The role of advice: an experimental study of Inertia and Reinforcement Heuristics in decision making

Submitted September 2022, in partial fulfilment of the conditions for the award of the degree MRes Economics

Yifan Li

20370414

Supervised by

Abigail Barr

Martin Sefton

School of Economics

University of Nottingham

31/08/2022

Abstract

Social learning, a phenomenon that people observe the behaviour of others to make better decisions, is common in many social and economic situations. However, people learn from others not only by observing their actions but also by seeking their advice, usually from those non-experts such as friends, relatives or families. In this paper, we design and trial a novel experiment to study the effect of advice on two types of decision biases: inertia and reinforcement heuristics. Given a small sample of participants(n=20) in this pilot study, we draw no conclusion from the data analysis. Instead, this paper focuses more on the research questions, experimental design and methods we will use to analyse the data in the formal experiment.

Keywords: Bayesian updating, Inertia, Reinforcement heuristics, Naive advice, Social learning

Acknowledgements

I am deeply grateful for the support and encouragement I have received in the completion of my dissertation. I would like to express my gratitude to my supervisors, Professor Abigail Barr and Professor Martin Sefton, for their guidance and suggestions in the writing of this dissertation, especially their reassuring and encouraging words when I struggled. To Jose Guinot Saporta, the manager of the Cedex lab, I cannot thank him enough for his unwavering support throughout the whole period of experiment preparation and conducting the formal experiment, for continuing to instruct me and providing me constructive advice on my first laboratory experiment. To all the teaching and supporting staff at the School of Economics, I thank them for their work and I have deeply benefited from the knowledge I have learned during my time at the school.

To all my classmates, this journey would be a lot lonelier without their camaraderie and friendship. I'd like to express my thanks to Lucas, Malena and Tong, for supporting my experiments and giving me sincere and helpful advice to improve my study. To all of my friends, I would like to thank them for supporting and participating in my experiments and for always being the first to lend a helping hand when I am struggling.

Lastly to my family, this is not possible without their unconditional support, love, and patience throughout my study and life.

Contents

1 Introduction

In this paper, we design and trial a novel experiment to investigate the effect of advice on two types of decision biases in the individual's decision-making. When it comes to making decisions, in addition to relying on one's own information and experience, people also exhibit a tendency to observe the behaviours of others and then imitate or copy them in order to achieve a more favourable outcome(Bandura et al., 1961). People tend to choose a certain restaurant just because they see a lot of customers inside(Fishman et al., 2019), or buy a certain stock because other stockholders are buying in large numbers(Yao et al., 2014). The tendency to learn from one another by observing their actions is usually called social learning(Bandura, A., & Walters, R. H, 1977).

At the same time, although social learning is a common phenomenon in people's social life, much of social-learning literature are not comprehensive and accurate enough to reflect real social behaviours(Çelen & Kariv, 2005). One of the most important parts is that it does not take into account the fact that people learn by seeking advice as frequently as they learn by observing. High school graduates apply to universities not only by observing what schools their peers are applying to but also by receiving outside advice on which to apply. Consumers buy handbags not only by observing which one the best-seller is but also by asking for recommendations about which handbag their outfit matches better. Furthermore, since not everyone has thirty years of experience working in a university admissions committee or is a top designer, people determine their choices in many cases by relying on advice from non-experts such as families, friends, or even ordinary netizens who have never met. This type of advice, as opposed to professional advice, is named naive advice by Schotter(2003).

Much research has displayed the decision-improving and welfare-improving effects of naive advice in many different game settings. Schotter and Sopher(2003) show that advice boosts the coordination between game participants in coordination games. Chaudhuri & Graziano(2006) find that advice helps sustain high contribution and less free-riding in the public goods game, which leads to higher welfare for everyone involved. Studies on Tournament Games(Merlo and Schotter, 2003) and social-learning games(Çelen et al., 2002; Çelen and Kariv, 2004; Iyengar and Schotter, 2008) also found the beneficial effect of advice on decision-making. However, the findings on the role of advice are not always positive and the possible reasons vary, such as conflict of interest(Mullainathan et al., 2012; Anagol et al., 2017) and low decision quality of advice-giver(Van Swol & Ludutsky, 2007; Li & Zhang, 2022).

Given the previous studies on advice, the effect of advice on decision-making in different contexts is inconsistent and how naive advice affects people's decisions is remained obscure. Thus, the main aim of the research laid out in this thesis is to empirically test the influence of advice on decision biases in the individual's decision-making. More specifically, we are interested in the effect of advice on Inertia and Reinforcement Heuristics in decision-making. Inertia, a term that is widely used to describe the phenomenon that people tend to make the same choice all the time, and reinforcement heuristics, capturing the tendency to repeat the previous successful choice and change the choice that led to failure in the past, are both ubiquitous decision biases exist in people's decision-making process and have been already investigated empirically by many studies(Tykocinski & Pittman, 1995, 1998; Charness and Levin, 2005; Achtziger et al., 2015). To investigate how advice affects these two decision biases, we choose a controlled laboratory experiment as our method since experimental data can more directly reflect and measure the biases in people's decision-making.

