constrains. Friedenberg et al. (2017) use data from Kneeland (2015) to distinguish between Rationality Bounds and Cognitive Bounds. Rationality bounds correspond to rationality orders: "Ann has a rationality bound of level m if she is rational, thinks Bob is rational, thinks Bob thinks she is rational, and so on, up to the statement that includes the word "rational" m times, but not further". Cognitive bounds are defined similarly, but without imposing rationality: "Ann has a cognitive bound of k if she has some belief about how Bob will play the game, if she believes Bob has some belief about how she will play, and so on, up to the statement that includes the word "belief" k times, but no further". Our analysis differs in two ways. First, we study two types of cognitive bounds: Best Response Iteration ability and Belief Formation ability. Second, their analysis focuses on the idea that subjects might have low order beliefs about rationality but high order beliefs about other type of behavior (e.g. heuristics), while we are interested on how cognitive bounds affect higher order rationality.

Jin (2016)'s motivating question and experimental strategy are closer to the present paper. Jin alters Kneeland's Ring Games by making them sequential: in each treatment, one of the players moves second, after observing the choices made by other players. This design is similar to the one used in our second experiment, except for the identity of the player who moves second and our use of "rational" computers. The main difference between both approaches is that our objective is to alter higher order beliefs about rationality for all subjects, while Jin changes first and second order beliefs about rationality only for subjects that are at least third order rational. Nevertheless, the two papers are closely related, and her results should be seen as complementary to our second experiment. She finds that about half of the subjects classified as third and fourth order rational are responding to their beliefs about others' strategies, rather than having reached their own cognitive bound.

The idea that individuals' apparent strategic sophistication depends on their beliefs about others' play is explored in Alaoui and Penta (2016), where agents' level of interactive thinking is modeled to respond to costs and benefits. The cost of performing more rounds of belief formation depends on players' own cognitive limitations, while the benefit depends on the payoffs of the game and their beliefs about the sophistication of other players. Alaoui and Penta test experimentally the comparative statics properties of their model. Cognitive bounds studied in this paper can be interpreted in Alaoui and Penta's model as the maximum steps of reasoning subjects are willing to do, while fixing payoffs across games and having subjects believe their opponents are "more sophisticated than themselves".[2]

[2]Brandenburger and Xiaomin (2015) present an epistemic approach of how individuals might

In addition to Alaoui and Penta (2016), other recent papers propose experimental designs that relate subjects' apparent sophistication to the perceived cognitive ability of their opponents (e.g Gill-Prowse, 2016; Georganas et al., 2015; and Palacios-Huerta and Volij, 2009, Agranov et al., 2012). These papers consistently find that some subjects play more sophisticated strategies when they believe their opponents to be more intelligent, but they do not attempt to estimate or explain cognitive bounds for subjects that have reached them.

1.2 Ring Games

Keeland (2015) presented a novel setup which allows to estimate higher order rationality without the need to impose strong restrictions on the behavior of irrational players, as is the case in *level-k* models. Kneeland studies a family of dominance solvable simultaneous games, called Ring Games. These games are similar to the example used in the previous section, except that Carlos' payoffs may depend on what Ann chooses.

A N-player Ring Game is a simultaneous game in which the payoff of each player is affected by the next player's action, and the playoff for the last player is affected by the first player's action. Figure 1.1, RG-1, presents an example of a 3-player Ring Game. The first matrix represents the payoffs of Player 1, which depend

build higher order beliefs, drawing on neural evidence. They find that the number of cases individuals have to consider, and possibly store, grows exponentially with the belief order. This result could explain why subjects seem to stop at a small finite orders.

on her choice, either a or b, and the choices made by Player 2, either c or d. Player 2's payoff, represented in the second matrix, depends on her choice, either c or d, and the choice made by Player 3. Finally, Player 3's payoff, represented in the third matrix, depend on her choice and Player 1's choice. In order to be able to identify higher order rationality, two additional restrictions are imposed on the payoff matrices: the last player has a dominant strategy, and each action of the other players is a best response to some beliefs. In the example, a (b) is a best response to c (d) for Player 1; c (d) is the only best response to e (f) for Player 2; and e dominates f for Player 3. Notice that by labeling the top action (i.e. a, c, and e) as "going to the movies" and the bottom action (i.e. a, c, and e) as "going to the stadium", RG-1 becomes a representation of the game used in the previous section.

P1's payoff

		P2: c	P2: d
P1	a	1	0
	b	0	1

P2's payoff

		P3: e	P3: f
P2	c	1	0
	d	0	1

P3's payoff

		P1: a	P1: b
P3	e	1	1
	f	0	0

Figure 1.1: RG-1

The game is solvable by a process of iterated deletion of dominated strategies, and each step of this process has a direct relation to different rationality orders. First, notice that if Player 3 is rational she will choose e, since it dominates f. Second, as long as Player 2 is rational and believe Player 3 is rational (i.e. second order rationality), Player 2 will believe that Player 3 will choose e, and best respond by choosing c. Third, as long as Player 1 is third order rational, she will choose a as this is her only best response to Player 2 choosing c.

Choice data from subjects playing RG-1 alone would not be enough to correctly estimate individuals' order of rationality because higher orders nest lower ones, and thus different orders can lead to the same behavior. For example, while only a is consistent with third order rationality, it is also consistent with lower orders. Player 1 could be irrational and choose this action, or believe that Player 2 is irrational and plays c. In other words, observing b would allow us to conclude that Player 1 is not third order rational, but observing a would not be enough to conclude that the opposite is true. To deal with this identification problem Kneeland includes a second game, represented in Figure 1.2 (RG-2). This game is identical to the first one, except that the payoff matrix of Player 3 has been relabeled so that now f is the dominant strategy. By changing what the rational strategy for Player 3 is, we also change the second order rational strategy for Player 2 from c to d, and the third order rational strategy for Player 1 from a to b. Identifying rationality orders is now possible assuming the following exclusion restriction, borrowed from Kneeland (2015):

ER (K): *Subjects satisfying only lower order rationality do not respond to changes in higher order payoffs.*[3]

It is now possible to distinguish, say, third order and second order rationality by observing choices in both games: while a third order rational player must play a when choosing as Player 1 in the RG-1 and b in RG-2, a second order rational player should not react to changes in the payoffs of Player 3, since she does not take them

[3] See Kneeland (2015) for an in depth discussion about the validity of this assumption.

		Position		
		Player 1	Player 2	Player 3
Rationality	*R*1	(a, a) or (b, b)	(c, c) or (d, d)	(e, f)
	*R*2	(a, a) or (b, b)	(c, d)	(e, f)
	*R*3	(a, b)	(c, d)	(e, f)

Table 1.1: Rationality profiles
Note: Predicted strategies for each rationality order and Player, for (RG-1,RG-2). Source: Kneeland (2015)

into account when trying to predict Player 2's action. Thus, if a subject is not third order rational we expect her to play either a in both games or b in both games.

P1's payoff

		P2	
		c	d
P1	a	1	0
	b	0	1

P2's payoff

		P3	
		e	f
P2	c	1	0
	d	0	1

P3's payoff

		P1	
		a	b
P3	e	0	0
	f	1	1

Figure 1.2: RG-2

To summarize, individual rationality orders can be estimated by observing choices of subjects playing as each player in both RG-1 and RG-2, and looking for the Player number at which they switch from choosing the same action in both games to choosing the sequentially rational strategy. Choice profiles for first, second, and third order rationality are shown in Table 1.1, with a pair (x, y) representing what individuals will play in RG-1 and RG-2, respectively. As an example, a second order rational subject is expected to play the rational actions as Player 1, (e, f); the second order rational actions as Player 2, (c, d); and a constant action as Player 1, either (a, a) or (b, b).

Throughout this paper, we are going to use variations of the same 4-player

Ring Games used by Kneeland (2015). These games, which we will call G-1 and G-2, are represented in Figure 1.3. They follow the same structure as RG-1 and RG-2, with Players 1 to 3 having no dominated strategy, and Player 4 having j and l as the dominant strategy in G-1 and G-2, respectively. The addition of an extra player allows us to estimate up to fourth order rationality, while having an extra strategy reduces the probability of incorrectly assigning rationality orders.

The sequentially rational strategies for these games are (a, c) for Player 1, (e, d) for Player 2, (g, h) for Player 3, and (j, l) for Player 4. Assumption $ER(K)$ implies that a kth order rational subject will play the same action in G-1 and G-2 when playing as Player 1 to Player $(4-k)$, and she will play the sequentially rational strategies when playing as Player $(5-k)$ to Player 4. For example, a second order rational subject will play (a, a), (b, b), or (c, c) as Player 1; (d, d), (e, e), and (f, f) as Player 2; (g, h) as Player 3; and (j, l) as Player 4.

G-1

P1's Payoff

P2

		d	e	f
	a	8	20	12
P1	b	0	8	16
	c	18	12	6

P2's Payoff

P3

		g	h	i
	d	14	18	4
P2	e	20	8	14
	f	0	16	18

P3's Payoff

P4

		j	k	l
	g	20	14	8
P3	h	16	2	18
	i	0	16	16

P4's Payoff

P1

		a	b	c
	j	12	16	14
P4	k	8	12	10
	l	6	10	8

G-2

P1's Payoff

P2

		d	e	f
	a	8	20	12
P1	b	0	8	16
	c	18	12	6

P2's Payoff

P3

		g	h	i
	d	14	18	4
P2	e	20	8	14
	f	0	16	18

P3's Payoff

P4

		j	k	l
	g	20	14	8
P3	h	16	2	18
	i	0	16	16

P4's Payoff

P1

		a	b	c
	j	8	12	10
P4	k	6	10	8
	l	12	16	14

Figure 1.3: G-1 and G-2

1.3 Experiment 1

The purpose of Experiment 1 is to test whether subjects' estimated rationality is a consequence of cognitive constrains associated with their ability to compute iterated best responses. The identification strategy is based on comparing behavior in two treatments. The first treatment, called the Human Treatment (E1-H), is a replication of Kneeland's experiment.

Each subject plays G-1 and G-2 four times without feedback, each time as a different Player. Subjects are classified as $R1$, $R2$, $R3$, or $R4$ if their choice profiles are an exact match to a rationality order. Additionally, if the profile is not an exact match but it is one mistake away from an order, the subject is also assigned to it. If the action profile is one mistake away from two rationality orders, the lowest one is used. Allowing for one mistake, 95% of all possible combinations of action for each player and game do not satisfy any of these requirements, which suggests that classifying random behavior to an order is unlikely. Our final conclusion does not change if we instead use a strict classification, although the proportion of subjects estimated to be either irrational or not satisfying $ER(K)$ (i.e. $R0$ classification) increases from 6.25% to 25.78%. [4]

In order to test whether these estimates of higher order rationality are the result of cognitive bounds limiting subjects ability to compute iterated best response,

[4] These same arguments are valid for all estimations in this paper. Kneeland (2015) has a more detail discussion of this topic.

we designed a second treatment that does away with higher order beliefs and manipulates subjects' first order beliefs about opponents' rationality, without reducing the need to calculate iterated best responses. For this treatment, called the Computer Treatment (E1-C), we modify G-1 and G-2 in two ways. First, we use a sequential version of these games, with each subject always as the first mover. Second, subjects' opponents are replaced by "rational" computers who always best respond to previous choices.

As a clarifying example, suppose a subject is assigned to play as Player 2 in either G-1 or G-2. The timing of the game she would play is presented in Figure 1.4.

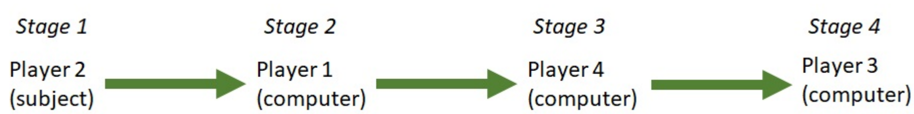

Figure 1.4: Timing of a subject playing E1-C as Player 2

The subject, which in this case is playing as Player 2, is always the first mover. In Stage 2, the computer "observes" her choice and then chooses as the player whose payoff depends on Player 2's choice, which in this case is Player 1. Next, the computer chooses as Player 4, since Player 1's action affects Player 4's payoff. Finally, the computer chooses as Player 3, since its payoff depend on Player 4's action. The transformation of the game into a sequential one eliminates the need for higher order beliefs since after the subject has chosen, each other player can see what the previous player has chosen. In the simultaneous game played in E1-H, for Player 2 to choose the iterated dominance solution (e, d), she must believe that Player 3 believes that

Player 4 is rational. In this sequential game played in E1-C, believing that both Player 3 and Player 4 are rational is enough.

The second modification fixes subjects' first order beliefs about others' rationality. Subjects are told that computers will observe choices made by previous players and always "choose the option that maximizes their earnings", a statement made concrete via examples and a short quiz. The use of computers in this way is meant to generate first order beliefs of rationality.

Given these two modifications, behavior cannot be driven by subjects' inability to form higher order beliefs or by them believing that their opponents are not rational. But notice that the number of iterations of best responses required to find the backward induction solution in this game is the same as the one required to find the sequentially rational solution in E1-H.

In the example, Player 2 still needs to iterate best responses at least 3 times. That is, if the subject plays (e, d) as Player 2, we can conclude that she is able to compute her best response to Player 3's best response to Player 4's best response. Using the same argument, subjects playing the equilibrium solution as Player 1, Player 2, and Player 4 must be able to iterate best responses at least 4, 2, and 1 times, respectively. We formalize this argument by defining the concept of Best Response Iteration ability:

$BRIk$: *We say that a subject has a Best Response Iteration ability of k (i.e. BRIk) if she is able to compute what is her best response to a player's best response to a player's best response... and so on, up to when the term "best response" is repeated k times, but not further.*

In general, the action profile in E1-C for a subject that can compute k, but not $k + 1$, best response iterations ($BRIk$) will be the same as a kth order, but not $k + 1$th order rational subject (Rk) in E1-H.[5] The exclusion restriction for this experiment is:

ER (E1): *if subjects' higher order rationality estimates from a Ring Game are determined solely by the number of best response iterations they can compute, eliminating higher order belief requirements by making the game sequential and having them play against rational computers will not change their choice profile.*

Using the same method as for higher order rationality, we will classify subjects as $BRI1$, $BRI2$, $BRI3$, and $BRI4$ if they are a perfect match with an action

[5]Note that we are using a conservative approach to estimating $BRIk$: although for Player i in an N-player Ring Game the backward induction solution requires at most N best response iterations, the minimum requirement is the $N + 1 - i$ iterations. In the example, Player 2 might look for her best response to Player 3's best response to Player 1's best response, which implies 3 iterations; or find her best response to Player 3's best response, which implies 2 iterations.

profile; or if they have no perfect match but they are one mistake away from one.

Implementation

The data was collected in the Columbia Experimental Lab for Social Science (CELSS) at Columbia University. 128 subjects were recruited through campus and participated in 8 sessions of 16 subjects each. For half of the sessions, subjects played the simultaneous treatment E1-H first and the sequential treatment E1-C second, while for the other half the order of treatments was reversed. Each subject played a total of 16 rounds without feedback (two as each of the 4 Players in G-1 and G2, for each treatment). The order in which they were assigned to each role was randomized. Before starting each treatment, instructions were read out loud and subjects had to complete a short quiz in order to better help them understand the game. This information is included in the appendix.

1.4 Results

Figure 1.5 presents estimated rationality orders in Kneeland (2015) and treatment E1-H. As expected, we find a similar distribution of estimated orders. A Kolmogorov-Smirnov two-sided test is unable to reject the null hypothesis of equal distributions with a p-value of 0.41.[6] In our experiment, 19% of the subjects are classified as

[6] Bootstraps were used to calculate p-values for the Kolmolorov-Smirnov test because of the discreteness of the distribution. The null hypothesis is that all data came from the same distribution. The alternative hypothesis is that data for T_H and Kneeland came from different, unknown distributions. A description of the method can be found in Abadie (2002). All KS tests described later in this paper follow the same methodology.

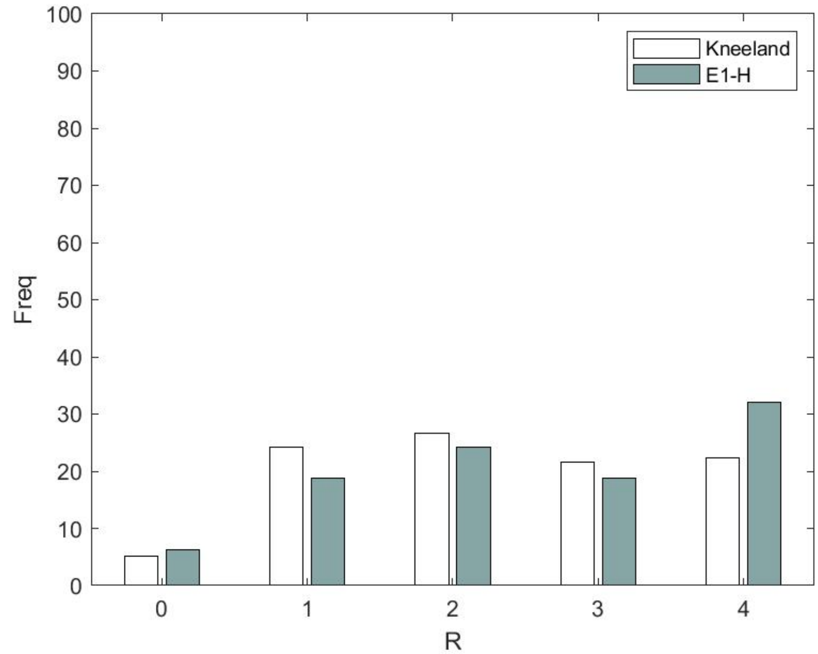

Figure 1.5: Frequencies of R and BRI

rational but not second order rational ($R1$), 24% as second but not third order rational ($R2$), 19% as third but not fourth order rational ($R3$), and 32% as fourth order rational or higher ($R4$). The remainder 6% did not fit into any rationality order and are assigned as irrational ($R0$).

Figure 1.6 reports the proportions of subjects classified to each rationality order in E1-H and their best response iteration ability in E1-C. The difference in behavior across treatments is striking: 9% of subjects cannot iterate best responses more than once, 4% can do two iterations, 6% can do three iterations, and 74% can iterate best responses at least four times. This suggest that, for a big portion of our subjects pool, Best Response Iteration ability by itself does not explain higher order rationality.[7]

In order to better understand the data, Figure 1.7 presents individual estimates of higher order rationality and best response iteration ability for every subject. Within-subject analysis has the added difficulty of having to deal with learning derived from playing a treatment second. Subjects represented by a circle participated in sessions in which E1-H was played first and E1-C was played second, while those represented by a cross participated in sessions with the reversed order. Even though subjects received no feedback, learning effects are present in our data, although this is not surprising given that the experimental design involves subjects playing similar games across treatments. When comparing orders across sessions, we find that on average higher order rationality estimates increases by .21 if E1-H is played second, while Best Response Iteration ability estimates increases by 0.43 if E1-C is played second. This means that sessions where E1-H was played first are biased towards accepting the hypothesis that subjects can iterate more best responses than their higher order rationality, while sessions where E1-C was played first are biased towards rejecting it. Learning should also be taken into account when trying to understand observations that lay below the 45° line in Figure 1.7.[8] These observations cannot be explained by our theory, since they refer to subjects whose BRI is at least one order lower than their R, and being able to iterated best responses k times is a necessary

[7]The p-value of a Kolmogorov-Smirnov test for these distributions is less than 1 percent.

[8]In this figure, we added random noise to each observation for visibility. Observations "below the 45° line" refer to those for whom BRI is lower than R.

condition for being kth order rational. But notice that all these observations are subjects that were estimated to be irrational in E1-C ($BRI0$), or they played E1-H after E1-C, suggesting that learning effects are responsible for these counterintuitive cases.

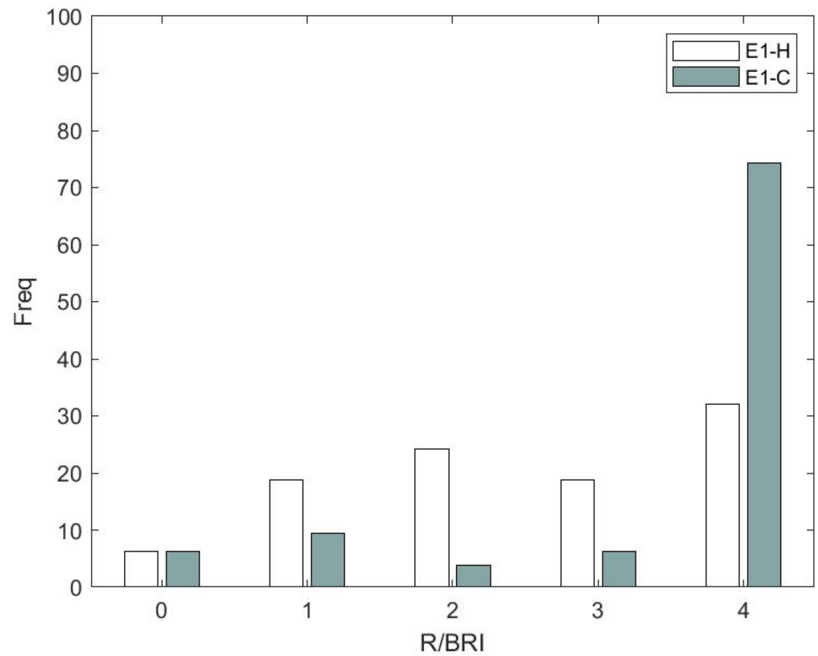

Figure 1.6: Frequencies of R and BRI

By looking at sessions where either E1-H or E1-C was played first, we can obtain upper and lower bounds on the proportion of subjects whose higher order rationality can be explained by their ability to iterate best responses. To calculate these bounds we exclude subjects who are estimated to be $R0$. These subjects are either not rational, or have a decision process not consistent with higher order rationality and our exclusion restriction $ER(K)$. We also exclude subjects classified as both $R4$ and $BRI4$, since we cannot discard the possibility that their higher order

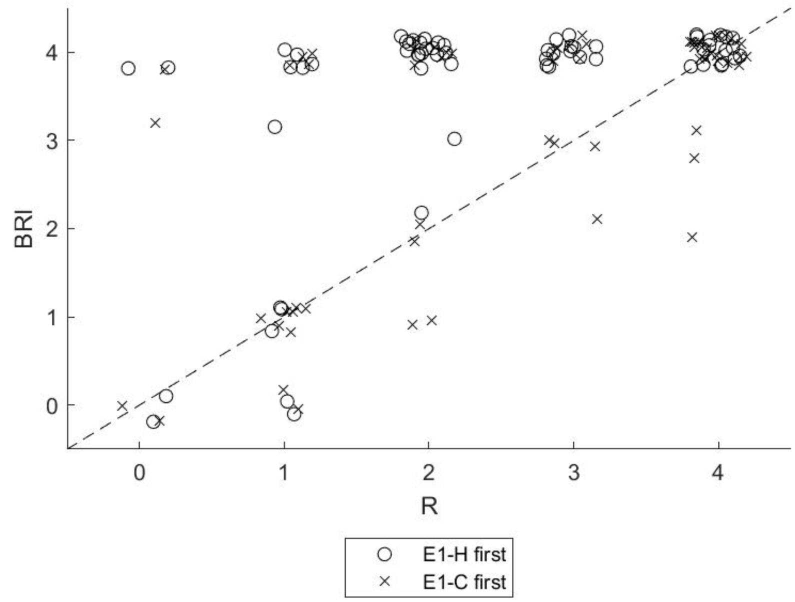

Figure 1.7: Individual estimates of R and BRI in Experiment 1
Note: subjects are represented by circles if they played E1-H first and E1-C second, and by crosses if they played E1-C first. A random noise $\epsilon \in [-0.1, 0.1]$ was added to estimates in order to differentiate subjects with the same R and BRI.

rationality and ability to iterate best responses are even higher. We conclude that limits on the ability to compute iterated best responses, by itself, can explain higher order rationality for 14% to 53% of subjects.