In this paper, the experimental design of studying Inertia and reinforcement heuristics is mainly based on the posterior probability task in Charness and Levin(2005) and Alós-Ferrer, Hügelschäfer and Li(2016) respectively, combined with the social-learning experiment adjusted from the experiment of Çelen, Kariv and Schotter(2010). In the experimental design, two participants form a team and make decisions under uncertainty. The decision problem is, in both the Inertia study and Reinforcement heuristics study, to choose either of the two urns containing 6 balls (with the colour of either black or white) to draw a ball and more black balls means higher earnings in the experiment. The configuration of two urns in the Inertia task and Reinforcement Heuristics task follows the designs of Charness and Levin(2005) and Alós-Ferrer, Hügelschäfer and Li(2016). The uncertainty comes from the existence of two unknown world states that appear randomly, with the combination of the 6 balls in the two urns changing according to the state of the world. Hence, subjects were unable to tell which of the two urns had more black balls at first. In each task, participants have two opportunities to draw a ball from the urn and learn the colour of the ball, which also allows the subject to make use of the feedback to update their beliefs about the state of the world.

There are 3 treatments in our experimental design(see section 3 for more information). In the *baseline* treatment, each team member makes their decisions independently without any type of interaction. In the *observation* treatment, team members are assigned to two different roles, where one of the team members can observe another member's choice of the second draw by receiving a message indicating which urn another member chose in the second draw. Similarly, in the *advice* treatment, team members with different roles also have a one-way interaction where one can send a message to another. The main difference between the *advice* treatment and the *observation* treatment is that the message containing the advice or suggestion on the choice of the second draw in the *advice* treatment instead of containing the choice itself in the *observation* treatment.

Compared to the research of Çelen, Kariv and Schotter(2010) and many other studies on social learning and advice(see Section 2 for more information), possible concerns in the experimental design are the perceived quality of observation and advice. Most previous studies use either inter-generational games(Schotter and Barry Sopher, 2003; Celen and Kariv, 2004) where each participant(except the one in the first generation)can only observe the choice or receive advice one time from the immediate predecessor and give advice once to the immediate successor, or other types of games(Merlo and Schotter, 2003;Steinel et al., 2007; Li & Zhang, 2022) where each participant interact with strictly stranger in each round of the experiment. In our experiment, two participants, with fixed roles(either Player A or Player B), stay in the same group during the whole experiment and play the game for several rounds. Thus, with repetition and the feedback of each draw, the action-observers or advice-receivers have the chance to learn whether their partner's choices or suggestions are optimal and helpful, which further influences the effect of observation and advice on decision-making. Meanwhile, advisors or the observed will care about both payment incentive and their reputation when interacting with the same partner for multiple rounds. In addition to this, the posterior probability task in the experiment only involves the process of belief updating with new information instead of the variety of strategies existing in other studies based on strategic games, such as coordination games(Schotter and Barry Sopher, 2003) and cooperation games(Chaudhuri & Graziano, 2006). With binary choices(either correct or wrong) in the posterior probability task, our experiment makes it simpler and more straightforward to study the effect of advice in decision-making.

Based on the features of the inter-generational games, Çelen, Kariv and Schotter(2010) argue that both pieces of messages in the *advice* treatment and the *observation* treatment should be equally informative in equilibrium because the advice from advice-givers should be identical to their real choices. However, they also find the case that advice-givers offer advice different from their actions. In our experiment, this may also happen for advice-giver who don't follow Bayesian updating decision rules. For these participants, there may exist hedging where advice-givers chose one urn but suggest their partners choose another urn, to increase the probability of winning(drawing black balls) under uncertainty. Since this paper is a pilot study and only one participant out of 20 was found to have this hedging strategy, this paper will not pay much attention to the hedge.

In this paper, we pilot the experimental design with a small sample(n=20) to investigate the hypothesis about the effect of social learning and advice on decision errors. Given the sample size, the conclusions we draw from data analyses are not reliable and appropriate. What we find is that decision errors, in all three treatments, are more likely to occur when decision biases and Bayesian updating conflict, which is consistent with previous studies. However, we find not enough evidence that the existence of observation or advice is effective in reducing such errors. A possible explanation for the insignificant difference between the 3 treatments is that the number of subjects was insufficient to draw a valid conclusion about the differences. While we find that there exists a learning effect, the results of data analysis are inconsistent. Another finding which we consider of interest is the effect of the quality of observed decisions and advice. Good-quality observation leads to a much lower likelihood of errors. Nevertheless, there is no evidence for the same result with the existence of advice. Meanwhile, quality also is positively related to the following behaviour of the advice-receiver. Receiving good advice largely increases participants' likelihood of following what advice suggests, but we don't find enough evidence for the positive effect of good-quality observation.

The rest of this paper is organized as follows. In the next section, we summarize more studies on two decision biases, social learning and naive advice. In Section 3 we describe the hypotheses and more detail about the experimental design. The data analysis and results are contained in Section 4, and the conclusions and further discussion appear in Section 5.

2 Literature Review

2.1 Reinforcement Heuristic

In an individual's decision-making under uncertainty, how people update their beliefs by processing and integrating new information is important for making the optimal choice. Following the Bayesian updating rule, rational decision-makers should make use of Bayesian rules to combine previous beliefs with new information when making decisions. However, much experimental research finds evidence that people often pay too much attention to new information but tend to ignore prior information and belief, thus their decisions deviate from what the Bayesian updating rule expects. There are many sources where deviation comes from. 'Reinforcement heuristic', where one tends to repeat the previous successful choice and change the choice that led to failure in the past, is one of the ubiquitous sources of people's biased decisions in the processes of decision-making and human learning.

Reinforcement has gained wide attention in psychology(Thorndike, 1911; Barto & Sutton, 1997) and neuroscience(Holroyd and Coles, 2002; Schönberg et al., 2007;) for a long time. In game theory, scholars also built models to simulate the learning process through reinforcement(Börgers and Sarin, 1997; Erev and Roth, 1998) and treated it as a low-rational behaviour rule. The reinforcement heuristic is one of the simple versions of reinforcement that can be plainly described as a strategy of 'win-stay lose-shift' in human decision-making. This 'win-stay lose-shift' heuristic has been proven to occur very fast in decision-making(Holroyd and Coles, 2002; Schultz, 1998) and is always perceived as an 'automatic process' due to its immediacy, unconsciousness and efficiency. Thus, people in the real world frequently use it to take a shortcut when making decisions, such as investment decisions (Kaustia & Knüpfer, 2008; Chiang et al., 2011) and saving behaviour(Choi et al., 2009).

To study the Reinforcement heuristics in-depth, Charness and Levin(2005) design a binary-choice and belief-updating game to see whether there exist reinforcement heuristics in the individual's decision-making. They find that subjects do follow the reinforcement heuristics when making choices and have higher error rates especially when the expected choices based on reinforcement heuristics are different from those choices based on Bayesian updating rules. That is, people tend to choose sub-optimal choices when reinforcement heuristics are conflicting with optimal Bayesian decisions. Also, they don't find significant evidence to prove that error-rates decline with many-times repetitions. Based on the design from Charness and Levin(2005), Achtziger and Alos-Ferrer(2014) extend the research by changing the information structure. They reconfirm previous findings on error rates. Meanwhile, they also find a small but significant learning effect in the experiment that error rates drop slightly over time. And, subjects make more errors when they make fast decisions in conflicting conditions, which is consistent with previous research in psychology. After that, Achtziger et al., (2015) and Alós-Ferrer et al(2022) further find a negative relationship between high incentives and people's rational choices. As incentive increase, people tend to rely more on the reinforcement heuristic in decision making, which leads to more errors. Alós-Ferrer & Ritschel(2018) find the presence of reinforcement when participants play 3x3 normal form games. Thus, although reinforcement heuristics are widely used and provide a quick shortcut in making choices, it sometimes leads to deviation from optimal choice, especially when this heuristic is conflicting with the rational Bayesian updating process.

2.2 Inertia

Inertia, a term borrowed from physics, is widely used to describe a phenomenon that things without external force will tend to move in the same direction. Inertia has been studied and found evidence in many social and economic fields, such as organizational management(Hodgkinson, 1997; Sull D. N., 1999; Tripsas and Gavetti, 2000), organizational development (Weick & Quinn, 1999), culture and social inertia (Bourdieu, 1985) and political voting(Schram & Sonnemans, 1996). In recent years, Inertia is also found to be built into human behaviours. Corstjens & Lal(2000) show the presence of brand inertia in consumption behaviour. Handel(2013) also finds that consumer inertia about health insurance is common and leads to low welfare of people. Madrian & Shea(2001) analyze the effect of automatic enrollment on 401(k) plans and find that people tend to stick to the default option of automatic enrollment and saving rate due to inertia.

In game theory on decision-making, Schotter and Sopher(2003) investigate coordination convention in the 'inter-generational games' and find that inertia of social conventions exists in both equilibrium and detrimental conditions. Tykocinski & Pittman(1995, 1998) find that when the more attractive option of two options is forgone, participants will not switch to the less attractive option(even though it still has a positive benefit for participants in an absolute sense) due to the fear of regression. Although not the main research question, Charness and Levin(2005) also spot consistency in their binary-choice game. That is, subjects who spontaneously choose the sub-optimal choice in the first place tend to choose the same choice again, compared to those forced to choose the sub-optimal choice. Akaishi et al.(2014) conclude that there exists decision inertia on ambiguous stimuli when people have no performance feedback. Erev & Haruvy(2016) and Alós-Ferrer, Hügelschäfer and Li(2016) also argue that people are often reluctant to change their current choice and have a tendency to maintain the status quo, even current choice is sub-optimal or detrimental. As for the reason for the presence of inertia, Gal(2006) argued that inertia is characterised by a trade-off between not only loss and gain, but also status-quo and change. Erev & Haruvy(2016) state that inertia is reasonable when the cost of deciding is higher than the expected benefit of making a new choice, especially for the decision reached after many times deliberation.