Another way to deal with learning is to debias data from treatments that were played second. Take for example a subject i that played E1-H first and then E1-C. Her Best Response Iteration ability when playing E1-C second, $BRI_{2,i}$, confounds the counterfactual she would have had if she had played this treatment first, $BRI_{1,i}$, with learning effects, x_i.

$$BRI_{2,i} = BRI_{1,i} + x_i$$

where x_i is assumed to follow a Poisson process $\mathcal{P}(\lambda)$ independent across subjects. An interpretation of this process is that subjects discover insights about the game which increase their BRI by one while listening to the instructions for the next treatment, and these insights have a fixed probability of happening at each point in time that is independent of when the last insight was discovered. Poisson processes have been used to model dynamic learning though information acquisition (e.g. Wald (1997), Che and Mierendorff (2019)). In this literature, players receive signals which help them learn what the state of the world is or which of various decisions to take. More closely related, Keller and Rady (2015) study learning in a model where agents obtain lump-sum payoffs which arrive with according to a Poisson process if they decide to engage in experimentation.[9] In hour story, these lump-sum payments would resemble possible payoffs subjects get from increasing their understanding of the game (i.e. higher BRI and R), similar to Keller and Rady's example where academics advance their knowledge of a topic though breakthroughs in their research.

Although we assume the learning process to be i.i.d. across subjects, we allow it to vary across treatments since, as shown above, average BRI increases by more if E1-C was played second than average R if E1-H was played second. Taking this differential learning effect across treatments into account, we assume that if

[9] See also Hörner and Skrzypaz (2016) for a review of the Strategic Experimentation literature.

subject i plays treatment $\tau \in \{\text{H,C}\}$ second she improves her understanding of the game, which results in an increase in her R/BRI by $x_i \sim \mathcal{P}(\lambda_\tau)$. The parameter λ_τ measures the expected number of insights (i.e. total increase in R/BRI) that a subject will have when playing treatment τ second.

We start by calibrating λ_τ to match the first moment of the distribution of rationality orders and best response iteration, so that average R (BRI) in sessions where $\tau = H$ (C) was the second treatment played equal average R (BRI) in sessions where it was played first, which are unbiased, plus expected learning. In our previous example where i played E1-C second, we would have

$$\hat{\lambda}_{\text{C}} = E(x_i) = \overline{BRI}_2 - \overline{BRI}_1$$

where $\overline{BRI}_1$ and $\overline{BRI}_2$ are average BRI across subjects that played treatment E1-C first and second, respectively. Next, we run 10000 Monte-Carlo simulations of x_i and debias i's $BRI_{2,i}$. Finally, we record for each iteration whether i's (debiased) ability to compute iterated best responses is greater than her rationality, and calculate how often this happened.

Summarizing, a step-by-step explanation of the process for debiasing Best Response Iteration ability for subject i, who played E1-C after E1-H, would follow these steps:

1. *Estimate $BRI_{1,j}$ for all subjects j that played treatment E1-C first, and compute the sample average, $\overline{BRI}_1$.*

2. *Estimate* $BRI_{2,j}$ *for all subjects* j *that played treatment E1-C second, and compute the sample average,* $\overline{R}_2$*.*

3. *Calibrate* $\hat{\lambda}_C = \overline{BRI}_2 - \overline{BRI}_1$*.*

4. *For iterations* $n = 1, ..., 10000$*:*

    4.1 *Simulate learning* $\hat{x}_i^n$ *using a Poisson distribution with parameter* $\hat{\lambda}_C$*, and debiase best response iteration ability:* $B\hat{R}I_{1,i}^n = BRI_{2,i} - \hat{x}_i^n$*. What we obtain is an estimate of* i*'s* BRI *for the counterfactual situation where she plays E1-C first.*

    4.2 *Record whether debiased* BRI *is higher than* R*:* $\hat{y}_i^n = \mathbb{1}_{\{B\hat{R}I_{1,i}^n > R_{1,i}\}}$

5. *Calculate the fraction of iterations where* BRI *was higher than the debiased* R *estimate:* $\hat{y}_i = \sum_{n=1}^{10000} \hat{y}_i^n / 10000$*.*

Estimates $\hat{y}_i$ can be interpreted as the probability that subject i who played E1-H first and E1-C second would have had a BRI greater than her R in the counterfactual scenario where she played both treatments first, under the assumption that the Poisson learning process is well specified. We use the same process for subjects who played E1-H after E1-C, except that debiasing is done on R instead. Then, by averaging $\hat{y}_i$ across subjects, we obtain an estimate on the probability that a random subject has a BRI which is grater than her R. Using this method, we find that higher order rationality can be explained by limits on the ability compute iterated best responses 44.7% of the time.[10]

1.5 Experiment 2

Experiment 1 allows us to test one specific cognitive constraint: subjects' ability to compute iterated best responses. However, as we discuss in the first section, higher order rationality also involves higher order belief formation. For example, think of a subject that can compute three iterations of best responses, and thus plays the iterated dominance solution when she is Player 2 in treatment E1-C. When she is Player 2 in E1-H, she has to do the same computations in terms of best responses, but she also has to form beliefs about Player 3's beliefs about Player 4's beliefs (i.e. 3rd order beliefs). Because subjects could have a Belief Formation ability that is lower than their Best Response Iteration ability, if this subject is $R2$ in treatment E1-H we cannot say for sure that her rationality order cannot be justified by cognitive constraints.

In light of this argument, we can reinterpret our main finding in the previous section -that around 45% of our subjects' choices in E1-H are due to their BRI- as a lower bound on the proportion of higher order rationality that can be justified by cognitive limitations. In Experiment 2 we test whether a more comprehensive

[10]A 5% confidence interval for this value is $[44.1, 45.2]$. We calculate this interval by Bootstrap. For each iteration $n = 1, ..., 10000$, we sample with replacement subjects from sessions where E1-H was played first and subjects from sessions where E1-C was played first. We then use their R and BRI to calculate this same estimate by Monte Carlo simulations of a Poisson learning process, as shown above, and pick the values that accumulate 2.5% of the estimates to either side of 45% to build the confidence interval.

measure of cognitive limitations that includes limits on the ability to form higher order beliefs can explain a bigger proportion of rationality orders.

In order to define what we mean by "cognitive limitations", we must first define Belief Formation ability:

BFk**:** *We say that a subject has a Belief Formation ability of k (i.e. BFk) if she is able to form kth but not $k+1$th order belief. That is, if she is able to form a belief of the sort "I believe that player X_1 believes that player X_2 believes that.. that player X_k will play strategy Y", but cannot form a belief that goes up to "player X_{k+1}".*

The maintained hypothesis we will use throughout is that forming higher order beliefs and iterating best responses are two distinct processes, so that cognitive constrains on higher order rationality are determined by the tighter one. As an example, think of an individual that can form 2nd order beliefs, such as "I believe that X_1 believes that X_2 will play strategy Y", and can also iterate two best responses in a sequential game like the one used in E1-C. In this case, we say that her cognitive bounds allow her to be second order rational. But this would not be true if she finds it harder to do both processes at the same time instead of each one separately, and thus cannot be second order rational.[11] Using this assumptions, we can define cognitive bounds on higher order rationality in the following way:

[11] Without this assumption we cannot separate between BRI and BF in terms of their effects on higher order rationality, and the notion of Cognitive Ability tested in Experiment 2 should be taken as a measure of the limits on the ability to form higher order beliefs **and** iterate best responses at the same time.

CAk: *We say that a subject who is $BRIr$ and BFs has a Cognitive Ability of k if $k = min\{r, s\}$. That is, Cognitive Ability is the minimum between two bounds: Best Response Iteration ability and Belief Formation ability.*

Experiment 2 uses another variation of Ring Games. As in Experiment 1, our identification strategy is based on comparing behavior in two treatments. The first treatment, called E2-H, serves the same role as E1-H, which is to estimate rationality orders. Taking these estimates as a baseline, the second treatment, E2-C, generates kth order beliefs of rationality for subjects estimated to be Rk. Our identification strategy is based on the idea that a subject whose rationality order was justified only by her kth order belief that players are not rational should react to this change, while a subject who is limited by either her Best Response Iteration or her Belief Formation ability (i.e. CAk subjects) should not.

In both treatments subjects play the same games. These games are a variation of G-1 and G-2 (see Figure 1.3) in which Player 1, Player 2, and Player 3 simultaneously choose an action in the first stage; while Player 4 is allowed to observe these actions before choosing herself in the second stage. Figure 1.8 presents the timing of the game. The only difference between treatments is that in E2-H every role is played by a human, while in E2-C Players 1 to 3 are humans and Player 4 is a "rational" computer best responding to Player 1's move.

For treatment E2-H, making Player 4 choose in the second stage does not change the theoretical predictions fo any rationality order. This is because Player

Figure 1.8: Timing for E2-H and E2-C

4 has dominant strategies (a, c), so as long as she is rational she will choose these strategies, independently of what others have played. Then, since making Player 4 choosing second does not change rationality requirements the same logic as is in E1-H and Kneeland can be applied to the other players for estimating subjects' higher order rationality. We use this timing instead of the simultaneous games played in E1-H in order to make both treatments in Experiment 2 identical, except for the identity of Player 4.

Using a rational computer in E2-C is meant to change higher order beliefs about rationality. As an example, imagine that a subject plays the sequentially rational strategies (g, h) when she is Player 3, and a constant strategy when she is Player 2. From these choices we say that this subject is second order, but not third order rational. The question that we would like to ask is, did she play a constant strategy as Player 2 because she believes that Player 3 believes Player 4 to be irrational? If the answer is yes, then once we make it common knowledge that Player 4 is a rational computer she should switch to playing the sequentially rational strategies as Player 2. In this case, we know that her higher order rationality is not determined by cognitive constrains.

If the answer to the question is no, there are two possible explanations.

The first one is that she is $CA2$ and her choices are justified by her own cognitive constrains on higher order rationality. A second possible explanation is that she believes that Player 3 is $CA1$. In both cases, changing the rationality of Player 4 should not affect her choices as Player 2. Exclusion restriction E2 formalizes this idea for other rationality orders.

ER(E2): *if subjects' higher order rationality estimates in a Ring Game with Player 4 as the second mover are determined solely by her higher order beliefs of irrationality, changing the identity of Player 4 to a rational computer will increase her choice profile's rationalizability order.*

Before we move on, there is one particular case we should analyze, which is when a subject is estimated to be $R1$ in E2-H, meaning that she played a constant strategy as Player 3. If this subject fails to play the sequentially rational strategy when choosing as Player 3 in treatment E2-C, it must be because she cannot find it (i.e. she is $CA1$). This is because her choices cannot be justify by her belief that others are $CA0$, since Player 4 is a rational computer. As shown later, a large fraction of the subjects that do not change their behavior between treatments are $R1$ in E2-H. As for the rest, we will take a skeptical view and assume that they have reached their limits on higher order rationality, as oppose to them believing that others have reached theirs. This way our estimates can be interpreted as an upper limit on the proportion of subjects whose higher order rationality can be justified by their cognitive limitations.

One problem when estimating CA in treatment E2-C is that we do not observe subjects' choices as Player 4, since this role has been replaced by a computer. Since this role only requires rationality, we solve this problem by completing subjects' choice profiles in E2-C with their choices as Player 4 in E2-H. Estimation of CA is then carried out using the same method as for higher order rationality: we classify subjects as CAk if their action profile in E2-C is a perfect match or one mistake away from the one of a kth order rational individual in E2-H.

Implementation

As for Experiment 1, the data for this experiment was collected in the Columbia Experimental Lab for Social Science (CELSS) at Columbia University. 96 subjects were recruited through campus and participated in eight sessions of 12 subjects each. For half of the sessions subjects played E2-H first, while for the other half they played E2-C first. Each subject played a total of 14 rounds without feedback (2 as each of the 4 Players in E2-H, and 2 as Players 1 to 3 in E2-C). Before starting each treatment, instructions were read out loud and subjects had to complete a short quiz in order to help them understand the game. This information is available in the appendix.

1.6 Results

Figure 1.9 presents estimated frequencies of each rationality order in E1-H and E2-H. Both games are identical from a theoretical point of view and we would expect both distributions to be similar. Yet subjects do worse in E2-H, which seems to indicate that they find this game more difficult.[12] Nevertheless, both treatments in Experiment 2 have the same timing -Players 1 to 3 moving first and Player 4 moving second- so any bias derived from it will influence E2-H and E2-C in the same way, without affecting our results.[13]

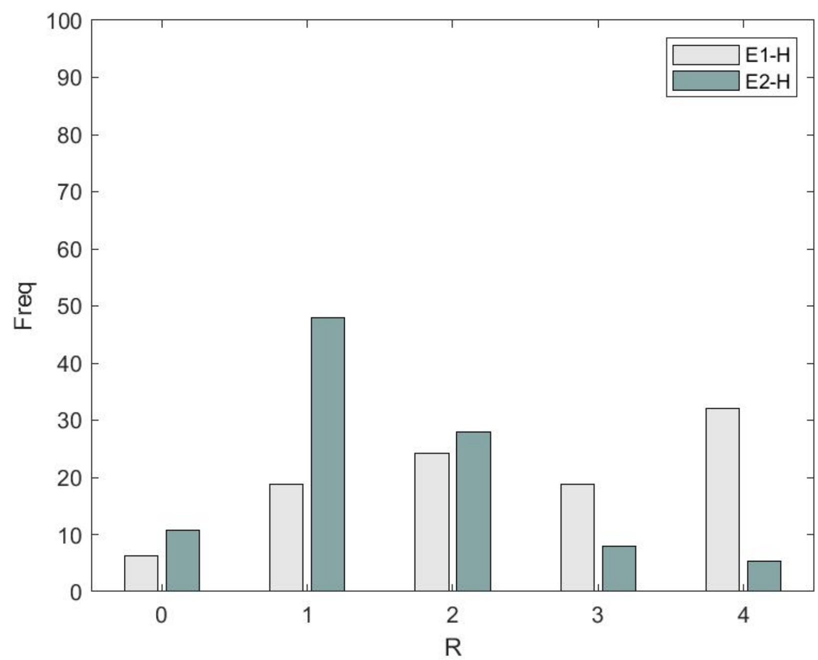

Figure 1.9: Frequencies of rationality orders in E1-H and E2-H

Figure 1.10 reports the distribution of rationality orders from E2-H and Cognitive Ability from E2-C. In the aggregate, subjects seem to react to having a

[12] A Kolmogorove-Smirnov two-sample test rejects the null of equality of distribution with a p-value of less than 1%.

[13] This might raise a concern with respect to the validity of Experiment 1. If adding sequentiality in treatment E1-C makes the game more difficult, we are underestimating the proportion of subjects with limits on their Best Response Iteration ability. However, if this was the case we would expect more subjects with lower BRI than their R (see Figure 7). Also, we should have subjects with R more than one order higher than their BRI in session were learning biases R estimates upwards (i.e. E1-C first). Both of these findings suggest that the added difficulty in Experiment 2 comes having some players choose simultaneously and some sequentially.

rational Player 4, which suggest that for many of them higher order rationality is justified by higher order beliefs of irrationality and not because of cognitive constrains.[14] When moving from E2-H to E2-C, the proportion of subjects classified as $R1$ and $R2$ decrease from 49% to 33% and from 30% to 27%, respectively; while proportion classified as $R3$ and $R4$ increase from 6% to 10% and from 5% to 22%, respectively.

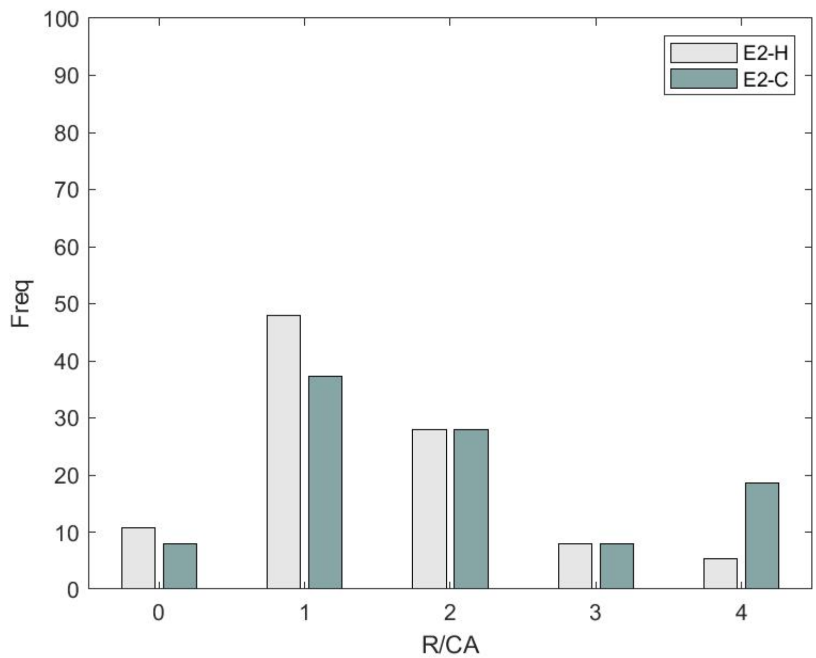

Figure 1.10: Distributions of R orders in $E1 - H$ and CA in $E2 - H$

Within subject analysis presents a more detailed picture. Figure 1.11 reports individual estimates of rationality orders and Cognitive Ability. 15% of the subjects were classified as either $R0$ or $CA0$. Since these subjects appear to be irrational or to make choices inconsistent with our model, we follow the same strategy as in

[14]A Kolmogorov-Smirnov two-sample test has a p-value of 3% to account for discreteness of the distribution.

Experiment 1 and exclude them from subsequent analysis. We also exclude subjects classified as $R4$ in both treatments, since 4-player Ring Games do not allow us to discard the possibility that their rationality order or Cognitive Ability is higher than 4. Of the remaining subjects, 43% have higher estimated CA than rationality order, 39% have the same one, and 18% a lower one.

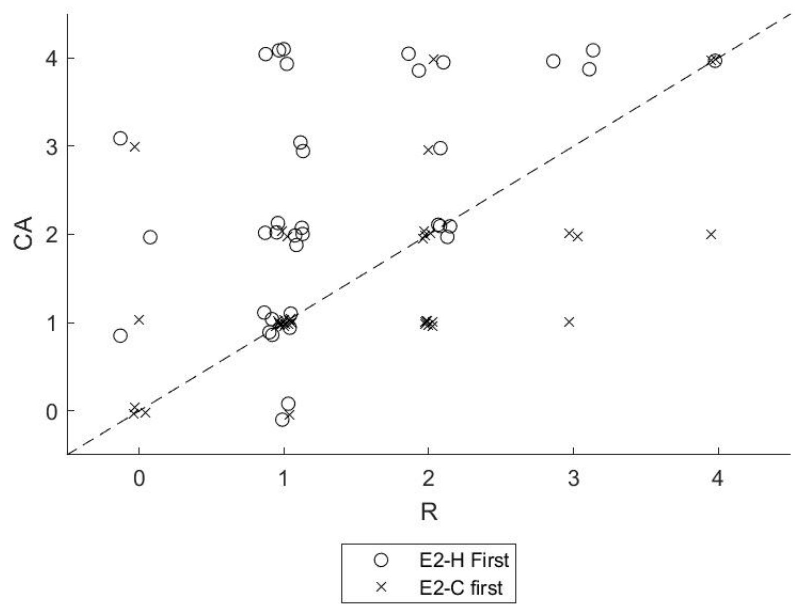

Figure 1.11: Individual estimates of R and CA in Experiment 2
Note:subjects are represented by circles if they played E1-H first and E1-C second, and by crosses if they played E1-C first. A random noise $\epsilon \in [-0.1, 0.1]$ was added to estimates in order to differentiate subjects with the same R and BRI.

Observations falling below the 45° line are counterintuitive, since according to our story CA is an estimate of the maximum order of rationality that subjects may have. However, all but two of these observations are from sessions where E2-H was played second, which suggests that learning might be driving this result. This conjecture is also supported by the fact that all these observations are one step away from the 45° line. As for the two observation that learning cannot explain, these seem to be of subjects that did not understand the game, as seen by their classification as $R1$ in E2-H and $CA0$ in E2-C.

By comparing distributions across sessions, we get more evidence of order effects: on average, R estimates are 0.2 orders higher in sessions where E2-H was played second, and CA estimates are higher by 0.8 in sessions where E2-C was played as the second treatment. Since we know the direction of the bias, we can calculate upper and lower bounds on the proportion of subjects whose higher order rationality cannot be justified by cognitive limitations, by treating sessions with the same treatment order separately. We find that rationality orders are not due to cognitive limitations for between 37.5% and 87.5% of the subjects.

A second method for dealing with learning is to debias our estimates using the same methodology as in Experiment 1, which involves calibrating a Poisson learning process where each subject receives "insights" about the game in-between treatments, which increase their R and CA. Using Monte Carlo simulations of the Poisson learning process, we remove the bias from our estimates and calculate the probability that a subjects' CA is greater than her R. From this procedure we estimate that higher order rationality cannot be justified by cognitive limits (i.e. Best Response Iteration or Belief Formation ability) about 70% of the time.[15]

[15] The bootstrapped 5% confidence interval is $[69.6, 70.4]$. See footnote 10 in Section 4 for further detail about the methodology used to find these numbers.

1.7 Conclusions

In this paper we use variations of Keeland (2015)'s Ring Games to study how much of people's failure to perform many rounds of iterated dominance reasoning in experimental studies can be justified by own limits on cognitive ability, as opposed to beliefs or higher order beliefs that others are not rational. We differentiate between two cognitively taxing processes in performing iterated dominance, Best Response Iteration ability and Belief Formation ability, and design two experiments in order to test how many of subjects' rationality orders can be justified by these constrains and how many by high order beliefs about irrationality.

We start by testing the hypothesis that subjects' choices are not consistent with high rationality orders because they are not able to perform more iterations of best responses. Experiment 1 takes advantage of the fact that transforming a Ring Game into a sequential game and replacing opponents by computers known to be rational eliminates the need to form high order beliefs and generates first order beliefs of rationality, without affecting the need to iterate best responses. After dealing with learning bias due to treatment-order effects, our preferred estimates suggests that subjects' high order rationality can be justified by limits on the ability to compute iterated best responses for about 44.7% of our subjects.

Next, we move to study the more comprehensive notion of Cognitive Ability, which includes both Best Response Iteration ability and Belief Formation ability. Experiment 2 studies whether higher order rationality is revealing own Cognitive

Ability by manipulating subjects' higher order beliefs about rationality. We replace one opponent by a rational computer, and then check whether this change increases rationality orders. Our preferred estimates suggests that about 70% of rationality orders are justified by cognitive limitations. Combining this result with our estimate from the first experiment, we can split this group into 44.7% of subjects who have trouble iterating more best responses,[16] and 25.3% who are limited by their ability to generate higher order beliefs.

This leaves about 30% of subjects who would typically be classified as having reached their cognitive bounds when in fact seem likely to behave as they do because of their beliefs (or higher order beliefs) about others' rationality.

[16] These subjects could have reached both their Best Response Iteration ability and their Belief Formation ability.

Chapter 2

Obfuscating Persuasion

2.1 Introduction

In the Bayesian Persuasion setting introduced by Kamenica and Gentzkow (2011), a Sender persuades a Receiver to take an action by committing to a signal structure, which involves revealing signals about the state of the world. The example used by the authors of a Prosecutor and a Judge has two players, a Sender (him) and a Receiver (her); and two states of the world, *Red* and *Blue*. The Receiver has two available actions, *Red* and *Blue*, and her objective is to match her action to the state. The prior probabilities of each state are set so that without new information, the Receiver would find it optimal to choose *Blue*. The Sender's objective is for the Receiver to choose *Red*, independently of what the state of the world is. He can influence her decision by designing a signal structure that sends different messages to the Receiver with probabilities that depend on what the state of the world is. The main finding from this model is that with the optimal signal structure, the Sender can "persuade" the Receiver to choose his prefer action even more often than by fully revealing the state.

However, the optimal signal structure, as well as how much the Sender can benefit from persuasion, are partly dependent on having Receivers that are rational and Bayesian. The literature on Behavioral Economics has documented many ways in which agents behave suboptimally, including non-Bayes updating (e.g. prior omission bias, prospect theory), and non-rationality (e.g. McKelvey and Palfrey, 1995). With this in mind, in this paper we test the Bayesian Persuasion model in an experimental setting. Our main goal is to find out whether Receivers deviate from theoretical predictions. Specifically, we are interested in testing the validity of one of the main results from Kamenika and Gentzkow, which states that a game where Sender inform the Receiver with which probability each signal will occur conditional on the state is equivalent to one in which he informs her of the posterior probability of each state after each signal realization. Our hypothesis is that this equivalence between signal structures and posteriors does not hold in the data because subjects are not necessarily Bayesian. If this is correct, the Sender can benefit not only from choosing the signal structure, but also from choosing the right way to communicate information and take advantage of Receivers' non-Bayes belief updating.

Our experiment tests the Bayesian Persuasion model using three treatments, which are theoretically equivalent, since they only differ in the way information is transmitted. These treatments vary from one in which subjects can find the correct posterior probabilities by counting how many balls of a color are in a jar, to one in which signal structures are presented using probabilities (e.g. 33%) which subjects then have to update in order to find posteriors.