Similar to reinforcement heuristics, inertia is also a product of efficiency. It can lead to quite "rational" choices for the decision maker when people's preferences are unchanged, there exist costs for change or people face uncertainty with the consequences of other choices(Anderson, 2003). However, from previous literature on organization, society and individual, most of the time, the existence of inertia drive people to deviate from the optimal choice or even make a detrimental choice.

2.3 Social learning and Naïve Advice

In real-world economic life, people make decisions not only relying on their own beliefs but also being affected by others. People seek and evaluate more information from others to help them make better decisions. The two most common ways to assist one in making more sensible choices are observing the behaviour of others or getting advice directly from them.

People's behaviours in many situations are influenced when they can see others' decisions. In the studies of daily consumption, an individual's movie consumption is found highly related to others' consumption of movies(Moretti,2011; Gilchrist & Sands, 2016). Chen et al.(2011) also conclude that the desire to consume is further enhanced by observing more sales when people shop online. Duflo & Saez(2002;2003) find out that the choices of enrollment in retirement saving plans are highly affected by the decisions of their colleagues in the same department. In terms of stock market participation, Hong et al.(2004) argue that "social" people are more likely to enter the stock market when more of their neighbours or friends participate in the stock market.

There are broad and varied findings about the relationship between an individual's decisions and the behaviour of others. Banerjee(1992) and Gale(1996) come up with models capturing the behaviour of making decisions by observing others' choices and argue that this may lead to the inefficiency of the equilibrium due to a failure of exploiting one's own information in a more rational way when people are making decisions. Difference evidence for this kind of social learning has also been obtained in the laboratory. Anderson and Holt(1997) conduct an experiment where participants made decisions sequentially and people making the decision later could see their predecessor's choice. Although not all participants make use of their own information perfectly as Bayesian decision-makers, most of them use information efficiently and only choose to follow the decisions of others when it is rational. Similarly, Çelen & Kariv(2004) extend on the experiments of Anderson and Holt(1997) and also find that people tend to follow others' decisions and this kind of behaviour turn out to give them correct choices most of the time. While, the results of Feri et al.(2011) show that when 3 players are playing the Chinos game in sequence, deviations from optimal choices are significant and the probability of making a mistaken choice for a player increases with the probability of the mistakes of the predecessors.

Another common way to help oneself make decisions is to seek the advice of others. Although advice from experts in the relevant fields is usually more constructive, most of the time we get advice from those non-experts such as friends and relatives. Advice has been widely studied in many research fields, such as machine learning(Maclin & Shavlik,1996), marketing sales(Chevalier & Mayzlin, 2006) and psychology(Harvey & Fischer, 1997; Yaniv, I. (2004). These studies have similar results in that people are generally willing to take advice and make use of it in conjunction with people's own experiences to improve judgment accuracy.

In the study of game theory, Schotter(2003) name this word-of-mouth advice as 'naïve advice'. Naïve advice' has been proven to act as a powerful role in helping people update their beliefs and make a decision closer to the prediction of what rational decision-maker would choose, in many experimental studies based on different kinds of strategic games. Schotter and Barry Sopher (2003) design an intergenerational Ultimatum Game where both senders and receivers get advice from their last generation. They find that advice is overwhelming in affecting subjects' behaviour, especially for senders. The experiments show that senders follow advice closely and both the variability and amount of offers decrease significantly. Tournament Game(Merlo and Schotter, 2003), Public Good Game(Chaudhuri & Graziano, 2006), Negotiation Game(Steinel et al., 2007) and social-learning game(Celen et al., 2002; Çelen and Kariv, 2004; Iyengar and Schotter, 2008;) have also been widely studied to find the positive impact of advice on decision-making in different game settings. In terms of why advice can enhance decision-making abilities, Schotter(2003) and Iyengar and Schotter(2008) provide a possible explanation that both giving and accepting advice causes a decision maker to focus their attention on the problem in a way that leads to greater learning or better information processing.

However, a recent study on the effect of advice on decision-making (Li & Zhang, 2022) displays different results. Li & Zhang (2022) conduct experiments based on that of Charness and Levin (2005) to study how advice influences different decision rules in decision-making tasks. They find that the decisions of the advice-receiver are not improved with advice because the advisors are less likely to follow the Bayesian-updating decision rule and tend to give bad suggestions that are not Bayesian-updating optimal choices. Thus, neither the advisors nor the receivers have improved in decision-making with the existence of the advice.

As both observing others and getting advice from others have been widely proven to have an impact on people's decision-making, some scholars begin to delve into these two ways and compare their effects on making decisions. Schotter and Sopher(2003) study the impact of both observing the predecessors' choices and having advice from the predecessors in inter-generational coordination games. The results demonstrate that 'Naïve advice' plays a strong role in helping people make better decisions in creating social conventions by achieving a specific type of coordination faster, than those subjects who only rely on history(by observing predecessors' choices). Çelen et al., (2010) also find that the presence of advice improves the accuracy of the decisions of advice-receivers, but not when the predecessor's action is observed. The reason is that participants are more willing to follow the advice and the tendency to follow increased over time, but they disagree more often with the action they observed and tend not to follow.

Through the review of related literature, we can know that reinforcement heuristics and inertia are often relied upon in people's decision-making, yet their presence can lead to irrational, non-optimal choices in certain situations, especially when they contradict Bayesian decision rules. Little research has been done in the literature on how to reduce the reliance on these two decision biases. Also, since the impact of social learning and advice varies in different contexts and game settings, this paper investigates how the presence of observation and advice influences people's decision-making in posterior probability tasks. Our study contributes to a large literature on advice and social learning by extending the relevant study to different game settings and also provides more hints on how advice and social learning affect people's decision-making.

3 The Hypotheses and the Experiment

3.1 Hypotheses

According to Bayesian updating rules, the basic premise is that all the posterior probability tasks designed in the experiments are choice problems that have right and wrong answers in the second choice. Furthermore, when choices relying on decision biases are different from the choice suggested by Bayesian updating rules, these choices are treated as decision-making errors. Many previous studies have proven the positive effect of naive advice, while the impact of social learning was found to vary in different laboratory studies. Based on the findings of Schotter and Sopher(2003) and Çelen et al.(2010), we make hypotheses on the effect of observation and advice on decision-making.

H1.1: Participants who can observe their partner's choices make fewer decision errors than participants who complete the tasks independently.

H1.2: Participants who receive their partner's advice make fewer decision errors than participants who complete the tasks independently.

H1.3: Participants who receive others' advice make fewer decision errors than participants who can observe their partner's choices.

Although the findings about whether there is a learning effect in previous studies on Inertia and Reinforcement heuristics vary, studies on social learning and advice have indicated that people can learn better with the existence of advice or by observing other's decisions(Merlo & Schotter, 2003; Schotter & Sopher, 2007; Steinel et al., 2007), because people tend to focus more on the problem and deal with the problem more carefully when people can see other people's choices or receive others' advice(Schotter, 2003). Thus, we make a hypothesis here.

H2: Participants who can observe others' choices or receive others' advice exhibit the learning effect that the likelihood of making errors drop over the periods.

One of the main concerns in the experiment is the effect of the quality of others' choices or advice. While many scholars have found that people's decisions can be positively influenced by good behaviour or good advice and vice versa, different studies have differed on the likelihood of people following the advice or choices of others. Merlo & Schotter(2003) and Celen & Kariv(2004) find that people learn faster and better when they observe a good partner than watch a bad one. Anderson and Holt(1997) conclude that people only follow the decisions of others when they are rational. However, Li & Zhang(2022) and Feri et al.(2011) both argue that people tend to follow others even when receiving bad advice or observing bad actions, which makes them worse off. According to previous studies on the effect of quality, we make hypotheses here.

H3.1: Participants who observe good choices or receive good advice make fewer errors than participants who observe bad choices or receive bad advice.

H3.2: Participants who receive good advice or observe good action are more likely to follow it than participants who receive bad advice or observe bad action.

3.2 Experimental design

We conducted three lab-based sessions in the Cedex lab of the University of Nottingham. Twenty university students(13 females and 7 males) participated in the experiment. Each session lasted for 25 minutes on average. Participants were invited to the lab and were given handout instructions explaining the details of the experimental set-up. Participants were also required to correctly answer control questions testing their comprehension of the experiment before the start of each session.

The experiment mainly follows the two-draw decision game(Charness and Levin, 2005) where optimal decisions are derived based on the Bayesian updating rule. In each round, there are two tasks, the Reinforcement Heuristics task(Task1) and the Inertia task(Task2). In each task, there are two urns, the Left Urn and the Right Urn. Under two possible world states(Up and Down), each urn contains 6 balls which are either black or white.

Before the treatment, two anonymous participants are randomly paired to form a two-member group without rematching. Each treatment contains 12 rounds of the game. In each round, one member in each group is assigned the role of 'advisor' and another one is assigned to be 'receiver'. To avoid frame effect, 'advisor' is named as 'Player A' and 'receiver' is named as 'Player B'. Since two different tasks need to be completed in each round, half of the groups in each session are asked to finish Task 1(Reinforcement Heuristics) first and then Task 2(Inertia), while another half need to finish Task 2 first and then Task 1. After assigning roles, both participants are asked to make a binary choice to choose an urn(either Left Urn or Right Urn) to draw a ball in each task. Figure 1 shows the screen that participants faced when making the first draw in Task 1(Reinforcement Heuristics task). They can see the urn configurations in Task 1 and make their decisions by clicking two possible options. Then participants see the colour of a randomly extracted (with replacement) ball from the urn they chose. After the feedback of the first draw, participants are asked to make a second decision of choosing an urn and draw a ball(also with feedback). After finishing the first task, both participants in each group continue the second task with a similar process. Thus, each round of the game contains four draws of balls, two draws in each task. Before the start of each task, the world state(Up or Down) randomly changes with a 50% of chance and members in the same group always face the same world state in each task.