When looking at Receivers' choices, our first finding is that they generally use cutoff strategies: they have a cutoff posterior, and their choice only depends on whether their updated beliefs about the state fall above or below this posterior. Their belief updating process, however, is biased in that after a "red" signal they overestimate the probability of the state being *Red*. The main result from this paper follows from this finding: Senders benefit from "muddling the water" and making it hard for Receivers to find the correct posteriors, thus taking advantage of their bias in the belief updating process. In particular, we find that the optimal way to communicate information about the state for the Sender is to use signal structures, instead of the (theoretically equivalent) posteriors.

In our experiment, Senders are not able to use persuasion successfully. In fact, their payoffs are comparable to what they would obtain by perfectly revealing the state of world. This result is consistent with previous studies which find that Senders tend to do poorly at first, and only with enough time and feedback are able to take advantage of persuasion.

To our understanding, there exist only three previous studies which bring Bayesian Persuasion to the laboratory. Frèchette et al. (2018) study a unifying experimental framework for different models of communication. These models share the same payoff structure for Senders and Receivers, but vary in whether messages can be verified and in how much commitment the Sender has. By changing these aspects they are able to test games that range from cheap talk, *à* la Crawford and Sobel (1982), to a Bayesian Persuasion setup similar to the one presented in Treatment 3 in our experiment. Nguyen (2017) studies Bayesian Persuasion using a framework similar to our first treatment, although limiting the available signal structures to 4. Au and Li (2018) incorporate reciprocity in the model by allowing the Receiver's utility from choosing each action to depend on how "kind" the Sender was when choosing a signal structure. Our robustness experiment in Section 2.4 tests whether social preferences, of which reciprocity would be a particular case, play a role in Receivers' behavior. We find no evidence in support of the hypothesis that Receivers take into account Senders choices in their decisions, or when looking at how they update their beliefs. However, Receivers perform better in this experiment, suggesting that they find it easier to think when faced with a decision problem than in a strategic setting, where they may have to form beliefs about Senders' strategies. One distinction from this study and previous ones is that in our experiment subjects play the game without feedback, since we are not interested on investigating whether behavior converges to equilibrium, but on whether subjects deviate from theoretical predictions in one-shot games. A second, more significant difference is that to our knowledge, in all existing papers the way the Sender communicates its chosen signal structure to the Receiver is fixed, in contrast to the focus of our study.

The paper is organized as follows. In the next section we introduce the Bayesian Persuasion model and present theoretical results to be tested with our data. Section 2.3 explains our main experiment. In Section 2.4 we discuss our results from this experiment and from a robustness experiment designed to check for social preferences. Section 2.5 concludes.

2.2 Bayesian Persuasion: Theory and predictions

In this paper we will use a particular case of the Bayesian Persuasion model similar to the example of a Judge and a Prosecutor in Kamenica and Gentzkow (2011). The game has two agents, a Sender (him) and a Receiver (her). The state of the world, ω, can be either *Red* or *Blue*, with prior probability $\mu \equiv p(\omega = Red)$. The objective of the Receiver is to guess the state of the world: she gets a payoff of 1 if her guess $g \in \{Red, Blue\}$ is right, and 0 otherwise. Notice that with this payoff structure, the optimal strategy is to guess the most likely state. The prior probability on the state of the world is such that without any information, the Receiver maximizes her expected utility by guessing $g = Blue$. Throughout this paper we are going to use $\mu = 2/5$.

The Sender, on the other hand, wants the Receiver to guess *Red*, independently of the state of the world: he gets a payoff of 1 if $g = Red$ and 0 if $g = Blue$. In order to persuade the Receiver, he can communicate information by providing a signal $s \in \{r; b\}$. This signal is a random variable correlated with the state of the world.[1] Specifically, we define a signal structure E by the probabilities with which each of the two signals are sent, conditional on the state: $E = \{p(s = r|\omega = Red); p(s = r|\omega = Blue)\}$. The signal structure is chosen by the Sender before the state of the world is realized, and it is common knowledge. The

[1] K-G show that any equilibrium strategy where the Sender uses more than two signals is equivalent one where only two are used. Our experimental design restricts the signal space to two. Au and Li (2018) analyze the case of more than two signals in an experimental setting.

timing of the game is as follows:

1. Sender chooses a signal structure E.
2. The state of the world ω is realized.
3. Receiver observes a signal $s \in \{r, b\}$, with the identity of the signal determined according the signal structure chosen in 1.
4. Receiver guesses $g \in \{Red, Blue\}$.
5. Payoffs are realized.

Optimal experiment and observational implications

In this section we will derive the Bayes Nash Equilibrium, together with observational implications to look for in the data.

In order to simplify our analysis, we follow K-G and change the problem from one in which the Sender chooses a signal structure E, to one in which he chooses which Bayes Plausible posteriors he wants the Receiver to have. Take any signal structure $E = \{p(s|\omega = Red); p(s|\omega = Blue)\}$. We say that E *induces* posterior p if p is derived from E by using Bayes Rule. That is,

$$p = p(Red|s) = \frac{p(s|Red)\mu}{p(s|Red)\mu + p(s|Blue)(1-\mu)}$$

We say a duple of posteriors $\{p(\omega = Red|s = r); p(\omega = Red|s = b)\}$ together with a distribution over signals $\{p(s = r); p(s = b)\}$ is Bayes Plausible if

$$p(s = r)p(\omega|s = r) + p(s = b)p(\omega|s = b) = p(\omega) \quad \forall \omega \in \{Red; Blue\} \tag{BP}$$

K-G showed that there is a one-to-one mapping between duples of posteriors and signal distribution, and signal structures. That is, any signal structure will generate posteriors and signal distribution that satisfy Bayes Plausibility. Conversely, any set of posteriors and signal distribution that are Bayes Plausible can be derived using a unique signal structure and the prior via Bayes Rule.

The treatments used in this experiment are going to serve as a test of this result, which derives from assuming Bayesian agents. Observationally, it implies that as long as subjects are using Bayes Rule to update their beliefs and believe others are doing so too, presenting an experiment where the Sender chooses and the Receiver observes signal structures is equivalent to one in which the Senders chooses and the Receiver observes posteriors. Result BP-1 summarizes this idea.

BP-1 (Bayesian Agents): *The following games are observationally identical:*

1. A game where the Sender chooses a signal structure $E = \{p(s|\omega = Red); p(s|\omega = Blue)\}$.

2. A game where the Sender chooses (Bayes Plausible) duples of posteriors $\{p(\omega =$

$Red|s = r); p(\omega = Red|s = b)\}$ together distributions over signals $\{p(s = r); p(s = b)\}$.

We can now characterize the Perfect Bayes Equilibrium. Without loss of generality, we will assume that $p(\omega = Red|s = r) \geq p(\omega = Red|s = b)$. This assumption is similar to the notion of straightforward signals in K-G and the Revelation Principle (Myerson, 1979). It states that after receiving a r signal, the updated posterior on the state being *Red* must be higher than after receiving a b signal.[2]

The optimal strategy for the Receiver is to guess the state that she thinks is more likely, and play any pure or mixed strategy if her posterior is equal to 1/2. This is because her expected utility from guessing g after observing signal s is equal to her posterior on the state being equal to her guess, $p(\omega = g|s)$.

BP-2 (Rational Receivers): *A Receiver with posterior $p(\omega = Red|s)$ after observing signal s has a cutoffs strategy, with the cutoff being at 1/2:*

$$g = \begin{cases} Red & \text{if } p(\omega = Red|s) > 1/2 \\ Blue & \text{if } p(\omega = Red|s) < 1/2 \\ Red \text{ or } Blue & \text{if } p(\omega = Red|s) = 1/2 \end{cases}$$

[2]Straightforward signals must also satisfy $p(\omega = Red|s = R) \geq 1/2$ and $p(\omega = Red|s = b) \leq 1/2$, so that they can be interpreted as recommendations to take an action.

The final step in order to characterize the equilibrium is to find the optimal induced posteriors, which are the Sender's best response to the Receiver's strategy. From BP-2 we know that the only way for the Receiver to choose $g = Red$ is to have a posterior greater or equal to 1/2. By analyzing Bayes Plausibility when $\omega = Red$ we can see that with a prior $\mu = 2/5$ it is impossible to induce posterior beliefs that satisfy this condition for both signals, since the prior has to be equal to a weighted average (with signal probabilities as weights) of these posteriors:

$$p(s = r)p(Red|s = r) + (1 - p(r))p(Red|s = b) = 2/5$$

Using this result, the optimization problem for the Sender can be transformed into one where he maximizes the probability of signal r, subject to Bayes Plausibility[3] and the incentive compatibility constraint $p(Red|r) \geq 1/2$:

$$\max_{p} \quad p(r)$$
$$s.t. \begin{cases} p(r)p(\omega = Red|r) + (1 - p(r))p(\omega = Red|b) = 2/5 \\ p(\omega = Red|r) \geq 1/2 \end{cases}$$

From the first constraint of the problem we can see that in order to maximize $p(r)$, the sender must choose a signal structure such that both posteriors on the state being *Red* are as low as possible.[4] Intuitively, any increase in the probability of *Red*

[3] If Bayesian Plausibility is satisfied for $\omega = Red$, it will also be satisfied for $\omega = Blue$.

[4] Notices that the left hand side of the Bayes Plausibility constraint is increasing in $p(r)$ under the assumption $p(\omega = Red|r) > p(\omega = Red|b)$. Using the Implicit Function Theorem we obtain that $p(r)$ is decreasing in both posteriors.

after signal s comes at the cost of decreasing $p(r)$. The second constraint is necessary for the Receiver to guess *Red* after observing signal r. This constraint holds with weak inequality, since in the only PBN the Receiver is choosing the Sender's preferred action upon indifference. As we will discuss later this may not be true in an experimental setting, which will result in the optimal signal structures inducing more certainty after signal r. Thus in the solution to this problem the Receiver believes that each state is equally likely after receiving signal r, and is certain that the state is *Blue* after receiving signal b. The next two results formalize these ideas, and together with Bayes Plausibility fully characterize the optimal signal structure for the Sender.

BP-3 (Optimal induced posteriors): *In the PBE, after observing r the optimal signal structure induces beliefs $p(\omega = Red|r) = 1/2$.*

BP-4 (Optimal induced posteriors): *In the PBE, after observing b the optimal signal structure induces beliefs $p(\omega = Red|b) = 0$.*

The PBE has two implications about how persuasion affects both players' payoffs. The fact that the Sender can do better than with no information transmission is straightforward: without any information revelation the Sender gets a payoff of 0, while with an information structure that perfectly reveals the state the expected payoff for the Sender is 2/5, which corresponds to the probability with which the incentives of both players are aligned. But the Sender can do even better by making signal r less precise and increase its occurrence. With the PBE signal structure described by BP-3 and BP4 the Receiver guesses *Red* 4/5th of the time, even though the state is *Red* half as often.

BP-5 (Sender's benefit from persuasion): *In the PBE, the Sender benefits from persuasion. With the optimal signal structure, the Sender's expected utility is equal to 4/5, compare to 0 with an uninformative signal structure and 2/5 with a fully revealing signal structure.*

There is one final prediction to look for in the data, which is that in equilibrium the Receiver does not benefit from persuasion. This is directly related to result BP-3. To see how, note that since after observing r the Receiver is indifferent between both guesses and after observing b she guesses *Blue*, her ex-ante expected payoff in the PBE has to equal the one from guessing *Blue* after both signals, which is what she does with no information.

BP-6 (Receiver's benefit from persuasion): *In the PBE, the Receiver does not benefits from persuasion. With the optimal signal structure, the Receiver's expected payoff is equal to 3/5 (i.e.* $1 - \mu$*).*

Formally, the unique PBE of this game is:

PBE: *In the unique PBE with* $\mu = 2/5$*,*

- *The Sender chooses experiment* E^* *with* $\{p(r|Red); p(r|Blue)\} = \{1, 2/3\}$.

- *The Receiver updates her beliefs upon receiving signal s using Bayes Rule, and chooses $g = Red$ if her belief is greater or equal to 1/2, and $g = Blue$ otherwise.*

The next section describes our experimental setting.

2.3 Experiment

Our main experiment was run using the Mechanical Turk platform. We recruited two hundred subjects, half of whom were assigned to play as Senders while the other half played as Receivers. Each subject played three treatments, each one consisting of ten rounds of the game described in the previous section. To qualify to participate in the experiment, subjects were required to be in the 99 percentile of the "Hit Approval Rate", a measure of how well they did in previous studies.[5] Additionally, we restricted our pool to subjects that have participated in at least 50 previous studies. On average, Senders completed the study in 14 minutes and 9 seconds, earning $1.179 dollars; while Receivers took 10 minutes and 31 seconds, earning $0.955 dollars.

The three treatments differ only in the way the signal structure is presented, and thus should have similar behavior if BP-1 is true. Alternatively, if we find systematic differences across treatments we can use them to investigate how subjects update their beliefs and to ask whether Senders could take advantage of this behavior.

[5] Some studies conducted in Mechanical Turk have questions designed to test whether the subject is paying attention and understands the problem presented. In these cases, subjects can be rejected if they do not answer these questions correctly. Hit Approval Rate measures the percentage of previous studies a subject has not been rejected to.

In all treatments there are ten numbered balls. The first four of these balls are red, while the last six are blue (see Figure 2.1). One of the ten balls is randomly selected at the beginning of the Receiver's turn, and she is asked to guess its color. If she guesses right, she earns 70 points for the round. The Sender, on the other hand, earns 70 points for the round if the Receiver guesses *Red*, independently of the actual color of the selected ball. Final monetary earnings for the experiment consisted of 50c as a participation fee, plus 70c for each point the subject earned in one randomly selected round per treatment.

This scenario is meant to mimic the one presented in the previous section: the color of the randomly selected ball represents the state, ω, and thus the prior probability of the state being *Red* is equal to the number of red balls divided by the total number of balls, 2/5.

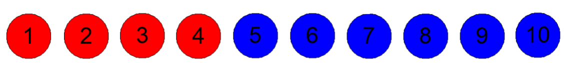

Figure 2.1: State of the world

In Treatment 1, Senders have to choose how to allocate all ten balls into two jars, Jar X and Jar Y. Receivers are shown how the balls are distributed into the two jars, and are told in which jar the randomly selected ball is. This setting is equivalent to the Bayesian Persuasion game from Section II, with the exception that Senders have a limited number of available signal structures to choose from. The number of red and blue balls in each jar represents the signal structure, while the identity of the jar with the chosen ball is the signal realization. We chose to identify the jars as X and Y in order to avoid any bias generated by an association with the colors of the balls, as might be the case with a “Red Jar" and a “Blue Jar". Figure 2.2 shows an example of an arrangement, where Jar X has three red and two blue balls, while Jar Y has one red and four blue balls.

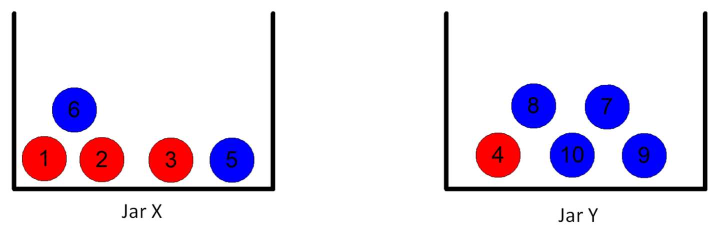

Figure 2.2: Treatment 1, example

Note that with this setup, it is easy to figure out what the posterior probabilities are. Say, for example, that the ball is in Jar X. Then, the probability of the ball being red is 3/5, since there are 5 balls in that jar and 3 of them are red. Similarly, if the ball is in Jar Y the probability of it being red 1/5. The signal structure generating these posteriors probability can also be recovered, although with more difficulty, by counting how many balls of a color are in each jar. In the example, conditional on a red (blue) ball being selected, the probability of a Jar X signal is 3/4 (2/6), since Jar X has 3 (2) of the 4 (6) balls of that color.

Treatment 2 is designed so that subjects find it easier to think in terms of signal structures than in terms of posteriors. For this part part of the experiment, each of the 10 balls is in a small jar, which is labeled either “X" or “Y". Once the Sender has chosen how to allocate the labels, the Receiver is shown all Jars’ labels, and whether the selected ball is in a jar labeled X or Y. Figure 2.3 shows an example, where three of the red balls are in jars labeled X and one in a jar labeled Y, while two of the blue balls are in jars labeled X and four in jars labeled Y. Notice that by grouping the balls by colors, subjects should find it easier to figure out the signal structure than the posterior. In this example, a subject can work out that a red ball is in a jar labeled as X with probability 3/4, since 3 of the 4 red balls are in jars with that label. Similarly, a blue ball is in a jar labeled as X with probability 2/6. From a theoretical point of view, this signal structure is equivalent to the one represented in Figure 2.2. Thus, according to theory, and more precisely result BP-1 Treatment 1, and Treatment 2 should yield the same results.

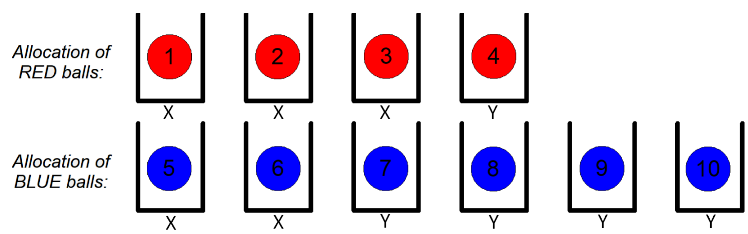

Figure 2.3: Treatment 2, example

The last treatment is similar to Treatment 2 in that subjects have to deal with signal structures as oppose to posteriors. In Treatment 3 Senders are shown two small jars, Jar X and Jar Y. At the beginning of the Receiver's turn, the chosen ball is placed into one of them. The Sender's decision problem is to choose with which probability the selected ball is placed into either jar, depending on its color. The Receiver is shown these probabilities, as well as whether the selected ball is in Jar X or Jar Y, before having to guess its color. Figure 2.4 shows an example, where if the randomly selected ball is red it is placed in Jar X with probability 75% and in Jar Y with probability 25%; while if the ball is blue it is placed in Jar X with probability

33% and in Jar Y with probability 67%. This example represents the same signal structure as the previous ones. Presenting the problem in this way has two distinctive features. First, it ensures that subjects have to use Bayes Rule to obtain the correct probabilities of each state. In Treatment 2 it is easier to see what the signal structure is, but it is still possible to find the posterior probabilities without using Bayes Rule, by counting how many red balls have a label and dividing that number by the total number of balls with that same label. Second, subjects have to think in terms of probabilities. This resembles more closely the way we thing about these types of problems in the theory, compare to using the number of balls in a jar as in the other two treatments.

Figure 2.4: Treatment 3, example

To make the treatments comparable, the available signal structures for the Senders were limited. They could choose to place a red ball in Jar X with probability 0%, 25%, 50%, 75%, or 100%. A blue ball could be placed in Jar X with probability 0%, 17%, 33%, 50%, 67%, 83%, or 100%. These probabilities generate the same signal structures as the ones from Treatment 1 and Treatment 2.

2.4 Findings

In this section we describe the findings of our experiments. We will be using results BP-2 to BP-6 from Section 2.2 as a guideline to contrast theoretical predictions with the data, while BP-1 will be tested by comparing the validity of these results across treatments. Our main hypothesis is that subjects are not Bayesian, which leads to systematic deviations from the theory that could be exploited by the Sender.

Subjects played ten rounds of the game in each treatment, for a total of 30 rounds. The order in which they played each treatment was randomized at a subject level. There is some evidence that behavior in these type of games can be skewed, depending on random realizations in previous rounds (Nguyen, 2017). This would be the case if, for example, a Receiver's guess depended on whether the randomly chosen ball was red or blue in previous rounds. For this reason and in order to mimic a series of one shot games, subjects played without feedback for the entire experiment. Our results do not change if we restrict our sample to the last 5 rounds of each treatment, suggesting that learning is small.

In the following analysis, we will be avoiding the use of "jars" and "labels" to refer to signal realizations, reverting to our previous nomenclature of "signal r" and "signal b". For each observation we choose signal r to refer to the jar or label that has the largest probability of the ball being red. For example, if one game has the signal structure presented in Figures 2.2, 2.3, and 2.4; we will refer to Jar X (or jar labeled X) as signal r since it implies a posterior probability of 3/5, compare to

1/5 in the other one. This allows us to use the same terminology as when explaining the the theory in Section 2.2 across all treatments.

We start our analysis by studying Receivers' choices and the validity of BP-2.

Receivers

A natural starting point in analyzing Receiver's behavior is to test whether choices are consistent with subjects having (weakly increasing) cutoff strategies, meaning that if they guess *Red* (*Blue*) when the posterior probability of a red ball is p, they will also guess *Red* (*Blue*) when the posterior is higher (lower) than p. This is a necessary condition for result $BP-2$ to be true.

For the most part, our data supports this notion. In Treatment 1, 74% of our subjects are consistent with having a cutoff strategy. This number decreases to 67% for Treatment 2 and 66% for Treatment 3.[6] One interpretation for this difference across treatments is that the first treatment only involves finding the correct guess given the posterior (i.e. checking whether the posterior is greater or smaller than 1/2), while the other two have the added difficulty of having to update the prior using Bayes Rule. This increase in difficulty could be leading to more mistakes when having to make a guess, as would be the case if subjects have limited attention or

[6] A bootstrapped K-S test with a null hypothesis that the fraction of subjects with cutoff strategies is the same between treatments generates p-values of 21% for Treatment 1 versus Treatment 2, 16% for Treatment 1 versus Treatment 3, and 82% for Treatment 2 versus Treatment 3.

effort that they have to distribute among both tasks.

A sharper prediction derived from BP-2 is that subjects should switch from choosing *Red* to choosing *Blue* when the posterior probability is equal to 1/2. To check whether this is corroborated by the data, we estimate the cutoff for each subject and each treatment. Let $p = (p_1, p_2, ..., p_{10})$ be the Bayes induced posterior probabilities of the ball being red that a subject encounters during a treatment, ordered in ascending order. If there exists a posterior p_j, so that the subject guesses *Blue* for every $k \leq j$ and guesses *Red* for every $k' \geq j+1$, we assign this subject to have a cutoff strategy. If either p_j or p_{j+1} is equal to 1/2, we estimate the cutoff to be $c = 1/2$. Otherwise, the estimated cutoff is $c = \frac{p_j + p_{j+1}}{2}$. Additionally, if the subject's choices do not fit this criteria but are one mistake away from it, we still assign her to have a cutoff strategy, and the estimation is done correcting that mistake.[7] Admitting one mistake increases the number of subjects with cutoff strategies by 6% in Treatment 1, and 7% in Treatments 2 and 3.

Figure 2.5 presents the cumulative distribution of cutoffs by treatment, conditional on having a cutoff strategy. The doted line is what the CDF would look like according to the theory, assuming that agents are Bayes Rational (i.e. BP-2). In Treatment 1, 70.27% of the subjects had behavior consistent with BP-2. Of the remaining, 12.16% had cutoffs below, and 17.57% had cutoffs above 1/2. Importantly, 96% have cutoffs between 0.3 and 0.7, suggesting that for the most part subjects

[7]When allowing for one mistake, simulations show that the probability of classifying random behavior as a cutoff strategy is around 1%.

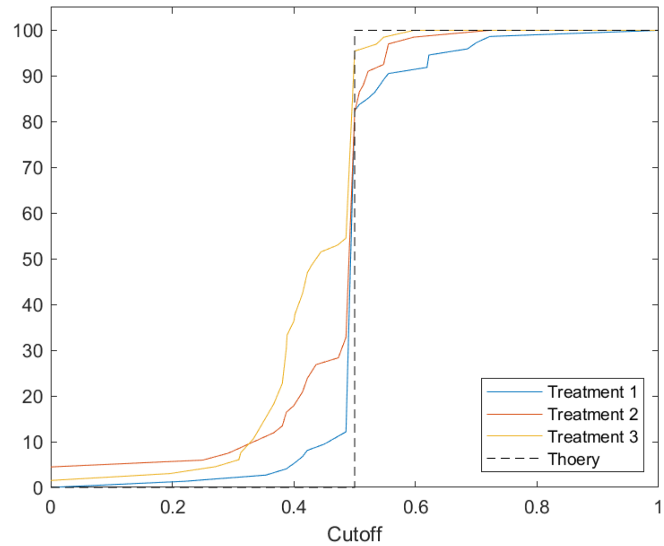

Figure 2.5: Cumulative Distribution of cutoffs
Cumulative distribution of cutoffs by treatments. The doted line represents our theoretical predictions derived from BP-2. A Kolmogorov-Smirnov test for difference in distribution has p-values of 0.7% for Treatment 1 vs Treatment 2, 0% for Treatment 1 vs. Treatment 3, and 0.7% for Treatment 2 vs. Treatment 3.

did not guess colors that were very unlikely. We conclude that subjects are (mostly) rational.

For the other treatments, rationality implies that subjects' choice profiles should be consistent with an increasing cutoff strategy, but the cutoff does not need to at 1/2. To be precise, if we only assume rationality, BP-2 could be interpreted using subjective probabilities. That is, BP-2 would say that an agent should guess ω if she *believes* state ω to be the most likely one. Only if agents are Bayesian (i.e. BP-1) we can treat these beliefs as the objective posteriors. Comparing cutoffs in Treatments 2 and 3 with Treatment 1 is thus an indicator of how subjects are updating their beliefs. When looking at Treatment 2, the proportion of subjects estimated to have a cutoff of 1/2 drops to 49%. More interestingly, the distribution of cutoffs is asymmetric, suggesting an upward bias in the belief updating process: 32% of the subjects have cutoffs lower than 1/2, compare to 18% with a higher one.

These results are even more pronounced in Treatment 3: 41% of our sample has a cutoff of 1/2, 54% has a lower one and only 5% a higher one. From these findings we can conclude that, in general, subjects seem to be overestimating the probability of the ball being red when updating their prior.

A possible explanation for these results is that subjects have base-rate omission bias, meaning that subjects underweight the prior probability of the state when constructing their posterior. This bias was first described in Kahneman and Tversky (1972 and 1973), and has been found in many different populations, such as doctors (Eddy, 1982); academic researchers (Bird, 2018), and law students (Eide, 2011).[8] We can illustrate this bias using the "false positive" story, introduced by Casscells et. al. (1978). A patient is tasted positive for a rare disease that is found in 1 of every 1000 people in the population. The test is known to be 100% accurate if the person tested has the decease, but it gives false positives with probability 1%. When asked what is the probability that the patient that tested positive has the disease, subjects tend to answer numbers close to 100%, even though Bayesian Updating would give a posterior of around 10%. This is because subjects tend to underweight the fact that a person having the disease is very unlikely in the first place (i.e. the prior). In our example, for every 1000 people, 999 do not have the disease, and of those 9.99

[8]See Benjamin et. al. (2019) for an extensive review of the literature.

(i.e. 1%) give a false positive. Thus, for every 1 sick person there are 9.99 healthy ones that test positive, resulting in a probability of having the disease conditional on testing positive of $1/(1+9.99)$, or around 10%.

A example of a model of base-rate omission biase used in the literature applied to our experiment would have subject i updating her belief after receiving signal $s = r$ in the following way:[9]

$$p_i(Red|r) = \frac{p(r|Red)\mu^{\alpha_i}}{p(r|Red)\mu^{\alpha_i} + p(r|Blue)(1-\mu)^{\alpha_i}} \qquad \alpha \in [0,1]$$

where α_i captures how much weight the subject puts in the prior when forming the posterior. If a subject has $\alpha_i = 1$, she is Bayesian, and she will guess *Red* only when state *Red* is more likely. In the other extreme, a subject that has $\alpha_i = 0$ would not take the prior into account and guess *Red* after r as long as $p(Red|r) > p(Red|b)$. However, notice that underweighting the prior would not have any effect on choices after observing signal b, as it would increase $p_i(Blue|b)$, which even for Bayesian agents is greater than 1/2. From Bayes Plausibility we know that after observing b the Bayes posterior (on the state being *Red*) cannot be grater than the prior, 2/5, so overestimating the probability of *Blue* after signal b would still lead to a posterior $p_i(Red|b) < 1/2$. But if a signal structure has a Bayesian posterior $p(Red|r)$ between 2/5 and 1/2, a subject with a low enough α_i would guess *Red* after r, resulting in an estimated cutoff lower than 1/2.

[9]This model is borrowed from Benjamin et. al. (2019) and Starting with Grether (1980)

A second explanation for the asymmetry in the distribution of cutoff strategies is that subjects have social preferences, and update in favor of the guess preferred by the Sender. This seems implausible given that social preferences would have to enter the updating process, and not just the utility function, in order to explain our results. This is because if Receivers took into consideration Senders' when making their decisions, we would expect Treatment 1 to have an asymmetric distribution as well. Nevertheless, we test this explanation with a robustness treatment is Section 2.4, and find no evidence to support it.

Finally, cutoffs are also concentrated around 1/2 for these treatments, suggesting that Receivers do not make obvious mistakes. Similar to Treatment 1, 94% and 95% of the cutoffs are between 0.3 and 0.7 in Treatments 2 and 3, respectively.

Sender-optimal signal structure

In this section we are going to derive the empirical Sender-optimal signal structure, defined as the one that maximizes the likelihood of a Receiver guessing *Red*. We will be focusing on how this structure changes across treatments, in order to find out up to what extend a Sender can benefit from manipulating the way he communicates information by taking advantage of the bias in the Receiver's belief updating process (i.e. in deviations from BP-1).

In Section 2.2 we characterized the PBE signal structure. This signal structure is based on two insights. The first one says that after observing a b signal Receivers should be sure that the state is *Blue* (BP-4). In our data, increasing the posterior probability of *Red* after observing b does increase the probability of a Receiver guessing *Red*, but this increase is not big enough to offset the decrease in the posterior after observing r plus the decrease in the probability of that signal occurring. Thus, result BP-4 is supported by our experiment.

The second insight is summarized by BP-3, which states that the optimal signal structure must leave the Receiver indifferent between both guesses after observing signal r. This result, however, is valid only when the Sender is unrestricted in the signal structures he can choose from, which implies that he can choose $p(Red|r)$ to take any value in the unit interval. In our experiments, the probability of sending each signal conditional on the state can only take one of five values when the state is *Red* and seven values when the state is *Blue*. In Treatment 1, for example, the Sender can only choose to have signal "Jar X" occur with probability $\alpha/4$ when the state is *Red*, by placing α red balls in Jar X and $4-\alpha$ in Jar Y. Because of this, the Sender cannot induce posteriors that are arbitrary close to 1/2: among signal structures that satisfy BP-4, the closest he can get from above is to induce $p(Red|r) = 4/7$.

The discretization of the space of probabilities eliminates the uniqueness of equilibria. Define $q \in [0;1]$ as the probability that the Receiver guesses *Red* when indifferent.[10] We now have two PBE, both of which induce a posterior 0 after b (BP-4), but differ in their induced posteriors after r. The intuition for this multiplicity

[10] More precisely, q is the mixed strategy of the Receiver after observing r with the signal structure that satisfies both BP-3 and BP-4. She may be playing another strategy when, for example, $p(\omega = Red|r) = 1/2$ and $p(\omega = Red|b) = 1/4$.

is as follows. Take a Sender that is thinking about inducing a posterior of 1/2. He can increase the probability of the Receiver guessing *Red* after observing r from q to 1 by inducing a larger posterior. But doing so decreases the likelihood of signal r happening. The direction in which this trade-off is resolved, which depends on whether q is high or low, will determine the equilibrium.

In the Sender preferred PBE (PBE-S), q is high enough so that the trade-off is resolved in favor of having signal r more often, and thus he induces a posterior of 1/2. In the Receiver preferred PBE (PBE-R), however, the opposite happens and the Sender is willing to choose a slightly more informative signal r at the cost of it being less frequent. The following results describe these two cases.

BP-3S (PBE-S): *If* $q > 7/8$*, the optimal signal structure for the Sender induces a posterior of* $p(\omega = Red|r) = 1/2$.

BP-3R (PBE-R): *If* $q < 7/8$*, the optimal signal structure for the Sender induces a posterior of* $p(\omega = Red|r) = 4/7$.

We can estimate q by looking at how often Receivers guess *Red* when the signal structure is $\{p(Red|r); p(Red|b)\} = \{1/2, 0\}$ and the realized signal is r. In Treatment 1, q is equal to 65%. The low likelihood of Receivers guessing *Red* when indifferent, compare to 85% with a posterior of 4/7, leads to an expected payoff for the Sender of 0.49 with the signal structure described in BP-3S, and a payoff of 0.63 with the one described in BP-3R. Thus, in Treatment 1 the empirically optimal signal structure is the one described by BP-3R. Interestingly, increasing the informativeness of signal r even further, so that the posterior is now 2/3, generates a similar expected payoff of 0.6, since the lower probability of this signal occurring is partly compensated with Receivers guessing *Red* even more often when they observe it.

A similar situation happens for Treatment 2. Receivers switch from guessing *Red* 68% of the time when the posterior is 1/2 to guessing 91% when the posterior is 4/7, with expected payoffs for the Sender of 49% and 63%, respectively.

In Treatment 3, however, the optimal signal structure is to generate a Bayes updated posterior of 1/2. As shown in the previous section, in this treatment a large proportion of Receivers have an upward biased posterior after observing r. A signal structure that induces indifference after signal r for Bayesian agents will then induce a posterior higher than 1/2 for these subjects. The probability of a Receiver guessing *Red* is 0.87 when the (Bayes) posterior is 1/2, which leaves the Sender with an expected payoff of 0.78. This probability increases only slightly to 0.89 when the posterior is 4/7, generating an lower expected payoff of 0.66.

Summarizing, the PBE signal structures seems to be the best choice for Senders when face with our Receivers, albeit not always for the same reasons as in the theory. As seen in Treatments 1, Receivers do not always break their indifference in favor of the guess preferred by the Sender, suggesting that even if the signal structure space was larger, leaving the Receiver indifferent between guesses after signal r (i.e. BP-3) might not be the optimal thing to do. Moreover, even when the PBE strategy from the continuous signal space is optimal, it is not because Receivers are breaking their indifference in favor of guessing red, but because they have a biased updating process, and their subjective beliefs are in fact greater than 1/2.

Finally, our main result derives from the fact that the payoff associated with the empirically optimal signal structure varies across treatments, and it is highest in Treatment 3. In other words, a Sender can benefit the most from persuasion if he uses signal structures to communicate information and if he is able to obfuscate the Receiver's belief updating process by communicating using probabilities, thus making it harder for her to find the correct posterior and taking advantage of her biased beliefs. In our experiment, a Sender who uses the empirically optimal signal structure (for each treatment) can persuade a Receivers to guess *Red* 78% of the time if allowed to communicate signal structures as in Treatment 3, compared to 63% of the time in Treatment 1 and Treatment 2.

Senders

In our data, Senders are not able to find the empirically optimal signal structures described in the previous sections. Figure 2.6 presents the cumulative distributions of average expected payoffs across rounds for each treatment. Although Senders are obtaining positive payoffs, they do not seem to be fully comprehending the benefits of persuasion. Average expected payoffs are comparable to what a Sender who chooses a signal structure at random would obtain. Expected payoffs averaged 0.31 for Treatment 1, 0.34 for Treatment 2, and 0.37 for Treatment 3; compared to 0.35, 0.34, and

0.42 using a random strategy. Furthermore, differences in payoffs across treatments seem to come mainly from differences in Receivers behavior. A Kolmogorov-Smirnov test for equality in distributions of signal structures cannot reject the null at a 10% level for any pair of treatments. The p-values are 3% for Treatment 1 vs Treatment 2, 4.5% for Treatment 1 vs Treatment 3, and 7.7% for Treatment 2 vs Treatment 3.

We interpret these findings as a refuting of BP-5: Senders are not benefiting from persuasion. In fact, not only they cannot find the empirically optimal signal structure, but they obtain similar expected payoffs than what they would have had with a perfectly revealing signal structure, which would have expected payoffs of 0.38, 0.3, and 0.47 in the three treatments. The validity of Result BP-6 is less clear as Receivers do benefit form Senders choices, although not much. Their average expected payoffs given Senders' signal structures are 0.67 for Treatment 1, 0.66 for Treatment 2, and 0.65 for Treatment 3; compare to payoffs of 0.6, 0.56, and 0.52 with an uninformative signal structure.

One explanation for this behavior is that the problem for Senders is harder than it is for Receivers. While a Receiver has to guess between one of two options, a Sender has 35 possible signal structures to choose from. Furthermore, Receivers are faced with a decision problem, while Senders' problem is a strategic one, adding the complexity of having to form beliefs about the other player's strategy. These results are supported by previous studies, who find that only with more rounds and with feedback Senders improve their play and partly converge to the optimal signal structure (Nguyen, 2017; Au and Li, 2018). Another possible explanation is that

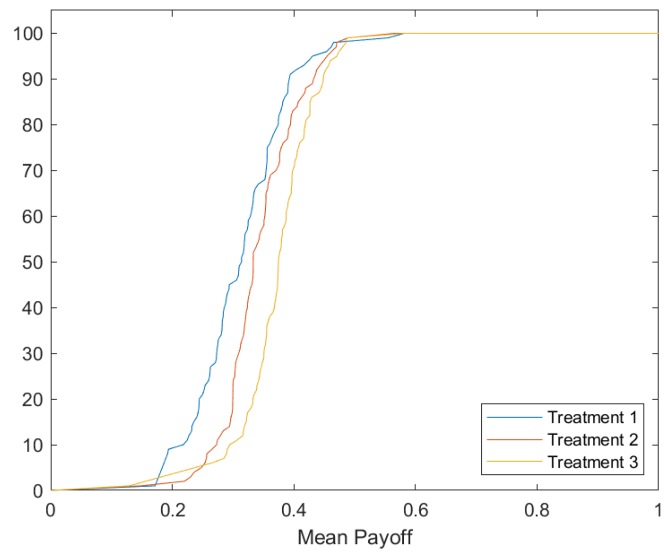

Figure 2.6: Cumulative Distribution of payoffs for Senders
Cumulative distribution of average expected payoffs by treatments. A Kolmogorov-Smirnov test for difference in distribution has p-values of 1.2% for Treatment 1 vs Treatment 2, 3.7% for Treatment 1 vs. Treatment 3, and 38% for Treatment 2 vs. Treatment 3.

monetary incentives were not high enough for Senders to put the necessary effort in finding the optimal strategy. We explore this idea in the appendix by running an experiment with higher payoffs and a possibly more sophisticated pool, but find little evidence in support of it.

Robustness: social preferences and reciprocity

In Section 4.1 we found that Receivers are overestimating the probability of the state being *Red* when updating their beliefs, as seen in lower estimated cutoffs in Treatment 2 and Treatment 3. Since guessing *Red* benefits the Sender, this behavior could be explained by a Receiver who has social preferences. For example, it might be that she would rather err in favor of the action preferred by the Sender when she is not very confident about the state of the world. Alternatively, the Receiver might label

Senders who choose signal structures the induce posteriors close to 1/2 as "unkind" and penalize them by guessing *Blue* more often. This hypothesis is studied in Au and Li (2018), who incorporate the model of reciprocity proposed by Falk and Fischbacher (2006) into the Bayesian Persuasion setting, and study its validity in an experimental setting.

As a robustness test, we run a second experiment without the role of Sender. This experiment is identical to the one used for our main results except for two changes. First, we modified the instructions, eliminating parts of two phrases where it was explained that the allocation of balls and labels was chosen by another subject, and one paragraph that described the payoff of the Sender.[11] The second modification is that in each treatment subjects played for eight rounds using the eight signals structures described in Table 2.1. The order in which they encountered each signal structure and the order of the treatments was randomized. We recruited a total of 150 subjects though the Mechanical Turk platform, using the same requirements as in the main experiment.

	S1	S2	S3	S4	S5	S6	S7	S8
$p(Red\|r)$	1	4/5	4/6	4/7	1/2	3/5	1/2	3/7
$p(Red\|b)$	0	0	0	0	0	1/5	1/4	1/3

Table 2.1: Induced posteriors
Induced posteriors of signal structure used in the experiment. For Treatment 2 and 3, these are the posteriors induced by a Bayes agent.

Table 2.1 shows Bayes induced posteriors for the eight signal structures used

[11] The instructions for both experiments can be found in the Appendix.

in this experiment. For the first five signal structures the Receiver is certain that the state is *Blue* after receiving signal b, thus satisfying BP-4, but the structures differ in how informative signal r is. S4 and S5 are the PBE structures satisfying BP-3R and BP-3S, respectively. S1 is a perfectly revealing structure. The last three signal structures were included to test result BP-4, and they induce a positive probability of *Red* after observing signal b. S6 and S7 are similar to the equilibrium signals, S4 and S5, in that they induce posteriors that are slightly higher and equal to 1/2. S8 is a the worst possible type of signal structure for the Receiver, as it always induces posteriors lower than 1/2.

Subjects seem to do better in this setup. In Treatment 1, 82% of our subjects made choices consistent with having a cutoff strategy, compare to 74% in our main experiment. The difference increases in the other two treatments: 89% and 81% of our subjects seemed to be using a cutoff strategy in Treatment 2 and Treatment 3, respectively, versus 67% and 66% in the main experiment. This could suggest that subjects find it more difficult to make decisions when playing with other people. One caveat is that part of this increase in the number of subjects assigned to have cutoff strategies might be due to them playing fewer rounds. As an illustration of this problem, the probability of a Receiver guessing at random to be estimated as having a cutoff strategy increases from 1% in the main experiment to 4% in this one. In order to make both experiments comparable, we can use choice data from the last 8 rounds in our main experiment. Notice that this is a conservative approach in favor of the null hypothesis that the proportion of subjects with a cutoff strategy is the same across treatments. When truncating the data in this way, we cannot reject the null for the Treatment 1, with a p-value of 60%. However, Treatment 2 and Treatment 3 significantly higher across experiments, with p-values of 2% and 1%. This suggests that subjects find strategic problems harder than decision problems, at least when updating beliefs about the state of the world.

The cumulative distributions of estimated cutoffs, conditional on having such strategy, is presented in Figure 2.7. For the most part, the patterns found in Section 2.4 are still present. In Treatment 1, 82.8% have a cutoff equal to 1/2, with 5.7% below it and 12.5 above it.

Omitting the Sender seems to have the effect of reducing biases coming from belief updating in Treatment 2, as is evident from a more symmetric distribution: 85.6% of the subjects had a cutoff of 1/2, with 8.5% below it and 6% above it. In Treatment 3, however, the asymmetry persist: 56.6% of the subjects had a cutoff of 1/2, 38.6% below it and 4.8% above it.

If social preferences were responsible the results from our main experiment, we would expect the bias to be reduced or disappear in both Treatment 2 and Treatment 3. The fact that the bias persists in the later suggests that our findings are more likely to be explained by deviations from Baeysian Updating unrelated to social preferences, such as the prior omission bias discussed above. In this experiment, the great reduction of the bias in Treatment 2 (and slight reduction in Treatment 3) may be explained by Receivers having an easier time solving a decision problem than a strategic problem, further supporting the idea that obfuscating the belief updating process can lead to biases which Senders can then take advantage of.

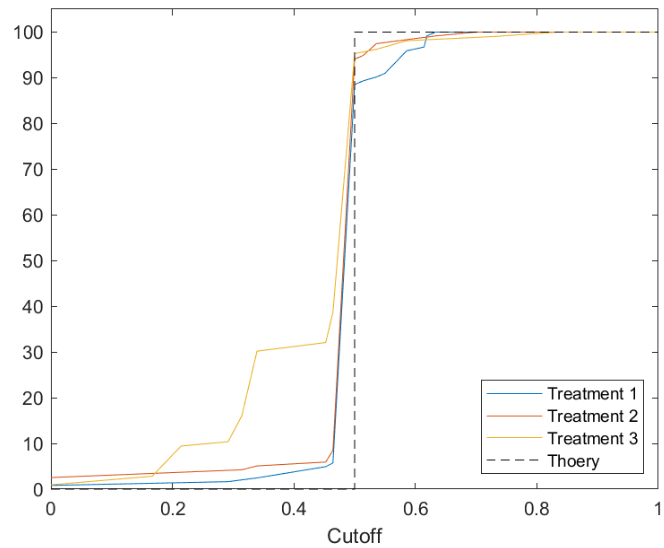

Figure 2.7: Cumulative Distribution of cutoffs
Cumulative distribution of cutoffs by treatments. The doted line represents our theoretical predictions derived from BP-2. A Kolmogorov-Smirnov test for difference in distribution has p-values of 9.7% for Treatment 1 vs Treatment 2, 0% for Treatment 1 vs. Treatment 3, and 0% for Treatment 2 vs. Treatment 3.

2.5 Conclusion

In this paper we test the model of Bayesian Persuasion in an experimental setting, using three treatments that vary the way in which signal structures are presented. Our first finding is that Receivers generally use cutoff strategies, as predicted by the theory. Their belief updating process, however, is biased towards overestimating the probability of the state being *Red* after observing a *Red* signal. Because of this bias, the isomorphism between signal structures and induced posteriors is not validated. Our main finding is that Senders should not only care about how much persuasion to do, but also about how they transmit the information. In particular, in our data Senders can significantly increase their benefit from persuasion by obfuscating the problem for Receivers, and use signal structures with probabilities instead of posteriors. These results partially persist when we eliminate the role of the Sender from the game in our robustness experiment, suggesting that they are not due to Receivers having social preferences.

Finally, we find that Senders are not able to use persuasion successfully, as their payoff is comparable and sometimes lower than what they would have gotten from full disclosure. We hypothesize that this is because their decision problem is harder than for Receivers, and the benefits of persuasion are also hard to understand. With feedback Senders would probably find the optimal persuasion strategy over time, as found in previous studies, but this would likely also decrease Receivers' bias as well. For example, Grether (1980) find that as subjects become more experienced, their base-rate omission bias is reduced, suggesting that feedback might decrease possible gains from obfuscating posterior probabilities by presenting signal structures as in Treatment 3.

Chapter 3

Multisensory Integration of Financial Data

with Zeyang Chen, Mark Dean, Harrison Hong, and Juanjuan Meng.

3.1 Introduction

Behavioral economics and finance largely emphasizes how computational limitations impede individual decision making under uncertainty. Individuals conserve on precious cognitive resources by adopting biased inference strategies or potentially even ignoring subsets of signals when making forecasts. This cognitive miser view, which goes back to (Simon (1982), Kahneman & Tversky (1973)) has been modeled or studied in a variety of ways such as rational inattention (Sims (2003), Woodford (2009), Caplin & Dean (2015), Peng & Xiong (2006)), reliance on simple models (Hong et al. (2007)) or other types of information costs (Gabaix (2014)). Regardless of the particular modeling approach, the underlying neuroscience assumption is that humans have limited neurons for computation.

Recent neuroscience research, however, points to another important and distinct mechanism associated with sensory limitations that impedes the accuracy of human inference (Raposo et al. (2012), Sheppard et al. (2013)). Experiments with rodents point to sensory bottlenecks in decision tasks such as judging whether pulses of light are long or short. Rodents do better when presented with both visual and redundant auditory signals (i.e. long (short) pulses of light correspond to drawn out (rapid) beeps of sound). When presented with redundant auditory cues, rodents simultaneously integrated both signals to perform better in these decision tasks than having either visual or auditory cues alone. The integrationof multisensory signals has been extended to humans in these simple tasks.

Through a variety of interventions for rodents, studies find that the reason is that neurons associated with sight and sound are dedicated and do not get redeployed for general cognition or computation (Licata et al. (2017)). That is, the performance in these decision tasks are not solely driven by bounded computation ability per se but by bottlenecks in the sensory architecture. As a result, expanding the range of signals beyond visual cues by leveraging auditory and potentially other sensory cues such as touch can lead to significant improvements in performance.

In this paper, we explore the importance of multisensory integration of data in financial markets. Financial market data are typically presented visually to investors, e.g. ubiquitous Bloomberg terminals present financial data on prices and trading volume typically as time series. Investors make inferences based on this visual presentation—be it to predict stock prices or liquidity the next second, the next day or the next year. Extrapolating from the neuroscience literature on multisensory integration — investors could make better predictions were some of this data also presented simultaneously as sound.

There is little research in financial economics in this vein except for the creative work of Coval & Shumway (2001), who find that elevated levels of sound in the futures trading pit contained valuable information for futures prices. Futures traders often report that background noise of others traders making bids conveyed valuable information, i.e. allowed them to make better inference regarding the arrival of news or liquidity.

To test for the value of multisensory integration in financial markets, we recruit subjects to play a simplified versions of a market trading game described in Hong & Rady (2002) — an informed trader or hedge fund (the subject) who knows the fundamental value of the stock one period ahead trades against market prices set by competitive market makers as in Kyle (1985). Whereas the market maker knows the liquidity of the stock, i.e. the amount of noise trading which can be drawn from a high or low liquidity distribution, the subject does not and have to infer it from past trading volume.