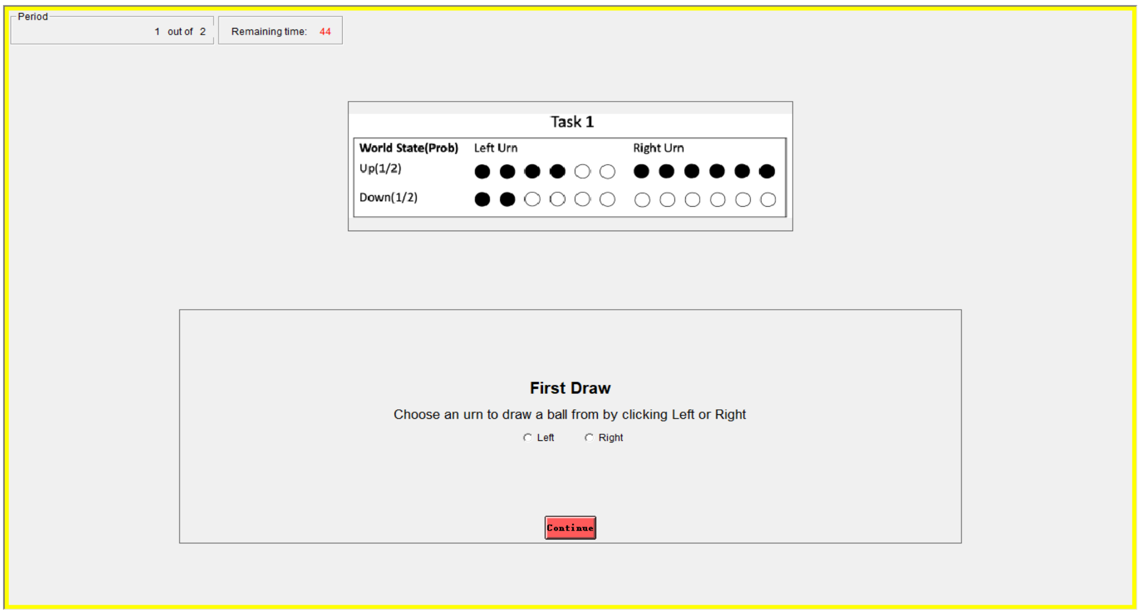

Figure 1: The screen participants faced when making the first draw in Task1

Participants are paid for drawing balls of a predefined colour(black). To make sure that earnings for members in the same group are identical, each time one of two group members draws a black ball, both members in the same group will earn 1 point. Thus, for participants in the same group, their monetary incentives should be the same, as they accumulate the same gains for each task and each round. The total earnings of the participants are the accumulated earnings over 12 rounds of the experiments. Before the first round, participants are asked to read the instruction and finish experiment-related questions to ensure they understand the whole process of experiments properly. At the end of the experiments, participants are asked to fill out questionnaires including questions about their basic characteristics, mathematical ability and knowledge of Bayesian updating rules.

In many previous studies on posterior probability tasks(Charness and Levin, 2005; Charness et al., 2007; Alós-Ferrer et al., 2016), the participants are not allowed to make free decisions in the first few dozen rounds and are forced to choose left in their first draw. These "forced" draws are usually designed for getting more data from the participants, especially in the case of the Reinforcement Heuristics task since choosing the right urn in the first draw provides a shortcut for participants to always choose the optimal choices in the second draw(see Section 3.2.1 for more information). The presence of the forced draws in the first few dozen rounds may cause participants to pay different attention to the left urn and the right urn, and may cause participants to over-guess or pander to the purpose of the researcher and the experiment, which is not conducive to explore the real and unaffected decision-making of the participants. Although Alós-Ferrer et al.,(2016) found that the distribution of errors is not significantly different in the case of forced draws and the case of free draws, we choose not to keep the design of forced draws in this paper due to the concern of its potential effects on participants' decision-making.

3.2.1 Decision tasks

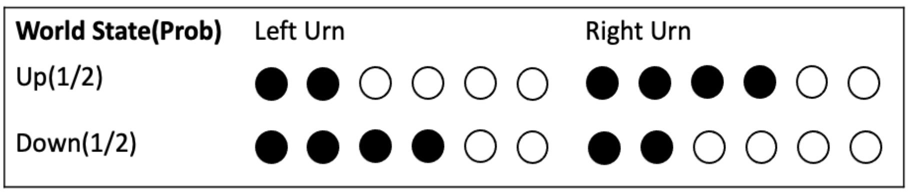

Figure 2: Urn configurations in the Inertia task

Figure 2 shows the design of the urn in the Inertia task, following that of Alós-Ferrer, Hügelschäfer and Li(2016). Each urn contains 6 balls in total. While in the Up state, Left Urn contains 2 black balls and the Right Urn contains 4 black balls. In the Down state, the number of black balls in the two urns is opposite to the Up state.