At the beginning of each round, prices and trading volume can be generated either from a high or low liquidity regime using formulas described in Hong & Rady (2002). Realized trading volume is the sum of the absolute values of the market order of the informed trader, the noise trades and inter-dealer trades that occur between market makers. The higher is trading volume, the more likely traders are in a high liquidity regime. But there is noise due to inter-dealer trades.

Given parameters for fundamental value and liquidity trades, we use the model to simulate paths of fundamental values, prices and trading volume or turnover using this model. In each round of the game, new paths of fundamental values, prices and trading volume are drawn, and subjects are then asked two questions. First, do they want to buy or sell the stock? The answer to this simply relies on comparing the terminal (oneperiod ahead) fundamental value and the current (last) value. This question is not our focus but serves two purposes: (1) as a check on whether the subjects understand what is going on; and (2) a placebo test of framing effects, which we discuss below. Indeed, we expect that students should almost always get this answer correct assuming they are paying attention and/or understand the set-up.

The second question is whether the market is hot (i.e. liquid) or cold (i.e. not liquid). The answer to this question depends on the level of trading volume. Students can receive either a standard visual cue in the form of a time series of volume on the screen, and/or a redundant auditory cue where we transform the time series of volume into sound. We consider several transformations of the visual trading volume cue to sound that exploits some of the unique qualities of sound to convey information. The accuracy of this question depends on the parameters of the model.

Subjects play many rounds of this game and randomly receive different treatments of trading volume signals. They get paid based on a randomly chosen round. The subjects start with 400 experimental coins. If they answer a question right, they get 100 points, while they loose 100 points if they answer it wrong. Summarizing, if they get both answers right in the same round then they obtain a maximum of 600 experimental coins (100 coins for each question). If they get both answers incorrect, then its 200 as they get deducted 100 points each question. If one is correct and the other wrong, then they get 400.

We can use a within-subject design where we have subject fixed effects to see if the accuracy of the subject varies depending on whether the subject is receiving visual, auditory, or both signals. We can also use a cross-subject design as we also record a number of demographic information about subjects including a variety of follow-up tests on attention and IQ for students that participated in different sessions.

We recruited students from Peking University (about 2/3 of our sample) and Columbia University (1/3 of our sample). There were a total of 12 sessions conducted between June 2018 and March 2019. These sessions are summarized in Table 2, where subjects in different sessions were confronted with either easy or hard discrimination tasks as to whether the market is hot or cold and potentially different presentations of sound.

The results are summarized in Table 4. When presented with just graph information, their average round bonus is 493 as opposed to 486 when presented with just sound information. There is a statistically significant difference between sound and graph, with the former under-performing the later — this is not surprising since graph is the more natural mode of stimulation when it comes to finance. Sound is nonetheless informative since it generates a round bonus that exceeds 400, which is what one might expect if the subject guessed randomly. That is, it is possible to present financial information using sound. More relevant for our paper is that when presented with graph and sound simultaneously, their round bonus is 501.85. As such, presenting both graph and sound is better than graph by about the same margin as graph being better than sound.

The accuracy rate (i.e. the probability of getting the answer right) for the buy/sell question is 0.77 regardless of the presentation of the market question in graph, sound or graph-and-sound mode. Remember that this is not obvious as there might be spillovers through salience effects whereby the presentation of sound might lead the subject to pay more attention. This is comforting as it tells us that introducing sound is not generating some salience effects per se. However, the accuracy rate for the market question is 0.73 for graph-and-sound, 0.70 for graph and 0.67 for sound. The round bonus differences are indeed only coming from the differences in these accuracy rates and not the accuracy rate of the buy/sell question. We show that these baseline results are robust across using within or across subject designs. The other interesting finding is that the effectiveness or value of multi-sensory integration relies on the informativeness of the sound signal. Under parameter values where it is very difficult to discriminate between hot versus cold regimes, adding a redundant sound signal generates no significant value — i.e. performance in the graph-and-sound mode is not better than graph. These findings suggest that better ways to transform financial data into sound might markedly improve our results.

Our paper proceeds as follows. We describe our experimental design in Section 2. We provide summary statistics in Section 3. Our main findings are in Section 4.

3.2 Experimental Design

Experimental mock-up

This experiment simulates a simplified version of a stock market modeled in Hong & Rady (2002). In each round of trading $t = 1, ..., T-1$, we use this model to simulate out prices and volume of a stock that the subject sees or hears before making inferences. In period T an informed trader, the subject, has to decide how much of the stock to buy or sell.

Fundamental value: We denote as V_t the fundamental value of the stock in period $t = 1, ..., T$. Given V_0 and σ_v, the fundamental value is assumed to follow a random walk:

$$V_t = V_{t-1} + v_t, \tag{3.1}$$

where $v_t \sim i.i.d.N(0, \sigma_v^2)$.

Informed trades: In periods $t = 1, .., T-1$, an informed trader observes v_t at $t-1$ and submits a market order of the form $S_t = \beta v_t$.

Noise or liquidity trades: There are also noise or liquidity traders who submit market orders of u_t, where $u_t \sim i.i.d.N(0, \sigma_u^2)$. We further assume that σ_u can be one of two values, set at the beginning of each round:

$$\sigma_u = \begin{cases} \sigma_{uH} & \pi = 1/2 \\ \sigma_{uL} & 1 - \pi = 1/2 \end{cases} \tag{3.2}$$

Stock Price: We denote P_t as stock price in period t, determined as in Kyle (1985). Price is set by competitive market-makers to equal the expect fundamental value conditioned on market orders from the informed trader and noise traders. That is, the stock price in current period equals the stock value in previous period added with two independent exogenous disturbances (the informed trader's market order (βv_t) and the noise trades (u_t)). Given V_0, P_0, σ_v and σ_u, the expression of stock price is:

$$P_t = V_{t-1} + \lambda(S_t + u_t), \tag{3.3}$$

where

$$\lambda = \frac{\sigma_v}{2\sigma_u},$$

and

$$\beta = \frac{\sigma_v}{\sigma_u}$$

Trading volume: We denote Q_t as trading volume in period t. The trading volume in the current period is defined as the absolute value of the sum of three independent exogenous disturbances: the informed trader's market order (βv_t), the noise trades (u_t), and inter-dealer trades $e_t \sim i.i.d.N(0, \sigma_e^2)$. Given σ_v, σ_u and σ_e, the expression of trading volume is:

$$Q_t = |\beta v_t + u_t + e_t| \tag{3.4}$$

Payoff: Subjects play the role of the informed trader in period T. We denote S_T as trading decision and π_T as payoff in the final period. $S_T \geq 0$ means to buy and $S_T \leq 0$ means to sell. The expression of payoff is:

$$\pi_T = S_T(V_T - P_T)$$

The decision S_T contains two dimensions, the sign and the absolute value of S_T. First, the sign of S_T should be the same as the sign of $(V_T - P_T)$. P_T is unknown to the trader before making a decision, but she knows that if she doesn't traded ($S_T = 0$) its expected value is just V_{T-1} (this can be seen by taking the expectation of equation 3, and using $E(u_t) = 0$). Therefore, if the stock value goes up in the final period, the correct decision is to buy, and vice versa.

Second, the absolute value of S_T is related to the value of σ_u. If $\sigma_u = \sigma_{uH}$, then the correlation between trading volume and stock price (i.e. λ) is low, so the trader can buy or sell a large volume without influencing the stock price too much. In the contrast, if $\sigma_u = \sigma_{uL}$, then the correlation between trading volume and stock price is high, and the trader's optimal choice is to only buy or sell a small volume. We call the former situation "busy market" and the latter situation "quiet market".

Task description

In this experiment, subjects play the role of an informed trader who will buy or sell some number of shares of a stock in period $T = 21$, as explained in the previous section. Subjects are informed that all trading information are randomly generated by computer, but we do not offer the formulas or parameters. Instead, we describe the intuition of the stock market and give them some tips to make decisions. In each round, they need to make two decisions: (1) To buy or to sell? (2) Is the market type busy or quiet?

After observing stock price P_t and fundamental value V_t information in the $1^{st} - 20^{th}$ period and stock value information in the 21^{st} period, subjects need to decide to buy or to sell in the 21^st period. The rule of thumb is as follows:

- $V_{21} \geq V_{20} \rightarrow$ buy
- $V_{21} \leq V_{20} \rightarrow$ sell

After observing trading volume Q_t in the $1^{st} - 20^{th}$ period, subjects need to judge whether the market type is busy or quiet. The thumb rule is:

- You think $Q_1, ..., Q_{20}$ is high $\rightarrow$ busy market
- You think $Q_1, ..., Q_{20}$ is low $\rightarrow$ quite market

These two questions are very different in the decision process. For the buying/selling question, the key information is not provided until the final period.

Once the information is disclosed, it is always a clear signal and indicates a certain answer. However, for the market type question, trading volume in each period is useful, but the other random disturbances apart from σ_u adds uncertainty to the signal.

At the beginning of each round, subjects receive 400 experimental coins as initial wealth. A correct choice is rewarded with 100 experimental coins and a wrong choice is punished by having to pay 100 experimental coins. Therefore, the payment in each round could be 600 (if both choices are correct), 400 (if one is correct, the other is wrong) and 200 (if both choices are wrong). In total, there are 30 independent rounds and no feedback throughout the experiment. Experimenter will randomly select one round as final payment.

Within-subject treatments

In the 30 rounds, historical trading volume information are presented randomly in three different forms, which respectively are:

- Graph mode. Trading volume presented by graph (a moving line chart).
- Sound mode. Trading volume presented by sound. We will describe the sound content and mapping method below.
- Graph & Sound mode. Trading volume presented by the combination of graph and sound at the same time.

The three modes are equivalent in the quantity of information presented, but differ in forms of expressions. In neoclassical economics, information is always well-received and understood no matter what form it is in. In other words, visual and auditory information are both perfectly informative and substitutes for each other. Therefore, the null hypothesis is that there is no difference in performance across the three treatments.

Besides, this variation in historical trading volume only affects the information in market type question. The buying/selling question remains the same across the three information modes above.

Our experiment has two aims. First, we want to measure the effectiveness in receiving visual or auditory information separately, by comparing the outcome with the theoretical predicted accuracy. Second, we want to study the relationship of visual and auditory senses, by combining graph and sound information in the third treatment.

Sound content

We have two ways of mapping trading volume into auditory form. The first is called "white noise", in which we convert trading volume into loudness of white noise. Considering auditory sense is a log form of loudness, our mapping function is in exponential form. We limit the maximum level of loudness to 10, in order to protect participants' hearing.

$$Loudness = 10 \times min(1, e^{\frac{0.693315Q_t}{2\sigma_u}} - 1)$$

The second is called "bubbling", in which we convert trading volume into the frequency of bubbling sound. We also take an exponential form to keep it consistent with the mapping into loudness. Then we round it down to integer times of bubbling in one second.

$$Frequency = floor(9 \times min(1, e^{\frac{0.693315Qt}{2\sigma_u}} - 1))$$

The type of market influence the auditory information through two channels. First, the market type parameter σ_u is in the mapping function, as shown above. Second, the market liquidity parameter σ_u also influence the value of trading volume, Q_t. Figure 1 shows the mapping of trading volume into auditory form for busy and quite markets, which depends on how we present this information and on whether it is a hard session, with $\sigma_{uL} = 8$, or an easy session, with $\sigma_{uL} = 3$ (more on this in the next section). Quiet market has a steeper slope and a narrower value range, both channels making it less distinguishable than busy market.

Readers can click the two buttons in Figure 2 to listen to examples of sound contents.[1]

[1]Enable Flash Helper Services in your computer if the audios do not work. You can right-click the bar to stop or play again

Cross-subject treatments

Apart from within-subject treatments, we have four other cross-subject treatments. Table 1 shows sample sizes in each subgroup.

Country: We carried out the same experiment in China (Peking University) and US (Columbia University) to compare the effects among two countries. We translated the same website into Chinese and English versions and offered purchasing-power equivalent payments to participants in two countries. In our sample, around two thirds of the participants are from Peking University and one third are from Columbia University.

Difficulty: There are two difficulty levels of market type questions. The parameters in "hard" sessions are: $\sigma_u H = 10$, $\sigma_u L = 8$, $\sigma_e = 3$. The parameters in "easy" sessions are: $\sigma_u H = 10$, $\sigma_u L = 5$, $\sigma_e = 3$. In order to describe the difficulty levels objectively, we simulated 5000 rounds based on Bayesian updating. Approximately, the correction rate is 100% in "easy" sessions and 88.6% in "hard" session. Theoretically, this variation should not affect performance in the buying/selling question.

Sound contents: As we have described, the first mapping of trading volume into sound is called "white noise", in which we convert trading volume into loudness of white noise. The second is called "bubbling", in which we convert trading volume into the frequency of bubbling sound.

Question sequence: In some sessions, participants answer the market type question first and the buying/selling question second. We swap the sequence in other sessions in order to see if there is any spillover effect of information.

Cognitive test

After 30 trading rounds, subjects need to take four cognitive tests. We order the four test according to their importance in the research.

Attention level test: In this set, subjects need to respond to color changes of two squares shown in the middle of the computer screen. When the game starts, the two squares randomly change their color every second, and they are asked to press the space bar when both of their colors match. If this set is selected in the end, the computer will randomly choose 10 rounds out of the 180 rounds, and they can get 10 experimental coins for each correct response in these 10 rounds.

Memory test: In this set, subjects are presented with 10 4-digit numbers, which they have 2 minutes to remember. They then have to recall these numbers in another 2 minutes (in any order). If this set is selected in the end, they can get 10 experimental coins for each correct recall. In this set, we do not allow using recording tools such as paper, pencil, cellphone, computer etc.

IQ test: In this set, subjects need to answer some logical questions. If this set is selected in the end, they can get 10 experimental coins for each correct answer.

Attention scale: In this set, subjects answer ten questions about their attention level in real life. If this set is selected in the end, they can get 100 experimental coins independently of their answers.

Experimental implementation

Up till now, we have 13 sessions and 232 subjects in total. Table 2 includes detailed information about experimental settings in each session. Since variations of these variables are not independent of each other, we must add them as controls in our regressions (or have Session fixed effects).

3.3 Summary Statistics

Individual Characteristics

Table 3 is the summary statistics of individual characteristics. The number of observations is not always 232 because our server happened to miss several records. Among all subjects, 39.4% are male. Their average age is 22.41. The rest of the variables come from the cognitive tests described above. Att_Level is the objective measure of attention level. Att_Scale is the self-reported scale survey of attention level. Memory is how many numbers they were able to remember. IQ is the score in logic-IQ questions. IQ_Min is the time they spent on IQ test.

The last four columns tests the differences of means of each variable by four cross-subject variations, so as to check the randomization. Columbia students are 1.92 years older than Peking students. Their attention level and IQ are also slightly lower than Peking students. Subjects in easy sessions are 1.13 years older and a little less attentive than their counterparts in hard sessions. There is no significant difference across two types of sound content. Subjects who takes market type questions first have somehow a poorer memory. These differences suggest that in the regression, controlling all these personal features will be more reliable than adding individual fixed effects.

Task Performance

Figure 3 shows average round bonus and accuracy in BS or MKT question in 3 information modes. In the graph, rounds in Graph & Sound mode gain the highest bonus and rounds in Sound mode gain the lowest bonus. Comparing to bonus in Graph mode, both are significant at 5%. There is no evident gap across three information modes in buying/selling question. Accuracy in market type questions contributes mainly to the bonus gap, which is also significant at 5%.

Table 4 shows the summary statistics of performances in detail. On average, subjects earn 494.9 experimental coins in each round. The overall accuracy is 77.0% in the buying/selling question and 69.9% in the market type question. In our experimental mock-up, the buying/selling question always has a definite correct answer, so theory predicts that its accuracy should be 100%. On the other hand, accuracy in market type question should be 100% in hard sessions and 88.6% in easy sessions if subjects can incorporate all the information presented and update their belief via Bayesian Rule. Not surprisingly, in the data these two rates are both much lower than the theoretical predictions.

The next two variables are time spent in each of the two questions. On average, subjects spent 10.82 seconds in each buying/selling question and 6.488 seconds in each market type question. This disparity may be due to the following. In the buying/selling question, the critical information is released in the last period, so subjects cannot come up with an answer until the last period appeared. In contrast, in market type question subjects receive useful information in each period and update their belief gradually, so they may arrive to the last period with a high level of certainty about the market type.

3.4 Results

Baseline Results

Table 5 shows regression estimates. Regressions on round bonus use OLS, and the ones on buying/selling or market type correction use Probit regression. All the standard errors are clustered at an individual level. Columns (1) (2) and (3) control for round and individual fixed effects. Columns (4) (5) and (6) drop individual fixed effects and add session fixed effects as well as control variables, including cognitive capacity, demographics and experience in financial markets. Considering differences in personal features, as we discussed above, these three columns are our preferred specifications.

Mode effect: The probability to give a correct answer in sound mode is $e^{0.170} = 0.84$ times of that in graph mode, which means auditory sense is less effective than visual sense in receiving information. The probability to give a correct answer in graph & sound mode is $e^{0.171} = 1.19$ times of that in graph mode, which means the combination of graph and sound information is more effective than graph information alone. Besides, information modes have no effect on accuracy of buying/selling question, so there is no evidence of spillover effect. The mode effects in market type questions lead to 7.489 less bonus and 7.995 more bonus separately.

Setting effect: Peking subjects did well in buying/selling question than Columbia students but they do not have difference in market type question. This could be related to a younger age, higher attention level and higher IQ in Peking sample. There is a large difficulty effect. The probability to give a correct answer in hard sessions is only $e^{0.925} = 0.40$ times of that in easy sessions. Moreover, difficulty in market type question does not influence accuracy in buying/selling question, which also suggests that these two questions are quite independent of each other. Sound content and question sequence do not affect the performance.

Time effect: The longer time spent on the question, the lower accuracy it is. A possible explanation is that the time spent reflects the relative difficulty of that round.

Heterogeneous Analysis

In this part, we analyze the heterogeneity across four variations in experimental settings by run subsample regressions as well as adding interactions terms into regressions.

Country. Table 6 shows in Columbia, subjects give a better performance in sound mode comparing to graph mode. In Peking, subjects perform better in multi-sensory mode. But there is no evidence of heterogeneity in mode effect across two countries.

Difficulty. Table 7 shows in easy sesssions, multi-sensory information works better than graph information alone. But in hard sessions, this effect does not exist. Moreover, in hard sessions, sound mode leads to a poorer performance. Besides, there is no evidence of heterogeneity in mode effect across two difficult levels.

Sound content. Table 8 shows comparing to graph mode, white noise will lead to a poorer performance in market type question than bubbling sound, which means white noise alone is less effective than bubbling sound. This suggests that auditory sense is better at judging the frequency of sound than the volume of sound.

Question sequence. Table 9 shows both auditory effect and multi-sensory effect only appear when market type question comes first.

3.5 Conclusion

In this paper, we devise an experimental trading game in which the subject, playing the role of an informed trader, has to infer the liquidity of the stock market based on a time series of trading volume. This time series is presented both using visual and auditory cues. The premise is that recent neuroscience research points to sensory bottlenecks as opposed to purely computational constraints in degrading performance—that is, we expect subject performance to improve markedly when provided not just visual but redundant auditory cues. Consistent with these sensory bottlenecks, we find that subjects performed better in their inference of the liquidity of the market when presented simultaneously with both visual and auditory cues as opposed to either cues alone. The improved performance is moderately robust even as we have only experimented with basic sounds. As such, we expect that there can be scope for alternative presentations of alternative sensory cues to markedly boost multisensory integration for financial data.

References

Abadie, A. (2002). Bootstrap Tests for Distributional Treatment Effects in Instrumental Variable Models. *Journal of American Statistical Association*, 97(457): 284-292.

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

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

Arad, A., and A. Rubinstein (2012). The 11-20 Money Request Game: A Level-k Reasoning Study. *American Economic Review*, 102(7); 3561-73.

Au, P. H., and K. K. Li (2018). Bayesian Persuasion and Reciprocity: Theory and Experiment. *Working paper.*

Benjamin, D., A. Bodoh-Creed, and M. Rabin (2019). *Base-Rate Neglect: Foundations and Implications. Working paper.*

Bird, A. (2018). Understanding the Replications Crisis as a Base Rate Fallacy. *British Journal for the Philosophy of Science*, 13.08.2018.

Brandenburger, Adam, and Xiaomin Li (2015). Thinking About Thinking and Its Cognitive Limits. *Working paper*, New York University.

Caplin, Andrew, and Mark Dean (2015). Revealed preference, rational inattention, and costly information acquisition. *American Economic Review*, 105(7), 2183–2203.

Coval, Joshua D., and Tyler Shumway (2001). Is sound just noise? *The Journal of Finance*, 56(5), 1887–1910.

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

Casscells, W., W. Schoenberger, and T. Graboys (1978). Interpretation by Physicians of Clinical Laboratory Results. *New England Journal of Medicine*, 299: 999-1001.

Camerer, Colin F., Teck-Hua Ho, and Juin-Kuan Chong (2004) A Cognitive Hierarchy Model of Games. *The Quarterly Journal of Economics*, 119(3) 861-898.

Crawford, V. P., and J. Sobel (1982). Strategic Information Transmission. *Econometrica*, 50(6): 1431-1451.

Eddy, D. (1982). Probabilistic reasoning in clinical medicine: Problems and opportunities. in *Judgment under uncertainty: Heuristics and biases*, ed. by D. Kahneman, P. Slovic, and A. Tversky, Cambridge University Press, Cambridge: 249-267.

Eide, E. (2011). *Two tests of the base rate neglect among law students.* Working Paper.

Frèchette, G. R., A. Lizzeri, and J. Perego (2018). Rules and Commitment in Communication. *Working paper.*

Friedenberg, Amanda, Willemien Kets, and Terri Kneeland (2017). Bounded Reasoning: Rationality or Cognition. *Working paper.*

GABAIX, XAVIER (2014). A sparsity-based model of bounded rationality. *The Quarterly Journal of Economics*, 129(4), 1661–1710.

Georganas, Sotiris, Paul J. Healy, and Roberto A. Weber (2014). On the Persistence of Strategic Sophistication. *Journal of Economic Theory*, 159(A): 369-400.

Gill, David, and Victoria Prowse (2016). Cognitibe ability, character skills, and learning to play equilibrium: A level-k analysis. *Journal of Political Economy*, forthcoming.

Goeree, Jacob K., and Charles A. Holt (2001). Ten little treasures of game theory and ten intuitive contradictions. *American Economic review*, 91, 1402-1422

Grether, D. (1980). Bayes Rule as a Descriptive Model: The Representativeness Heuristic. *The Quarterly Journal of Economics*, 95:537-746.

Ho, Teck-Hua, Colin Camerer, and Keith Weigelt (1998). Iterated Dominance and Iterated Best-response in p-Beauty Contests. *American Economic Review*, LXXXVIII: 947-969

Hong, Harrison, and Sven Rady (2002). Strategic trading and learning about liquidity. *Journal of Financial Markets*, 5(4), 419–450.

Hong, Harrison, Jeremy C. Stein, and Jialin Yu (2007). Simple forecasts and paradigm shifts. *The Journal of Finance*, 62(3), 1207–1242.

Hörner, J., and A. Skrzypacz (2016). "Learning, experimentation and information design". mimeo.

Jin, Ye (2016). Does Level-k Behavior Imply Level-k Thinking?. *Working Paper*, U.C. Berkeley.

Kahneman, D., and A. Tversky (2972). Subjective Probability: A Judgment of Representativeness. *Cognitive Psychology*, 3:430:454.

Kahneman, D., and A. Tversky (2973). On the Psychology of Prediction. *Psychological Review*, 80:237-251.

Kamenica, E., and M. Gentzkow (2011). Bayesian Persuasion. *American Economic Review*, 101(6): 2590-2615.

Keller, G., and S. Rady (2010). Strategic experimentation with Poisson bandits. *Theoretical Economics*, 5: 275-311.

Kneeland, Terri (2015). Identifying Higher-order Rationality. *Econometrica*, 93: 2065-2079.

Kyle, Albert S. (1985). Continuous auctions and insider trading. *Econometrica: Journal of the Econometric Society*, 1315–1335.

Licata, Angela M., Matthew T. Kaufman, David Raposo, , Michael B. Ryan, John P. Sheppard, and Anne K. Churchland (2017). Posterior parietal cortex guides visual decisions in rats. *Journal of Neuroscience*, 37(19), 4954–4966.

McKelvey, R. D., and T. R. Palfrey (1995). Quantal Response Equilibria for Normal Form Games. *Games and Economic Behavior*, 10(1): 6-38.

Myerson, R. B. (1979). Incentive Compatibility and Bargaining Problem. *Econometrica*, 47(1):61-73.

Nagel, R. (1995). Unraveling in guessing games: an experimental study. *American Economic Review*, 85 (5), 1313-1326.

Nguyen, Q. (2018). Bayesian Persuasion: Evidence from the Laboratory. *Working Paper*.

Palacios-Huerta, Ignacio, and Oscar Volij (2009). Field Centipedes. *American Economic Review*, 99(4): 1619-1635.

Peng, Lin, and Wei Xiong (2006). Investor attention, overconfidence and category learning. *Journal of Financial Economics*, 80(3), 563–602.

Raposo, David, John P. Sheppard, Paul R. Schrater, and Anne K. Churchland (2012). Multisensory decision-making in rats and humans. *Journal of neuroscience*, 32(11), 3726–3735.

Sheppard, John P, David Raposo, and Anne K. Churchland (2013). Dynamic weighting of multisensory stimuli shapes decision-making in rats and humans. *Journal of vision*, 13(6), 4–4.

Simon, Herbert A. (1982). Models of bounded rationality: Behavioral economics and business organizations. *Cambridge. MA: MIT Press*.

Sims, Christopher A. (2003). Implications of rational inattention. *Journal of monetary Economics*, 50(3), 665–690.

Snelgrove, R. and M. Havitz (2010). Looking back in time: the pitfalls and potential of retrospective methods in leisure studies. *Leisure Studies*, 32(4), 337–351.

Stahl, D., and P. Wilson (1994). Experimental Evidence on Players' models of other players. *Journal of Economic Behavior & Organizations*, 25 (3): 309-327.

Strzalecki, T. (2014). Depth of reasoning and higher order beliefs. *Journal of Economic Behavior & Organization*, 108: 108-122.

Wald, A. (1947). Foundations of a general theory of sequential decision functions. *Econometrica*, 279-313.

Woodford, Michael. (2009). Information-constrained state-dependent pricing. *Journal of Monetary Economics*, 56, S100–S124.

Appendix A

Instructions

In this section we present the instructions for Experiment 1 and Experiment 2 in Chapter 1. The instructions were read out loud, using a projector to show screenshots. Both experiments were divided into four sections, with two sections dedicated to each treatment. As an example, we use the instructions for sessions where the "Human Treatment" (E1-H and E2-H) was played first. In sessions where the "Computer Treatment" (E1-C and E2-C) was played first, Part 1 and 2 were switched be Part 3 and 4.

Experiment 1: instructions (E1-H first and E1-C second)

Thank you for agreeing to participate in this decision making experiment. During the experiment we require your complete, undivided attention. You may not open other applications on your computer, chat with other students, use your cellphone, etc.

You will be paid for your participation in cash at the end of the experiment. What you earn depends partly on your decisions during the experiment, partly on the decisions of others, and partly on chance.

The entire experiment will take place through computer terminals. It is important that you do not talk or in any way try to communicate with other participants during the experiment.

If you have any questions please raise your hand and your question will be answered out loud so everyone can hear. The experiment will be divided into 4 parts. We will now go over the instructions for the first part.

You will play 16 4-player games (4 per part). For each game, you will choose one of three actions. Each other participant in your game will also choose one of three actions.

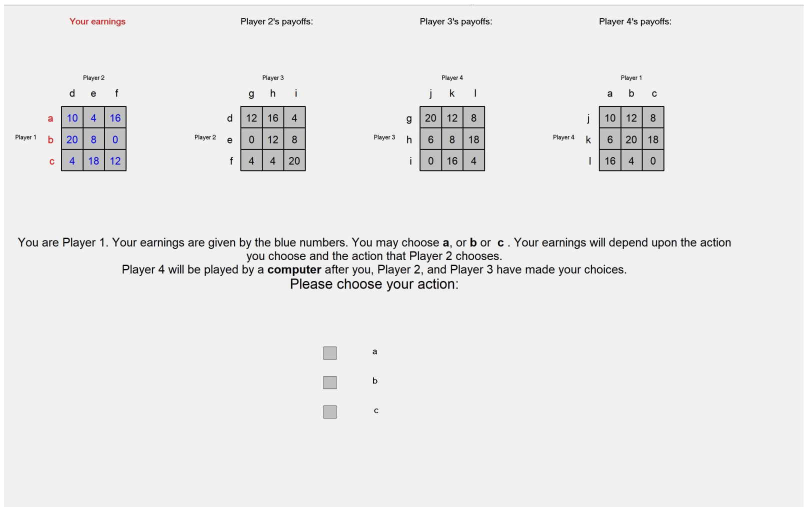

In this example, your are player 1, so your earnings will depend on the combination of your actions and player 2's action.

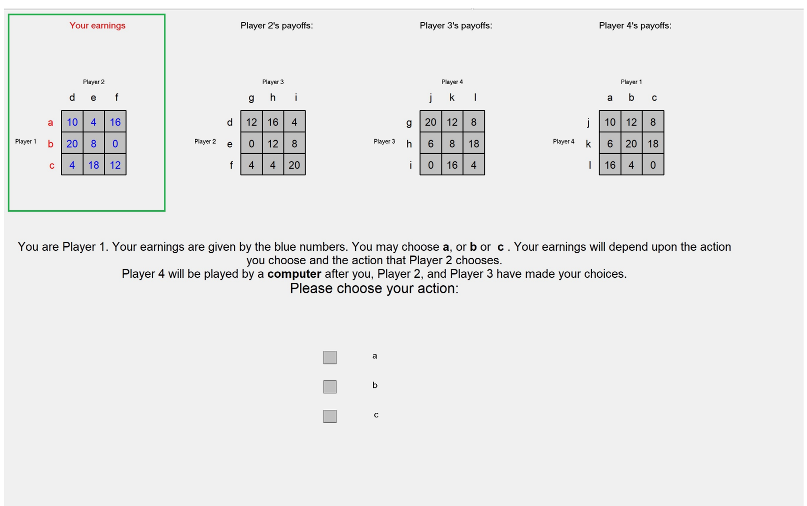

These earnings possibilities will be represented in a table like this one.

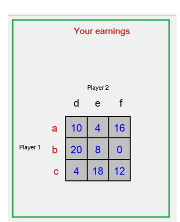

Your action will determine the row of that table and player 2's action will determine the column of the table. You may choose action a, b, or c and player 2 will choose action d, e, or f. The cell corresponding to this combination of actions will determine your earnings.

For example, in this 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.

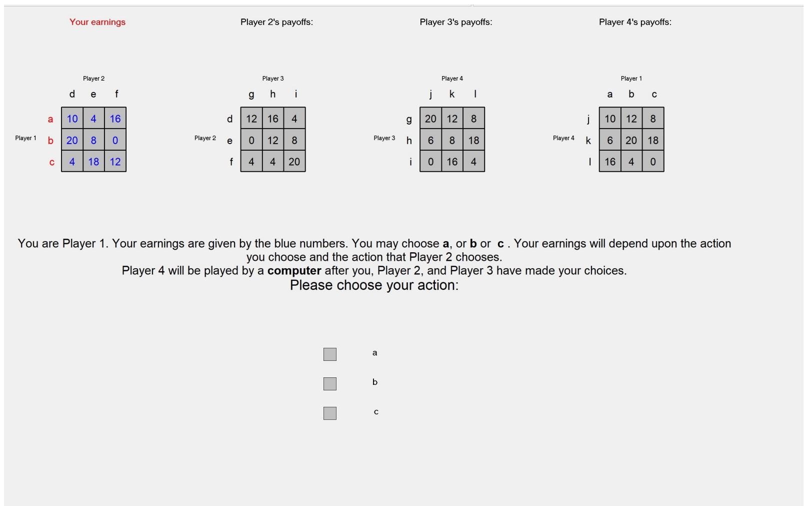

Player 2, Player 3, and Player 4's earnings are listed in the other three tables. Player 2 may choose action d, e, or f, Player 3 may choose action g, h, or i, and Player 4 may choose action j, k, or l. Player 2's earnings depend upon the action he chooses and the action player 3 chooses. Player 3's earnings depend upon the action he chooses and action Player 4 chooses. Player 4's earnings depend upon the action he chooses and the action you, Player 1, choose.

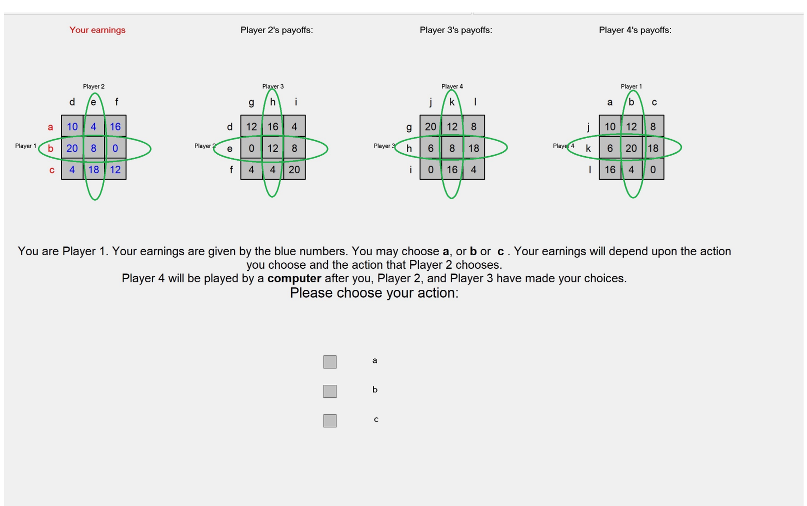

For example, if you choose b, player 2 chooses e, player 3 chooses h, and player 4 chooses k, then you would earn 8 dollars, player 2 would earn 12 dollars, player 3 would earn 8 dollars, and player 4 would earn 20 dollars.

The earnings tables will differ from game to game. So you should always look at the earnings tables carefully at the beginning of each game.

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 will be randomly selected for payment at the end of the experiment. Every participant will be paid based on their actions and the actions of the other players 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.

First Part

In this part of the experiment you will play 4 4-player game. When you start each new game, you will be randomly matched with different participants in this room and play as a different player. We do our best to ensure that you and your counterparts remain anonymous.

You will be informed of your payment at the end of the experiment. You will not learn any other information about the actions of other players in the experiment. The identity of your randomly chosen counterparts will never be revealed.

We will start with a second short quiz in order to help you understand the game.

Second Part

We will now proceed with the second part of the experiments. Good luck!

Third Third

In this part of the experiment you will play 4 4-player games. The other players will be played by computers. Computers will be allowed to observe the decision made by the player which is relevant for their earnings and then choose the action that maximizes their earnings.

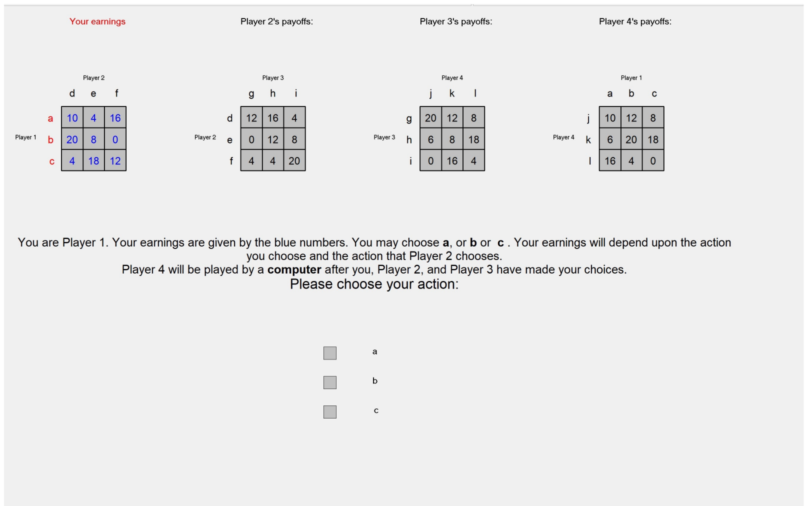

Since in the example you are Player 1, once you have made a decision a computer will see it and then choose as Player 4, since Player 4's earnings depend on Player 1's choice. The computer will choose the option that gives the highest earnings. For example, if you chose a the computer will choose l, which gives 16 dollars, since j would give 10 and k would give 6.

Once the computer has chosen as Player 4, a computer will see this choice and choose as Player 3 the option that gives the highest earnings. Finally, a computer will see this choice and choose as Player 2.

As before, you will not learn about your payoff and the choices made by the other players.

We will now proceed with the third part of the experiment. Good luck!

Fourth Part

We will now proceed with the fourth part of the experiment. For this part you are still playing with computers. Good luck!

Payment

This is the end of the experiment. A screen will now appear with your total earnings for this experiment, as well as your ID number. Please fill up and sign your receipt with this information. You will be called to the next room in the order of your ID numbers, where you will receive the amount that appears on your screen.

Remember that you are under no obligation to reveal your earnings to the other players.

Thank you for your cooperation.

Experiment 2: instructions (E2-H first and E2-C second)

Thank you for agreeing to participate in this decision making experiment. During the experiment we require your complete, undivided attention. You may not open other applications on your computer, chat with other students, use your cellphone, etc.

You will be paid for your participation in cash at the end of the experiment. What you earn depends partly on your decisions during the experiment, partly on the decisions of others, and partly on chance.

The entire experiment will take place through computer terminals. It is important that you do not talk or in any way try to communicate with other participants during the experiment.

If you have any questions please raise your hand and your question will be answered out loud so everyone can hear. The experiment will be divided into 4 parts.

You will play 14 4-player games. In each game, you will choose one of three actions. Each other player in your game will also choose one of three actions. Here is an example of what you could see during the game:

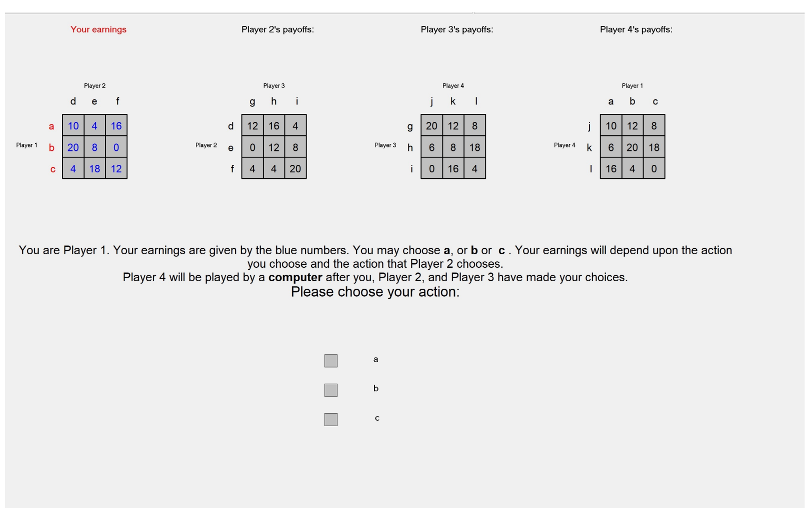

In this example, your are player 1, so your earnings will depend on the combination of your actions and player 2's action.

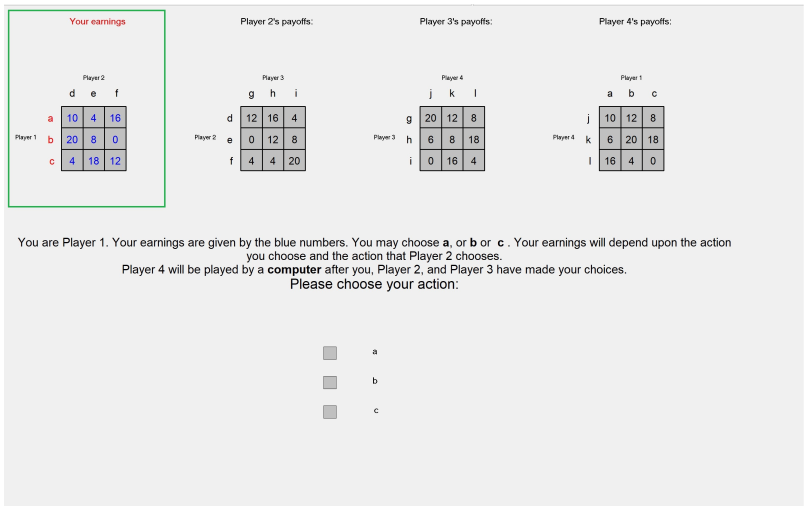

	d	e	f
a	10	4	16
b	20	8	0
c	4	18	12

	g	h	i
d	12	16	4
e	0	12	8
f	4	4	20

	j	k	l
g	20	12	8
h	6	8	18
i	0	16	4

	a	b	c
j	10	12	8
k	6	20	18
l	16	4	0

These earnings possibilities will be represented in a table like this one.

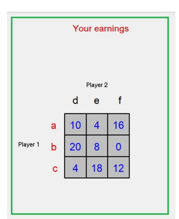

	d	e	f
a	10	4	16
b	20	8	0
c	4	18	12

You may choose action a, b or c; and Player 2 will choose action d, e, or f. Your action will determine the row of the table and player 2's action will determine the column of the table. The cell corresponding to this combination of actions will determine your earnings.

For example, in this 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.

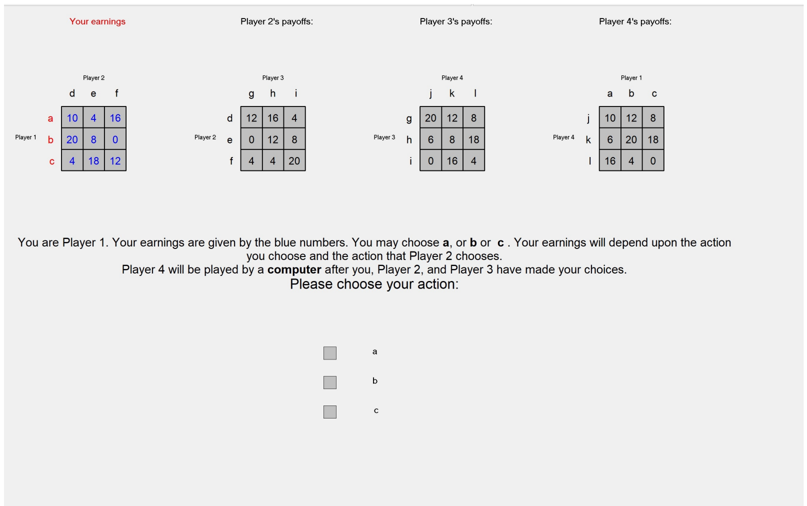

Player 2, Player 3, and Player 4's earnings are listed in the other three tables. Player 2 may choose action d, e, or f, Player 3 may choose action g, h, or i, and Player 4 may choose action j, k, or l. Player 2's earnings depend upon the action he chooses and the action player 3 chooses. Player 3's earnings depend upon the action he chooses and action Player 4 chooses. Player 4's earnings depend upon the action he chooses and the action you, Player 1, choose.

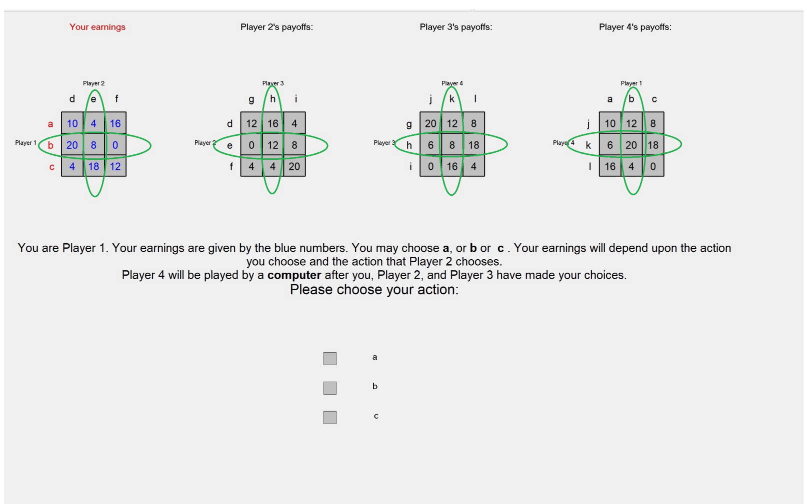

For example, if you choose b, player 2 chooses e, player 3 chooses h, and player 4 chooses k, then you would earn 8 dollars, player 2 would earn 12 dollars, player 3 would earn 8 dollars, and player 4 would earn 20 dollars.

The earnings tables will differ from game to game. So you should always look at the earnings tables carefully at the beginning of each game.

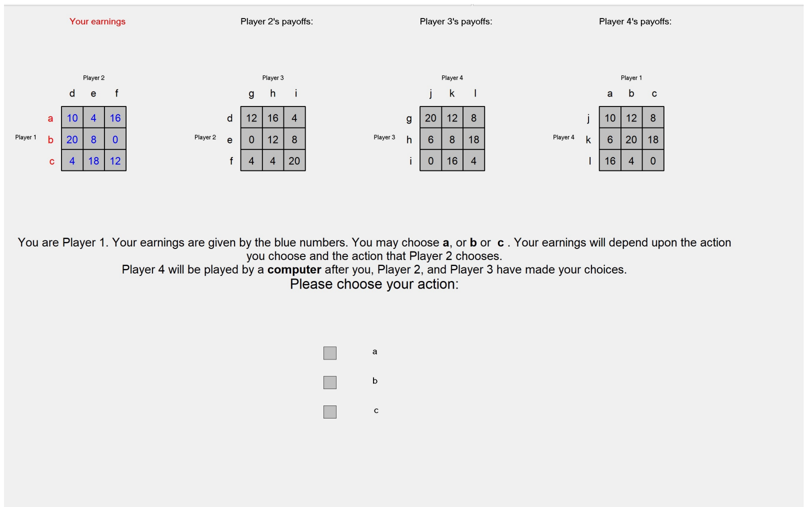

The timing of the game will be as follows. First, Player 1, Player 2 and Player 3 will choose an action. Player 4 will then observe these actions and decide which action to play himself.

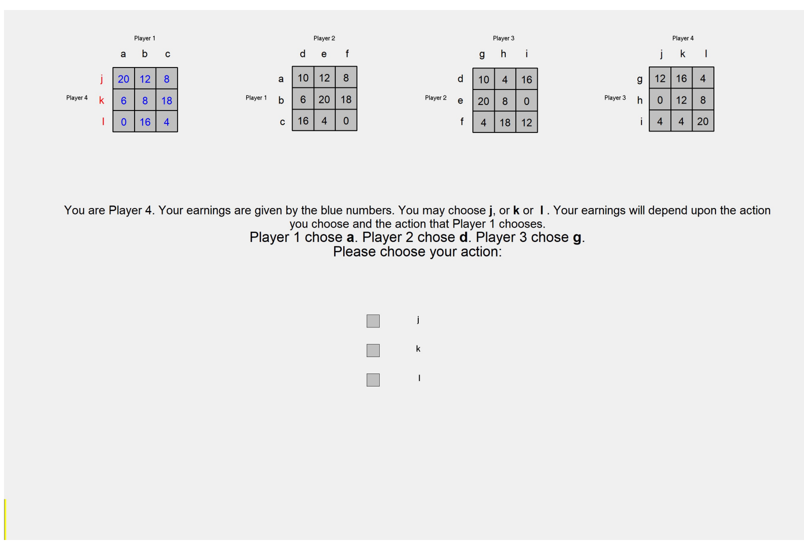

As an example, here is a screenshot of what Player 4 would see if Player 1 chose a, Player 2 chose d, and Player 3 chose g.

You will be require to spend at least 10 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 will be randomly selected for payment at the end of the experiment. Every participant will be paid based on their actions and the actions of the other players 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.

First Part

In this part of the experiment you will play 4 4-player games. When you start each game, you will be randomly matched with 3 other participants in this room and play as Player 1, Player 2, Player 3, and Player 4.

You will be informed of your payment at the end of the experiment. You will not learn any other information about the actions of other players in the experiment. The identity of your randomly chosen counterparts will never be revealed.

Second Part

We will now proceed with the second part of the experiment. As in the previous part, at the beginning of each game you will randomly matched with 3 other participants in this room to play as Player 1, Player 2, Player 3, and Player 4. Good luck!

Third Part

For this part of the experiment you will play 3 4-player games. When you start each game, you will be randomly matched with 2 other participants in this room and play as Player 1, Player 2, or Player 3. For this part, Player 4 will be played by a computer. After observing the decisions made by Player 1, Player 2, and Player 3; the computer will choose the action that maximizes its earnings.