Due to the concept of Inertia, people always stick to the same urn in both draws, no matter whether they win(draw a black ball) or lose(draw a white ball) in their first draw. In comparison, the Bayesian decision-makers will calculate the probability of each possible state following the Bayesian updating rule after having the feedback in the first draw. For example, when Bayesian decision-makers choose the left urn in the first draw and know the colour of the ball is black. With this new information(colour of the ball), they can update their belief about the probability of state and the participants will realise that it's more likely the state is Down. The probability of "Up" is $(1/2)(2/6)/((1/2)(2/6) + (1/2)(4/6)) = 1/3$, and 'Down' is $(1/2)(4/6)/((1/2)(2/6) + (1/2)(4/6)) = 2/3$. Thus, according to this new belief, Left Urn again delivers an expected winning chance of $(1/3)(2/6) + (2/3)(4/6) = 5/9$, while switching to the Right Urn delivers a smaller expected winning chance of $(1/3)(4/6) + (2/3)(2/6) = 4/9$. Thus, switching to the right urn in the second draw is a suboptimal choice with a lower expected payoff. Conversely, if the Bayesian decision-maker chooses the left urn in the first draw and knows the colour of the ball is white. The subject knows that it's more likely the state is Up. Because the probability of "Down" is $(1/2)(2/6)/((1/2)(2/6) + (1/2)(4/6)) = 1/3$, and 'Up' is $(1/2)(4/6)/((1/2)(2/6) + (1/2)(4/6)) = 2/3$. Thus, according to this new belief, the Left urn again delivers an expected winning chance of $(1/3)(4/6) + (2/3)(2/6) = 4/9$, while switching to the Right Urn delivers a higher expected winning chance of$(1/3)(2/6) + (2/3)(4/6) = 5/9$. Thus, switching to the right urn in the second draw is an optimal choice with a higher expected payoff. Drawing balls from Right Urn in the first draw follows the same calculating rule.

In sum, with the urn configurations shown in Figure 2, the inertia implying participants always choose the same urn conflicts with the Bayesian updating rule when participants lose(draw a white ball) in the first draw without switching to another Urn in the second draw, so lose-stay is not an optimal choice in this case. Meanwhile, inertia aligns with the Bayesian updating rule when participants win(draw a black ball) in the first draw and choose to repeat their choices in the second draw, so win-stay here is the optimal choice while win-switch is an irrational decision.

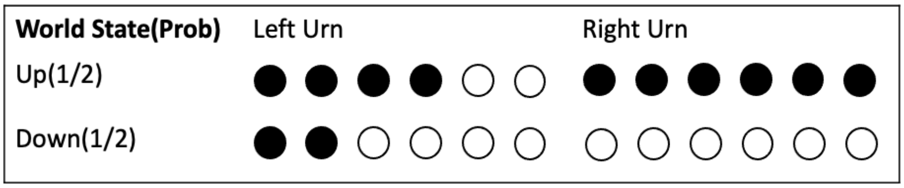

Figure 3: Urn configurations in the Reinforcement Heuristics task

Figure 3 shows the design of the urn in the Reinforcement Heuristics task following Charness and Levin (2005).

For participants who make decisions following reinforcement heuristics, they tend to follow the rule 'win-stay lose-go'. For example, if participants choose Left Urn first and draw a white ball, they tend to switch to Right Urn in the second draw. If they choose Left Urn first and draw a black ball, they tend to stick to the same urn in the second draw. While for the Bayesian decision-maker, when the first draw is black from the left urn, they will follow the Bayesian updating rules and realize it's more likely the state is Up. The probability of Up is $(1/2)(4/6)/((1/2)(2/6) + (1/2)(4/6)) = 2/3$ and the probability of Down is $(1/2)(2/6)/((1/2)(2/6) + (1/2)(4/6)) = 1/3$. According to this new belief, Left Urn again delivers an expected winning chance of $(1/3)(2/6) + (2/3)(4/6) = 5/9$, while switching to the Right Urn delivers a higher expected winning chance of $(1/3)0 + (2/3) = 2/3 = 6/9$. Thus, switching to the right urn in the second draw gives a higher expected payoff than staying at Left Urn. Following the same rule, it's optimal to stick to the Left urn in the second draw if subjects draw a white ball from the Left urn in the first place. Thus, reinforcement heuristics which imply subjects always choose win-stay and lose-go, are conflicting with the Bayesian updating rule in this setting.