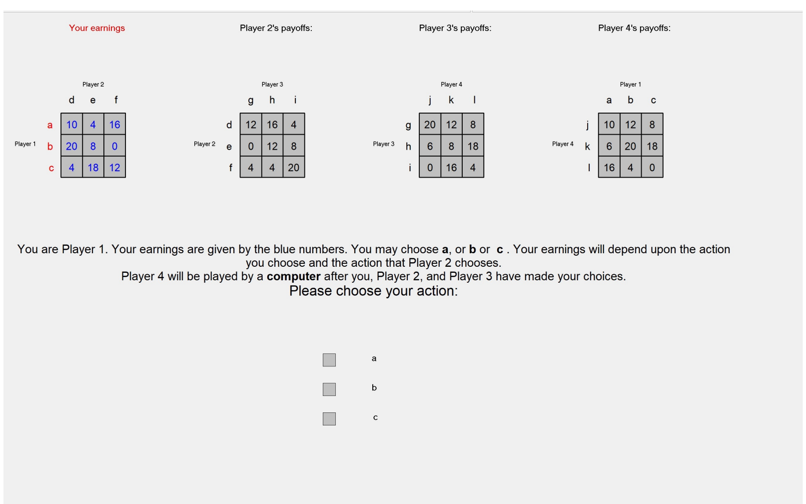

As an example, say you are Player 1 and you choose a, Player 2 chooses d, and Player 3 chooses g. The computer playing as Player 4 will observe these choices and choose l, which gives 16 dollars, since k gives it 6 dollars and j gives 10 dollars.

You will not learn about your payoff and the choices made by the other players.

We will now proceed with the third part of the experiment. Good luck!

Fourth Part

We will now proceed with the forth block of 3 games. For this part Player 4 will still be played by a computer. The computer will observe the decision of Player 1, Player 2, and Player 3 and choose the action that maximizes its earnings. Good luck!

Payment

This is the end of the experiment. A screen will now appear with your total earnings for this experiment, as well as your ID number. Please fill up and sign your receipt with this information. You will be called to the next room in the order of your ID numbers, where you will receive the amount that appears on your screen.

Remember that you are under no obligation to reveal your earnings to the other players.

Thank you for your cooperation.

Appendix B

This appendix refers to Chapter 2.

Robustness: Mechanical Turk

One concern when running experiments in Mechanical Turk is that subjects might not be paid enough to incentivize their decisions or that they are not sophisticated enough, and that results would improve if the experiment was run under more controlled conditions.[2] To address the first concern, for our Mechanical Turk experiments we chose monetary payments that were, on average, more than twice higher than those suggested by the platform. The second concern should be somewhat mitigated by the filters imposed on the quality of our subjects. It should also be noted that some of the subjects participating in this platform are both more dependent on income from experiments and more experienced in them compared to students from a university.

Nevertheless, we run a small batch of sessions using students from Columbia University, and set incentives to be ten times as high as for our Mechanical Turk sessions. We recruited 46 students across 5 sessions, and had them play the game in the Columbia University CELSS computer laboratory. This leaves us with 23 Senders and 23 Receivers, each one of them playing for 10 rounds and 3 treatments. At the end of the experiment, one round was randomly selected from each treatment to count for their monetary payment. Senders and Receivers were paid a show-up fee of 5 dollars, plus 7 dollars if they earned points for that round.

Receivers' choices were further away from the theory than in our main experiment, but the direction of the treatment effects was confirmed: subjects found Treatment 1 easier than Treatment 2, and Treatment 2 easier than Treatment 3. The proportion of subjects that were not classified as having a cutoff strategy was 52% for Treatment 1, 56% for Treatment 2, and 61% for Treatment 3. Figure 3.1 shows the distribution of cutoffs for those subjects that were estimated to have one. Although the size of our sample is small, the general patterns from Section 4 seem to hold. The first treatment has a symmetric distribution of cutoffs, with 18.2% having a cutoff of 1/2, 45.4% a lower one, and 36.4% a higher one. The other two treatments have a

[2]For example, Grether (1980) and Benjamin et. al. (2019) find that base-rate omission bias is reduced when financial incentives are increased.

more asymmetric distribution. The proportion of subject with a cutoff equal, lower, and higher than 1/2 was 20%, 70%, and 10% for Treatment 2; and 55.6%, 11.1%, and 33.3% for Treatment 3.

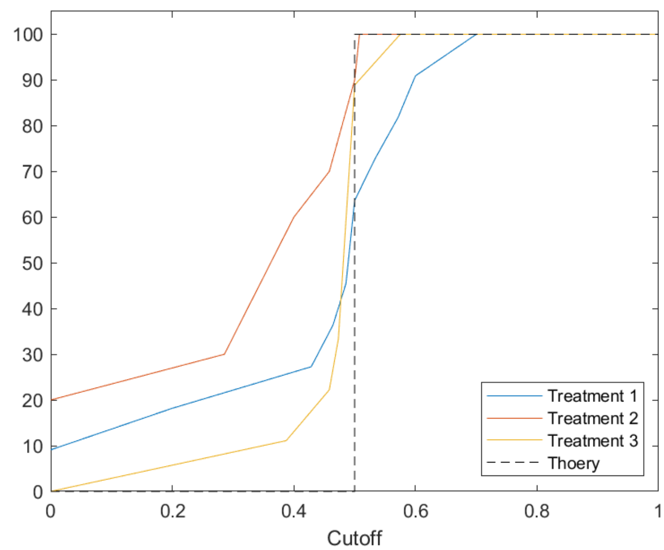

Figure 3.1: Cumulative Distribution of cutoffs
Cumulative distribution of cutoffs by treatments. The doted line represents our theoretical predictions derived from BP-2. A Kolmogorov-Smirnov test for difference in distribution has p-values of 15% for Treatment 1 vs Treatment 2, 55% for Treatment 1 vs. Treatment 3, and 8% for Treatment 2 vs. Treatment 3.

Instructions: Main Experiment, Senders

The following are the instructions for Senders in the main experiment. In this example, the subject is playing Treatment 1 first, Treatment 2 second, and Treatment 3 third. After reading the instructions for each treatment, subjects played the game as a Sender for ten rounds without feedback.

Treatment 1

This section will last for 10 rounds. In each round, there are 10 balls, numbered from 1 to 10. The first 4 of these balls are RED, while the other 6 are BLUE.

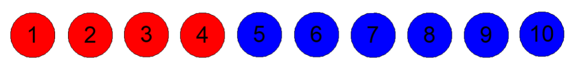

You will have to choose how to arrange these balls into two jars: Jar X and Jar Y.

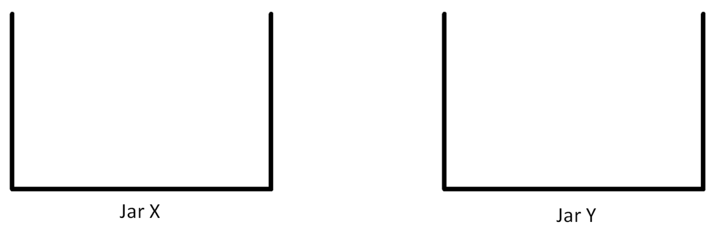

The picture below shows an example of an arrangement, where Jar X has 3 RED balls and 2 BLUE balls, while Jar Y has 1 RED ball and 4 BLUE balls.

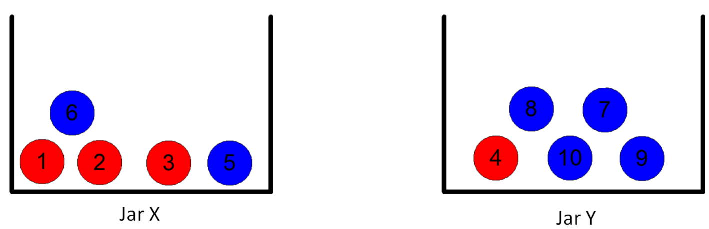

Once you've chosen an arrangement of balls into Jars, one of the 10 balls will be chosen with equal probability. Another MTurk Worker will then be shown how you chose to place the balls into jars and in which Jar the chosen ball is. With this information, he/she will have to guess the color of the chosen ball. This MTurk Worker will earn 70 point if the guess is right and 0 otherwise.

Your payoff will only depend on this guess. You will earn 70 points if the MTurk Worker guesses RED and 0 if he/she guesses BLUE; independently of whether the selected ball was RED or BLUE.

For example, say you choose this assignment and the randomly selected balls is Ball Number 2 (RED in Jar X).

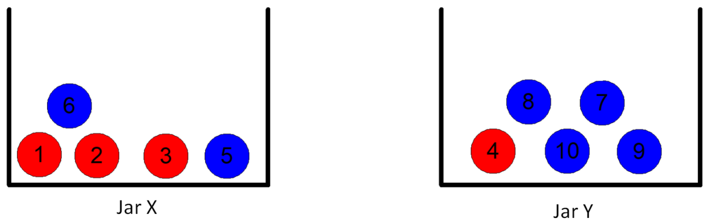

In this case, the MTurk Worker will only be told that the selected ball is in Jar X, and asked to guess the color of the ball. If he/she guesses RED you will earn 70 points, while if he/she guesses BLUE you will earn 0 points for this round.

Once you've completed the experiment, one round from this section will be randomly selected. Your final (bonus) payment will consist of once cent for each point you've earned in the selected round, plus your earnings in the other sections of the experiment. Once you've finished these 10 rounds, you will move to the next section.

REMINDER:
- Once you've decided how to arrange the balls into Jar X and Jar Y, a ball will be randomly selected.
- Another MTurk Worker will be shown your arrangement of balls into jars, and will be told in which Jar the selected ball is. He/she will then be asked to guess the color of the selected ball.
- If the guess is correct, he/she will earn 70 point.
- If the guess is RED you will earn 70 points, independently of the actual color of the selected ball.

Treatment 2

This section will last for 10 rounds. In each round, there are 10 balls, numbered from 1 to 10. The first 4 of these balls are RED, while the other 6 are BLUE.

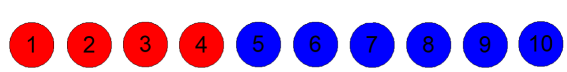

You will have to choose whether to place each of the 10 balls into a small jar labeled "X" or "Y". The picture below shows an example, where 3 RED balls and

2 BLUE balls are in jars labeled "X", while 1 RED ball and 4 BLUE balls are in jars labeled "Y".

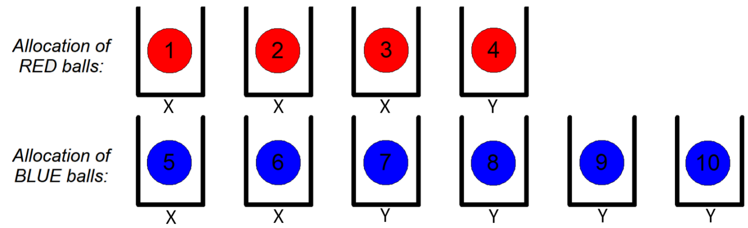

Once you've chosen an arrangement of balls into jar labels, one of the 10 balls will be chosen with equal probability. Another MTurk Worker will then be shown how you chose to place the balls into jars, and whether the chosen ball is in a jar labeled "X" or "Y". With this information, he/she will have to guess the color of the chosen ball. This MTurk Worker will earn 70 point if the guess is right and 0 otherwise.

Your payoff will only depend on this guess. You will earn 70 point if the MTurk Worker guesses RED and 0 if he/she guesses BLUE; independently of whether the selected ball was RED or BLUE.

For example, say you choose this assignment and the randomly selected balls is Ball Number 2 (RED in a jar labeled "X").

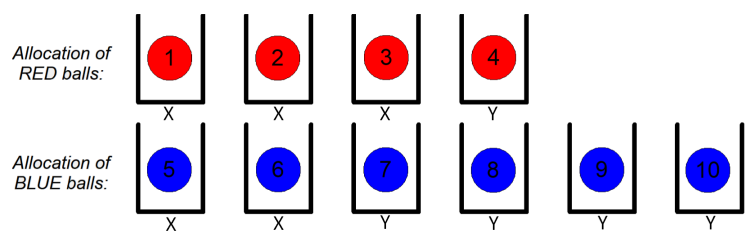

In this case, the MTurk Worker will be told that the selected ball is in a jar labeled "X", and asked to guess the color of the ball. If he/she guesses RED you will earn 70 point, while if he/she guesses BLUE you will earn 0 points for this round.

Once you've completed the experiment, one round from this section will be randomly selected. Your final (bonus) payment will consist of once cent for each point you've earned in the selected round, plus your earnings in the other sections of the experiment. Once you've finished these 10 rounds, you will move to the next section.

REMINDER:

- Once you've decided how to arrange the balls into jars labeled "X" and "Y", a ball will be randomly selected.
- Another MTurk Worker will be shown your arrangement of balls into jars, and will be told in which Jar the selected ball is. He/she will then be asked to guess the color of the selected ball.
- If the guess is correct, he/she will earn 70 point.
- If the guess is RED you will earn 70 points, independently of the actual color of the selected ball.

Treatment 3

This section will last for 10 rounds. In each round, there are 10 balls, numbered from 1 to 10. The first 4 of these balls are RED, while the other 6 are BLUE.

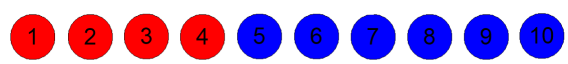

One of these balls will be randomly selected, and this ball will be put in Jar X or Jar Y. You will have to choose with which probability to put this ball into either jar, depending on the color of the selected ball.

The picture below shows an example. In this example, a RED ball is put in Jar X with probability 75% and in Jar Y with probability 25%; while a BLUE ball is put in Jar X with probability 33%, and in Jar Y with probability 67%.

Once you've chosen, one of the 10 balls will be chosen with equal probability. This ball will be put in Jar X or Jar Y, using the probabilities that you have chosen. Another MTurk Worker will then be shown with which probabilities you chose to place the ball into jars, and whether the chosen ball is in Jar X or Jar Y. With this information, he/she will have to guess the color of the chosen ball. This MTurk Worker will earn 70 points if the guess is right and 0 otherwise.

Your payoff will only depend on this guess. You will earn 70 points if the MTurk Worker guesses RED and 0 if he/she guesses BLUE; independently of whether the selected ball was RED or BLUE.

For example, say you choose this assignment and the randomly selected balls is Ball Number 3 (RED).

In this case, the ball will be put in Jar X with probability 75%, and in Jar Y with probability 25%.
The MTurk Worker will be told in which jar the selected ball is, and asked to guess the color of the ball. If he/she guesses RED you will earn 70 points, while if he/she guesses BLUE you will earn 0 points for this round.

Once you've completed the experiment, one round from this section will be randomly selected. Your final (bonus) payment will consist of once cent for each point you've earned in the selected round, plus your earnings in the other sections of the experiment. Once you've finished these 10 rounds, you will move to the next section.

REMINDER:
- Once you've decided with which probabilities to put a RED and BLUE ball into Jar X and Jar Y, a ball will be randomly selected. This ball will be put in Jar X or Jar Y, using your chosen probabilities.
- Another MTurk Worker will be shown your arrangement of balls into jars, and will be told in which Jar the selected ball is. He/she will then be asked to guess the color of the selected ball.
- If the guess is correct, he/she will earn 70 points.
- If the guess is RED you will earn 70 points, independently of the actual color of the selected ball.

Instructions: Main Experiment Receivers

The following are the instructions for Receiver in the main experiment. The instructions for our robustness treatment were identical, except that the parts highlighted in red color did not appear. In this example, the subject is playing Treatment 1 first, Treatment 2 second, and Treatment 3 third. After reading the instructions for each treatment, subjects played the game as a Sender for ten rounds without feedback.

Treatment 1

This section will last for 10 rounds. In each round, there are 10 balls, numbered from 1 to 10. The first 4 of these balls are RED, while the other 6 are BLUE.

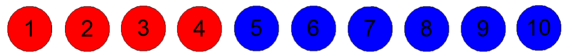

These 10 balls will be distributed among two jars: Jar X and Jar Y. The picture below shows an example of an arrangement, where Jar X has 3 RED balls and 2 BLUE balls, while Jar Y has 1 RED ball and 4 BLUE balls. The allocation of balls into Jar X and Jar Y will have been chosen by another MTurk worker, and it will differ from round to round, so you should play attention to it before making a decision.

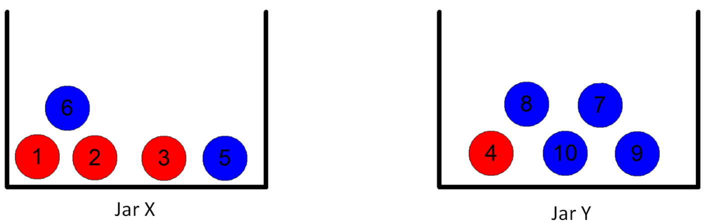

At the beginning of each round, one of the 10 balls will be chosen with equal probability. You will not be told which ball was selected, but you will learn the arrangement of balls into jars chosen by the MTurk Worker. You will also be told in which jar the chosen ball is. For example, if the balls are distributed as shown above and Ball number 3 was chosen, you will be told that the chosen ball is in Jar X, while if Ball number 8 was chosen you will be told that the chosen ball is in Jar Y.

With this information, you will be asked to guess the color of the chosen ball. Your payoff for the round will depend on whether your guess is correct. If you guess RED and the chosen ball was RED, or if you guess BLUE and the chosen ball was BLUE, you will earn 70 points for this round.

The other MTurk worker will earn 70 points if you choose RED, independently of the color of the selected ball. That is, he/she will earn 70 points for the round if you choose RED and the ball is RED, or if you choose RED and the ball is BLUE.

Once you've completed the experiment, one round from this section will be randomly selected. Your final (bonus) payment will consist of once cent for each point you've earned in the selected round, plus your earnings in the other sections of the experiment. Once you've finished these 10 rounds, you will move to the next section.

Now we will begin with this part of the experiment.

Treatment 2

In this section, you will play 10 rounds. In each round, there are 10 balls, numbered from 1 to 10. The first 4 of these balls are RED, while the other 6 are BLUE.

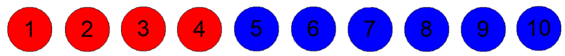

These balls will be distributed among small jars, each one of them fitting only one ball. Each jar will be labeled "X" or "Y". The picture below shows an example of an arrangement. In this example, 3 of the RED balls are in jars labeled X and 1 in a jar labeled Y; while 2 of the BLUE balls are in jars labeled X and 4 in jars labeled Y. The allocation of balls into jars labeled X or Y will have been chosen by another MTurk worker, and it will differ from round to round, so you should pay attention to it before making a decision.

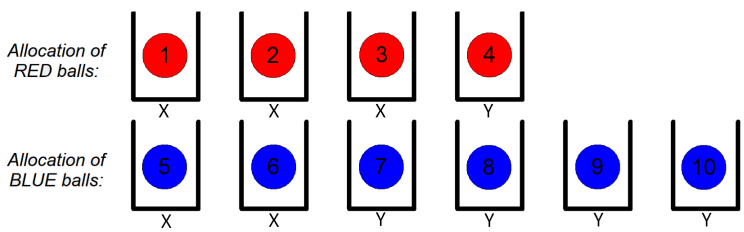

At the beginning of each round, one of the 10 balls will be chosen with equal probability. You will not be told which ball was selected, but you will learn the arrangement of balls into jars labeled X or Y chosen by the MTurk worker. You will also be told whether the selected ball is in a jar labeled X or Y. For example, if the balls are distributed as shown above and Ball number 3 was chosen, you will be told that the chosen ball is in a jar labeled X, while if Ball number 8 was chosen you will be told that the chosen ball is in a jar labeled Y. With this information, you will be asked to guess the color of the chosen ball. Your payoff for the round will depend on whether your guess is correct. If you guess RED and the chosen ball was RED, or if you guess BLUE and the chosen ball was BLUE, you will earn 70 points for this round.

The other MTurk worker will earn 70 points for this round if you choose RED, independently of the color of the selected ball. That is, he/she will earn 70 points for the round if you choose RED and the ball is RED, or if you choose RED and the ball is BLUE.

Once you've completed the experiment, one round from this section will be randomly selected. Your final payment will consist of one cent for each point that you've earned in the selected round, plus your earnings in other sections of the experiment. Once you've finished these 10 rounds, you will move to the next section.

Now we will begin with this part of the experiment.

Treatment 3

In this section, you will play 10 rounds. In each round, there are 10 balls, numbered from 1 to 10. The first 4 of these balls are RED, while the other 6 are BLUE

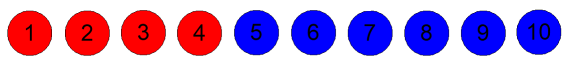

At the beginning of each round, one of the 10 balls will be chosen with equal probability. Next, the ball will be put into Jar X or Jar Y.

The probability that the selected ball is put into each jar depends on its color. The picture below shows an example of what these probabilities could be. In this example, a RED ball is put in Jar X with probability 75% and in Jar Y with probability 25%, while a BLUE ball is put in a Jar X with probability 33% and in Jar Y with probability 67%. These probabilities will have been chosen by another MTurk worker, and they will differ from round to round, so you should pay attention to them before making a decision.

Before you make your decision, you will learn the probabilities chosen by the MTurk worker, as well as whether the selected ball is in Jar X or Jar Y.

With this information, you will be asked to guess the color of the chosen ball. Your payoff for the round will depend on whether your guess is correct. If you guess RED and the chosen ball was RED, or if you guess BLUE and the chosen ball was BLUE, you will earn 70 points for this round.

The other MTurk worker will earn 70 points for this round if you choose RED, independently of the color the the selected ball. That is, he/she will earn 70 points for the round if you choose RED and the ball is RED, or if you choose RED and the ball is BLUE.

Once you've completed the experiment, one round from this section will be randomly selected. Your final (bonus) payment will consist of one cent for each point you've earned in the selected round, plus your earnings in other sections of the experiment. Once you've finished these 10 rounds, you will move to the next section.

Now we will begin with this part of the experiment.

Appendix C

Tables for Chapter 3.

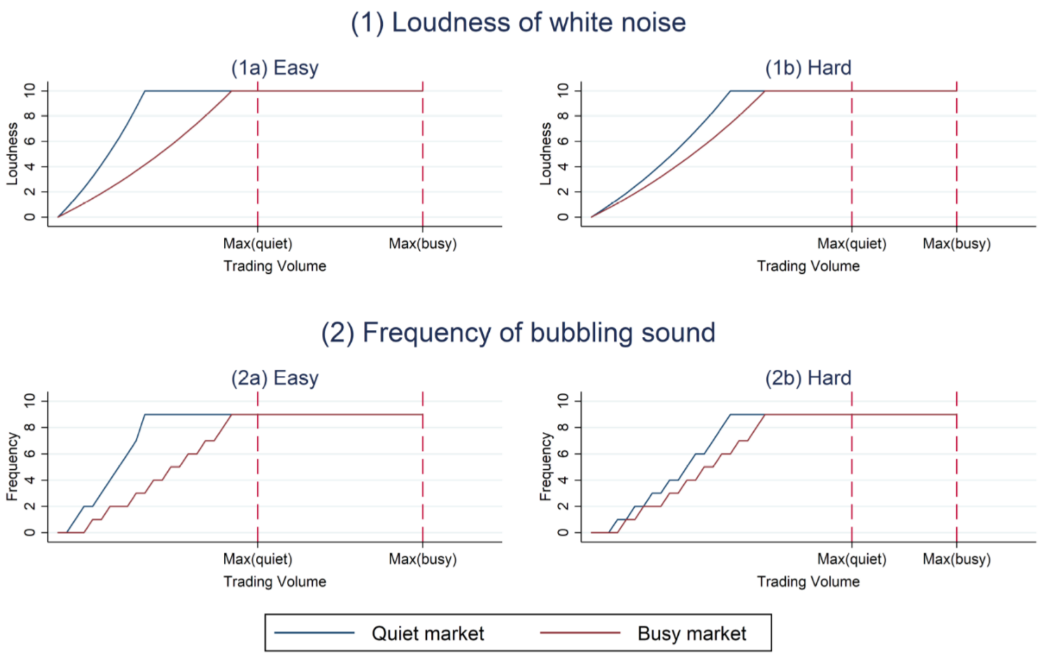

Figure 3.2: Mapping trading volume into auditory form

This figure shows how we map trading volume into auditory form in the experiment. There are two ways of mapping trading volume into auditory form. (1) The function of mapping trading volume into loudness of white noise is $Loudness = 10 \times min(1, e^{\frac{0.693315Q_t}{2\sigma_u}} - 1)$. (2) The function of mapping trading volume into the frequency of bubbling sound is $Frequency = floor(9 \times min(1, e^{\frac{0.693315Q_t}{2\sigma_u}} - 1))$. (1a) (2a) are cases of easy sessions, where $\sigma_{uH} = 10$ in busy market and $\sigma_{uL} = 5$ in quiet market. (1b) (2b) are cases of hard sessions, where $\sigma_{uH} = 10$ in busy market and $\sigma_{uL} = 8$ in quiet market. The market type parameter σ_u influence the auditory information in two channels. First, it is in the mapping function directly, as showed above. Second, the market σ_u also indirectly influence the value of trading volume, Q_t. Max(quiet) and Max(busy) are the average of maximum trading volume in two types of markets. In all four subfigures, quiet market has a steeper slope and a narrower value range, both channels making it less distinguishable than busy market.

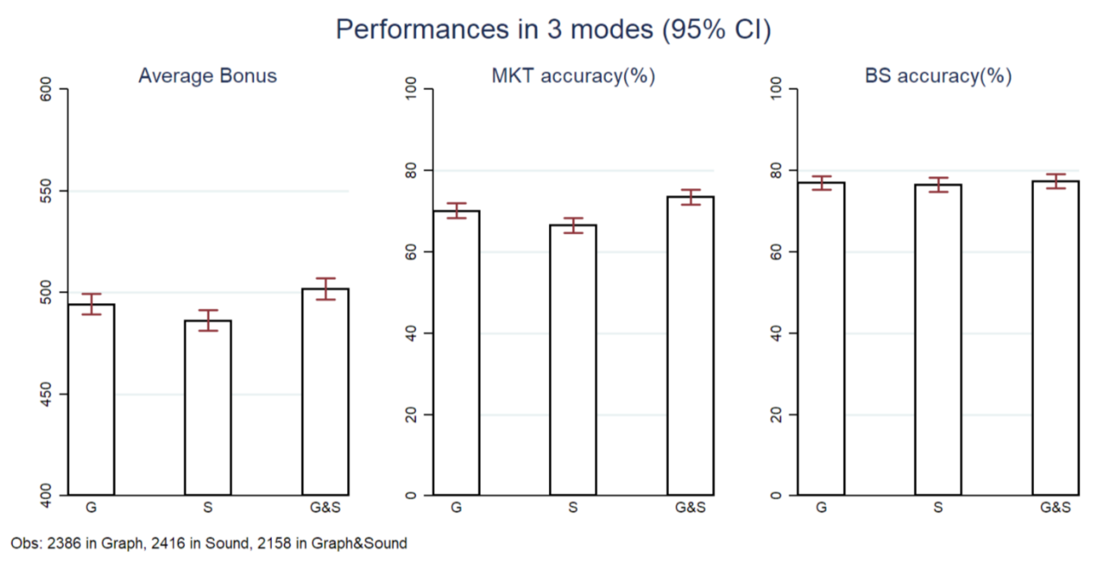

Figure 3.3: Performance in 3 information nodes (95% CI)

This figure compares performances across three informational modes. We first average the mean of round bonus. It could be 600 (if both choices are correct), 400 (if one is correct, the other is wrong) and 200 (if both choices are wrong). The relationship in three modes is Graph&Sound > Graph > Sound, both significant at 5% level. We next look into the accuracy in two market type question ans buying/selling question. The accuracy in market type question follows the same relationship as in average bonus, but there is no evident gap of accuracy in buying/selling question. This shows that the bonus gap is mostly driven by gap in accuracy in market type questions. There are 2386, 2416 and 2156 observations in Graph, Sound and Graph&Sound mode respectively.

Treatment	Group	# of Subjects	Total
Country	Columbia=0	78	232
	Peking=1	154	
Difficulty	Easy=0	103	232
	Hard=1	129	
Sound Content	Bubbling=0	112	232
	White noise=1	120	
Question sequence	Volume first=0	88	232
	Volume second=1	144	

Table 3.1: Cross-subject tratments and sample size

This table shows the composition of subjects in four dimensions: (1) Subjects are recruited from both Columbia University (US) and Peking University (China). (2) Two difficulty levels. In easy sessions, $\sigma_{uH} = 10$ in busy market and $\sigma_{uL} = 5$ in quiet market. In hard sessions, $\sigma_{uH} = 10$ in busy market and $\sigma_{uL} = 8$ in quiet market. (3) Two types of sound content, frequency of bubbling sound and loudness of white noise. (4) Two sequences of questions in each round, volume question first and buying/selling question second, or buying/selling question first and volume question first.

No.	Ses	Date	Country	Difficulty	Sound Content	Question sequence	# of Subjects
1	S1	2018/6/9	China	Easy	White noise	Volume first	24
2	S2	2018/6/9	China	Easy	White noise	Volume Second	22
3	S3	2018/9/22	US	Easy	White noise	Volume first	18
4	S4	2018/10/2	US	Easy	White noise	Volume Second	18
5	S5	2018/10/28	China	Hard	White noise	Volume first	18
6	S6	2018/10/28	China	Hard	White noise	Volume Second	20
7	N1	2019/2/23	China	Hard	Bubbling	Volume first	10
8	N2	2019/2/23	China	Hard	Bubbling	Volume Second	23
9	S7	2019/3/2	US	Hard	Bubbling	Volume first	18
10	S8	2019/3/2	US	Hard	Bubbling	Volume Second	14
11	S9	2019/3/6	US	Hard	Bubbling	Volume Second	10
12	N3	2019/3/16	China	Easy	Bubbling	Volume first	21
13	N4	2019/3/16	China	Hard	Bubbling	Volume Second	16

Table 3.2: Experimental settings in each session

This table shows the detailed information of each session. There are 13 sessions in total. We carryied out the research from June in 2018 to March in 2019. The last column shows the number of subejcts in each session.

Variable	Obs	Mean	SD	Difference in means			
				US-China	Easy-Hard	Bubbling-White noise	V first-V second
Att level	232	0.90	0.16	-0.05	-0.02	0.01	-0.03
Att Scale	232	2.35	0.61	-0.14*	-0.06	0.05	-0.06
Memory	224	6.44	2.10	-0.14	-0.03	-0.12	-0.64**
IQ	232	7.67	1.67	-0.50**	-0.40*	0.14	0.06
IQ Min	232	9.83	3.36	0.17	-0.10	0.37	0.07
Gender	231	0.39	0.49	0.07	-0.07	-0.02	0.07
Age	231	22.41	4.20	1.92***	1.13**	-0.72	-0.63
Edu F	229	3.29	1.34	1.03***	-0.28	0.25	0.45**
Edu M	229	3.11	1.36	1.05***	-0.35*	0.29	0.66***
Exper1	229	1.22	0.42	0.15**	0.07	-0.09	0.02
Exper2	227	1.25	0.46	0.15**	0.11*	-0.12*	0.05

*p<0.1,**p<0.05,***p<0.01

Table 3.3: Summary statistics of individual characteristics

This tables shows the observation, mean, standard deviation as well as difference in mean of individual characteristics. The number of observations in this table is not always 232 because our server happened to miss several records. The first five variables come from cognitive tests. Att_Level is an objective measure of attention level. Att_Scale is a self-reported scale survey of attention level. Memory is a measure of memory. IQ is the score in IQ questions. IQ_Min is the time they spent on IQ test. Among all the subjects, 39.4% are male. Their average age is 22.41. Edu_F and Edu_M are father's and mother's educations levels. Exper1 is a question about their experience in simulated trading. Exper2 is a question about their real experience in stock market. The last four columns tests the differences of means of each variable by four cross-subject variations, so as to check the randomization. Basically, Columbia students are 1.92 years older than Peking students. Their attention level and IQ are also slightly lower than Peking students. Subjects in easy sessions are 1.13 years older and a little less attentive than their counterparts in hard sessions. There is no significant difference across two types of sound content. Subjects who takes market type questions first have somehow a poorer memory. These differences suggest that in the regression, controlling all these personal features will be more reliable than adding individual fixed effects.

Variable	Obs	Mean	SD	Mean in 3 modes			Difference in means	
				Graph	Sound	Graph & Sound	Sound-Graph	Graph & Sound-Graph
Round Bonus	6960	493.85	124.61	494.30	486.26	501.85	8.04**	-7.55**
BS Correct	6960	0.77	0.42	0.77	0.77	0.77	0.00	0.00
MKT Correct	6960	0.70	0.46	0.70	0.67	0.73	0.04***	-0.03**
BS Time	6960	10.82	12.04	10.66	11.21	10.57	-0.55	0.09
MKT Time	6960	6.49	6.63	6.33	6.84	6.28	-0.51**	0.05

*p<0.1,**p<0.05,***p<0.01

Table 3.4: Summary statistics of performances

This table shows the observation, mean, standard deviation as well as difference in mean of subjects' performances. Round_Bonus is the bonus earned in each round. BS_Correct is a dummy of whether the answer is correct or not in buying/selling question. MKT_Correct is a dummy of whether the answer is correct or not in market type question. BS_Time is time spent on this buying/selling question (in seconds). MKT_Time is time spent on market type question (in seconds). We test the difference in mean across three informaion modes, as shown in the last two columns.

	(1)	(2)	(3)	(4)	(5)	(6)
	Round Bonus	BS Correct	MKT Correct	Round Bonus	BS Correct	MKT Correct
Sound	-7.210*	-0.000	-0.183**	-7.489*	-0.010	-0.170**
	(3.860)	(0.077)	(0.076)	(4.027)	(0.073)	(0.074)
Graph&Sound	8.065**	0.091	0.148**	7.995*	0.047	0.171**
	(3.764)	(0.081)	(0.074)	(4.057)	(0.079)	(0.074)
Country	47.949***	0.741***	0.401***	28.424**	0.754***	0.060
	(1.315)	(0.027)	(0.005)	(11.949)	(0.280)	(0.167)
Difficulty	-153.697***	-2.741***	-0.860***	-62.069***	-0.711	-0.925***
	(1.711)	(0.036)	(0.005)	(17.399)	(0.441)	(0.237)
Sound Content	-123.428***	-2.789***	-0.173***	-23.794	-0.565	-0.134
	(1.369)	(0.022)	(0.023)	(16.356)	(0.389)	(0.222)
Sequence	41.406***	1.035***	0.096***	6.709	0.173	0.011
	(1.107)	(0.017)	(0.027)	(10.978)	(0.243)	(0.156)
MKT Time	-1.025***		-0.033***	-0.866***		-0.031***
	(0.211)		(0.007)	(0.214)		(0.006)
BS Time	-0.532***	-0.012***		-0.668***	-0.014***	
	(0.143)	(0.003)		(0.128)	(0.003)	
Constant	608.199***	3.736***	1.135***	515.755***	2.263***	0.617*
	(9.008)	(0.199)	(0.170)	(24.913)	(0.669)	(0.363)
Round FE	Yes	Yes	Yes	Yes	Yes	Yes
ID FE	Yes	Yes	Yes	No	No	No
Session FE	No	No	No	Yes	Yes	Yes
Cognition	No	No	No	Yes	Yes	Yes
Demo	No	No	No	Yes	Yes	Yes
Experience	No	No	No	Yes	Yes	Yes
Observations	6960	6900	6930	6510	6510	6510
R2	0.106			0.042		
Pseudo R2		0.100	0.084		0.019	0.*

Table 3.5: Mode effect (Baseline=Graph)

This table run regressions on round-level. The dependent variables are Round_Bonus (the bonus earned in each round), BS_Correct (whether the answer is correct or not in buying/selling question), MKT_Correct (whether the answer is correct or not in market type question). Regressions on round bonus use OLS regression and buying/selling or market type correction use Probit regression. The independet variables are Sound (whether this round is in Sound mode), Graph&Sound (whether this round is in Graph&Sound mode), Country (Peking University=1, Columbia University=0), Difficulty (Hard=1, Easy=0), SoundContent (white noise=1, bubbling sound=0), Sequence (volume question second=1, volume question first=0), BS_Time (seconds spent on this buying/selling question), MKT_Time (seconds spent on market type question). Round FE is the fixed effect of round number. Session Fe is the fixed effect of session. Cognition are cognitive variables (Att_Level, Att_Scale, Memory, IQ, IQ_Min). Demo are demographic variables (Gender, Age, Edu_F, Edu_M). Experience are variables about subjects' financial experience (Exper1, Exper2). Detailed information about Cognition, Demo and Experience are covered in Table 3. All the standard errors are clustered by individual. Column (1)(2)(3) controls round and individual fixed effects. Column (4)(5)(6) drop individual fixed effects and add session fixed effects as well as control variables, including cognitive capacity, demographics and experience in financial market. Considering some differences in personal features, as we discussed above, we think the last three columns are more reliable estimations.

	Subsample: Columbia			Subsample: Peking			Full sample		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Round Bonus	BS Correct	MKT Correct	Round Bonus	BS Correct	MKT Correct	Round Bonus	BS Correct	MKT Correct
Sound	-13.998**	-0.054	-0.281**	-4.818	0.003	-0.117	-11.002*	0.023	-0.266**
	(6.539)	(0.132)	(0.126)	(5.034)	(0.090)	(0.092)	(6.197)	(0.123)	(0.120)
Graph&Sound	-1.823	-0.050	0.000	12.379**	0.099	0.252***	0.677	-0.013	0.035
	(6.617)	(0.148)	(0.125)	(5.028)	(0.097)	(0.091)	(6.413)	(0.138)	(0.122)
Sound*Country							5.074	-0.050	0.141
							(7.813)	(0.148)	(0.147)
Graph&Sound*Country							10.612	0.089	0.202
							(8.064)	(0.166)	(0.149)
Country							23.510*	0.746**	-0.049
							(12.260)	(0.290)	(0.193)
Difficulty	-33.947**	0.047	-0.878***	-30.590**	0.113	-0.835***	-62.237***	-0.713	-0.928***
	(16.018)	(0.411)	(0.135)	(12.109)	(0.305)	(0.146)	(17.358)	(0.441)	(0.237)
Sound Content				-12.842	-0.725****	0.214*	-23.761	-0.565	-0.133
				(11.292)	(0.269)	(0.124)	(16.319)	(0.389)	(0.222)
Sequence	7.958	0.050	0.164	-3.635	0.182	-0.225**	6.552	0.171	0.008
	(10.386)	(0.226)	(0.152)	(11.411)	(0.289)	(0.112)	(10.913)	(0.243)	(0.155)
MKT Time	-0.704**		-0.018**	-0.933***		-0.038***	-0.859***		-0.031***
	(0.346)		(0.008)	(0.289)		(0.008)	(0.215)		(0.006)
BS Time	-0.210	-0.008**		-0.862***	-0.018***		-0.666***	-0.014***	
	(0.176)	(0.004)		(0.158)	(0.003)		(0.128)	(0.003)	
Constant	508.907***	3.163*	-0.422	492.162***	2.373***	-0.023	519.222***	2.271***	0.689*
	(68.254)	(1.822)	(0.929)	(36.910)	(0.804)	(0.613)	(24.788)	(0.670)	(0.366)
Round FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Session FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Cognition	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Demo	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Experience	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2040	2040	2040	4470	4470	4470	6510	6510	6510
R2	0.045			0.051			0.042		
Pseudo R2		0.038	0.048		0.028	0.060		0.020	0.052

Table 3.6: Mode effect by country (Baseline=Graph)

In column (1)(2)(3) and (4)(5)(6), we replicate regression in Table 5 using subsamples of Columbia university (Country=0) or Peking University (Country=1) sessions. In column (7)(8)(9), we add the interaction term of Sound mode dummy and Country dummy (Sound*Country), as well as Graph&Sound mode dummy and Country dummy (Graph&Sound*Country). Others are the same as Table 5.

	Subsample: Easy			Subsample: Hard			Full sample		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Round Bonus	BS Correct	MKT Correct	Round Bonus	BS Correct	MKT Correct	Round Bonus	BS Correct	MKT Correct
Sound	-4.903	0.065	-0.217	-9.029*	-0.060	-0.146*	-4.661	0.084	-0.228*
	(6.229)	(0.106)	(0.136)	(5.283)	(0.100)	(0.088)	(6.022)	(0.102)	(0.129)
Graph&Sound	13.155*	0.157	0.274**	4.567	-0.022	0.120	12.987**	0.154	0.266**
	(6.874)	(0.140)	(0.137)	(5.118)	(0.093)	(0.089)	(6.424)	(0.133)	(0.125)
Sound*Difficulty							-4.899	-0.167	0.090
							(7.745)	(0.137)	(0.152)
Graph&Sound*Difficulty							-8.750	-0.191	-0.141
							(8.071)	(0.159)	(0.150)
Country	4.535	0.192	-0.080	25.283**	0.722**	0.044	28.351**	0.753***	0.059
	(17.424)	(0.437)	(0.219)	(12.500)	(0.294)	(0.173)	(11.928)	(0.279)	(0.168)
Difficulty							-57.390***	-0.588	-0.915***
							(17.954)	(0.442)	(0.259)
SoundContent	-19.681*	0.003	-0.687***	-21.410	-0.540	-0.093	-23.649	-0.562	-0.132
	(11.443)	(0.278)	(0.218)	(16.424)	(0.376)	(0.222)	(16.328)	(0.389)	(0.222)
Sequence	8.503	0.239	0.027	5.562	0.170	-0.012	6.604	0.171	0.010
	(14.318)	(0.342)	(0.198)	(11.461)	(0.257)	(0.152)	(10.945)	(0.243)	(0.155)
MKT Time	-1.308***		-0.038***	-0.650**		-0.026***	-0.868***		-0.031***
	(0.370)		(0.013)	(0.254)		(0.007)	(0.214)		(0.006)
BS Time	-0.488***	-0.012***		-0.781***	-0.016***		-0.666***	-0.014***	
	(0.174)	(0.003)		(0.183)	(0.004)		(0.128)	(0.003)	
Constant	442.889***	0.072	0.387	477.257***	2.328***	-0.319	512.431***	2.182***	0.604
	(65.621)	(1.542)	(1.039)	(28.006)	(0.783)	(0.416)	(25.003)	(0.668)	(0.371)
Round FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Session FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Cognition	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Demo	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Experience	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2790	2790	2790	3720	3720	3720	6510	6510	6510
R2	0.036			0.027			0.042		
Pseudo R2		0.022	0.048		0.029	0.017		0.020	0.052

Table 3.7: Mode effect by difficulty level (Baseline=Graph)

In column (1)(2)(3) and (4)(5)(6), we replicate regression in Table 5 using subsamples of easy (Difficulty=0) or hard (Difficulty=1) sessions. In column (7)(8)(9), we add the interaction term of Sound mode dummy and Difficulty dummy (Sound*Difficulty), as well as Graph&Sound mode dummy and Difficulty dummy (Graph&Sound*Difficulty). Others are the same as Table 5.

	Subsample: Bubbling			Subsample: White noise			Full sample		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Round Bonus	BS Correct	MKT Correct	Round Bonus	BS Correct	MKT Correct	Round Bonus	BS Correct	MKT Correct
Sound	-6.517	-0.080	-0.082	-7.623	0.068	-0.254**	-5.025	-0.078	-0.053
	(5.259)	(0.105)	(0.097)	(6.058)	(0.102)	(0.113)	(5.077)	(0.103)	(0.093)
Graph&Sound	6.185	-0.003	0.153	11.223*	0.129	0.193*	5.900	-0.024	0.164
	(5.438)	(0.103)	(0.104)	(6.069)	(0.118)	(0.108)	(5.332)	(0.101)	(0.103)
Sound*SoundContent							-4.843	0.130	-0.239*
							(7.612)	(0.138)	(0.140)
Graph&Sound*SoundContent							4.039	0.134	0.014
							(7.880)	(0.151)	(0.144)
Country	30.358**	0.882***	0.015	8.290	0.195	0.102	28.387**	0.752***	0.060
	(13.522)	(0.314)	(0.183)	(18.268)	(0.448)	(0.248)	(11.994)	(0.279)	(0.169)
Difficulty	-19.624	0.840**	-1.338***	-38.437***	-0.111	-0.930***	-61.982***	-0.705	-0.929***
	(17.562)	(0.401)	(0.263)	(10.808)	(0.234)	(0.165)	(17.444)	(0.441)	(0.239)
SoundContent							-23.279	-0.646	-0.053
							(16.852)	(0.393)	(0.237)
Sequence	9.742	0.283	-0.007	-0.665	0.234	-0.213*	6.644	0.171	0.010
	(11.131)	(0.256)	(0.150)	(11.661)	(0.289)	(0.119)	(10.932)	(0.243)	(0.155)
MKT Time	-0.484**		-0.028***	-1.295***		-0.032***	-0.870***		-0.031***
	(0.214)		(0.008)	(0.397)		(0.010)	(0.214)		(0.007)
BS Time	-0.849***	-0.017***		-0.487***	-0.012***		-0.666***	-0.014***	
	(0.186)	(0.004)		(0.162)	(0.004)		(0.129)	(0.003)	
Constant	486.780***	1.257	1.010*	429.043***	-0.112	0.325	515.327***	2.296***	0.585
	(36.198)	(0.989)	(0.594)	(54.709)	(1.077)	(0.766)	(25.087)	(0.669)	(0.367)
Round FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Session FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Cognition	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Demo	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Experience	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	3180	3180	3180	3330	3330	3330	6510	6510	6510
R2	0.050			0.043			0.042		
Pseudo R2		0.032	0.053		0.023	0.053		0.020	0.053

Table 3.8: Mode effect by sound content (Baseline=Graph)

In column (1)(2)(3) and (4)(5)(6), we replicate regression in Table 5 using subsamples of white noise (SoundContent=1) or bubbling sound (SoundContent=0)s essions. In column (7)(8)(9), we add the interaction term of Sound mode dummy and SoundContent dummy (Sound*SoundContent), as well as Graph&Sound mode dummy and SoundContent dummy (Graph&Sound*SoundContent). Others are the same as Table 5.

	Subsample: Volume first			Subsample: Volume second			Full sample		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Round Bonus	BS Correct	MKT Correct	Round Bonus	BS Correct	MKT Correct	Round Bonus	BS Correct	MKT Correct
Sound	-4.694	0.088	-0.183*	-8.407	-0.059	-0.152	-8.536	0.040	-0.242**
	(6.157)	(0.129)	(0.109)	(5.153)	(0.088)	(0.099)	(6.183)	(0.121)	(0.108)
Graph&Sound	6.186	-0.079	0.251**	8.731*	0.114	0.131	4.402	-0.073	0.192
	(7.106)	(0.137)	(0.123)	(4.946)	(0.097)	(0.094)	(6.902)	(0.134)	(0.119)
Sound*Sequence							1.673	-0.079	0.114
							(7.802)	(0.146)	(0.140)
Graph&Sound*Sequence							5.693	0.196	-0.033
							(8.292)	(0.161)	(0.147)
Country	28.421**	0.732**	0.118	12.337	0.345	-0.054	28.265**	0.746***	0.063
	(12.314)	(0.302)	(0.180)	(12.154)	(0.265)	(0.157)	(11.947)	(0.279)	(0.166)
Difficulty	-58.256***	-1.092**	-0.554**	-38.482**	-0.059	-0.858***	-61.913***	-0.703	-0.928***
	(20.176)	(0.506)	(0.223)	(16.597)	(0.376)	(0.223)	(17.417)	(0.441)	(0.237)
SoundContent	-16.801	-0.756***	0.185	-4.459	0.024	-0.113	-23.689	-0.558	-0.137
	(11.907)	(0.274)	(0.127)	(11.484)	(0.276)	(0.120)	(16.377)	(0.390)	(0.221)
Sequence							4.384	0.140	-0.019
							(11.747)	(0.253)	(0.182)
MKT Time	-1.390***		-0.049***	-0.670***		-0.024***	-0.868***		-0.031***
	(0.375)		(0.012)	(0.246)		(0.007)	(0.215)		(0.006)
BS Time	-0.480**	-0.013***		-0.875***	-0.017***		-0.667***	-0.014***	
	(0.211)	(0.005)		(0.161)	(0.003)		(0.128)	(0.003)	
Constant	554.370***	4.127***	-0.167	521.397***	2.028**	0.751	517.099***	2.283***	0.633*
	(43.107)	(1.102)	(0.546)	(40.994)	(0.873)	(0.710)	(25.139)	(0.678)	(0.366)
Round FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Session FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Cognition	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Demo	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Experience	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2400	2400	2400	4110	4110	4110	6510	6510	6510
R2	0.058			0.049			0.042		
Pseudo R2		0.036	0.073		0.024	0.055		0.020	0.052

Table 3.9: Mode effect by question sequence (Baseline=Graph)

In column (1)(2)(3) and (4)(5)(6), we replicate regression in Table 5 using subsamples of volume question first (Sequence=0) or volume question second (Sequence=1) sessions. In column (7)(8)(9), we add the interaction term of Sound mode dummy and Sequence dummy (Sound*Sequence), as well as Graph&Sound mode dummy and Sequence dummy (Graph&Sound*Sequence). Others are the same as Table 5.