In the Reinforcement Heuristic task, it's important to note that the number of black balls in the right urn is much different in each world state, either all black balls or all white balls. This setting gives participants a shortcut to find out what world state it is in each round if they always choose the Right urn in the first draw. If participants have the white ball from the right urn, then it's optimal to switch to the left urn in the second draw, because the white ball from the right urn implies the world state is Down. The expected winning chance is 1/3 for the left urn and 0 for the right urn. On the contrary, participants who draw a black ball from the right urn in their first draw should stick to the right urn, because the black ball implies the Up state and the expected winning chance is 2/3 for the Left urn and 1 for the Right urn. Thus, if participants draw their first ball from the right urn, reinforcement heuristics is consistent with the Bayesian updating rule because "win-stay" and "lose-switch" are optimal choices in this case.

3.2.2 Treatments

Treatments	Actions	Number of Subjects
BS	Making two draws independently	6
OB	Observing the partner's choice in the second draw	6
AD	Giving(receiving) advice in the second draw	8

Table 1: Treatments

There are 3 treatments in our experimental design. All three treatments pass the balance test($p > 0.1$) on the characteristics(gender, math ability and knowledge about the Bayesian updating rule) of participants.

In the *Baseline*(BS) treatment, all paired participants are assigned to either 'Player A' or 'Player B' and finish 12 rounds of tasks independently. The payments are identical within each group. Although there is no interaction within each group, the two participants are still paired together and form a group to make sure they have the same payment incentive compared to the participants in the other two treatments.

Compared to the *Baseline*(BS) treatment, the only difference between the *Observation*(OB) treatment occurs in the second draw. For both participants within a group, they finish their first draw independently same as those in baseline treatment. When it comes to the second draw, 'Player B'(receiver) is able to observe which urn 'Player A'(advisor) chooses in the second draw. Figure 4 shows the screen that Player B faced when making the second draw in OB treatment. Unlike the baseline treatment, player B receives a message about their partner's (player A) decisions in the second draw, in addition to the urn configuration and the decision box. After seeing what 'Player A' chose, 'Player B' is then asked to finish the second draw independently.

The *Advice*(AD) treatment is similar to the *Observation* treatment except observing is replaced by giving advice. The first draw is finished following the basic procedure for both participants within each group. Before the second draw, 'Player A'(advisor) is asked to send 'Player B'(receiver) a message about his/her own suggestion on which urn should be chosen in the second draw. Figure 5 presents the screen Player A faced when giving advice to their partner in the AD treatment. After seeing the colour of the ball they drew in the first draw, Player A is asked to make a suggestion to player B on the urn that

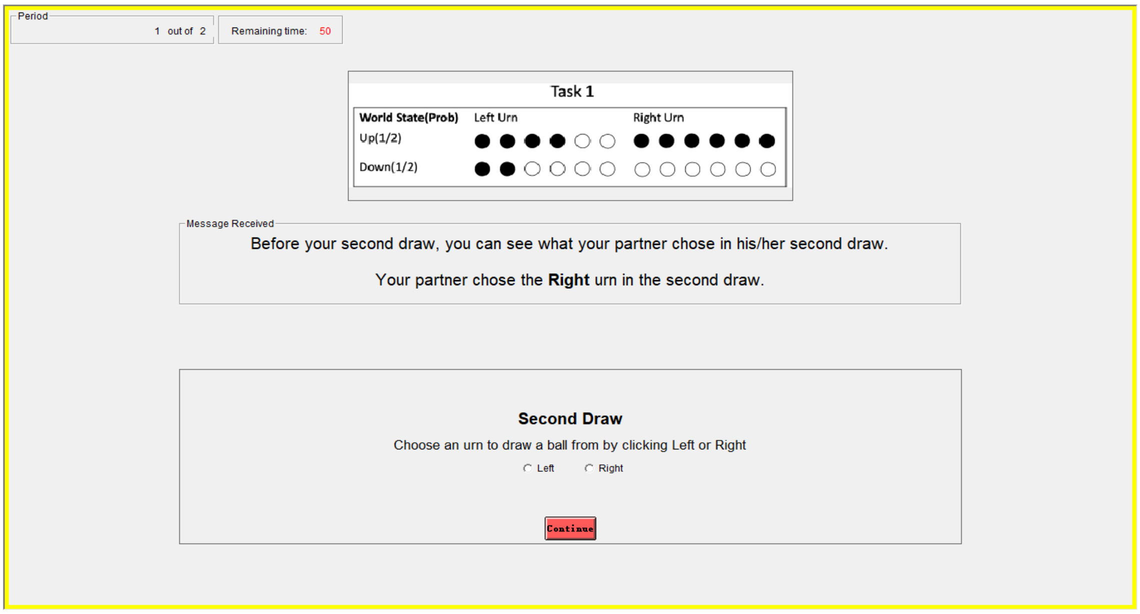

Figure 4: The screen Player B faced when making the second draw in Task1 in the OB treatment

should be chosen in the second draw. The advice is constrained in a pre-designed sentence where "Player A" can only choose 'Left' or 'Right' to be the answer filled into the blank of the message. For Player B, they can see the screen with a pre-designed message shown in Figure 6 in their second draw. After giving(receiving) the advice, 'Player A'('Player B') then finish the second draw and get feedback.

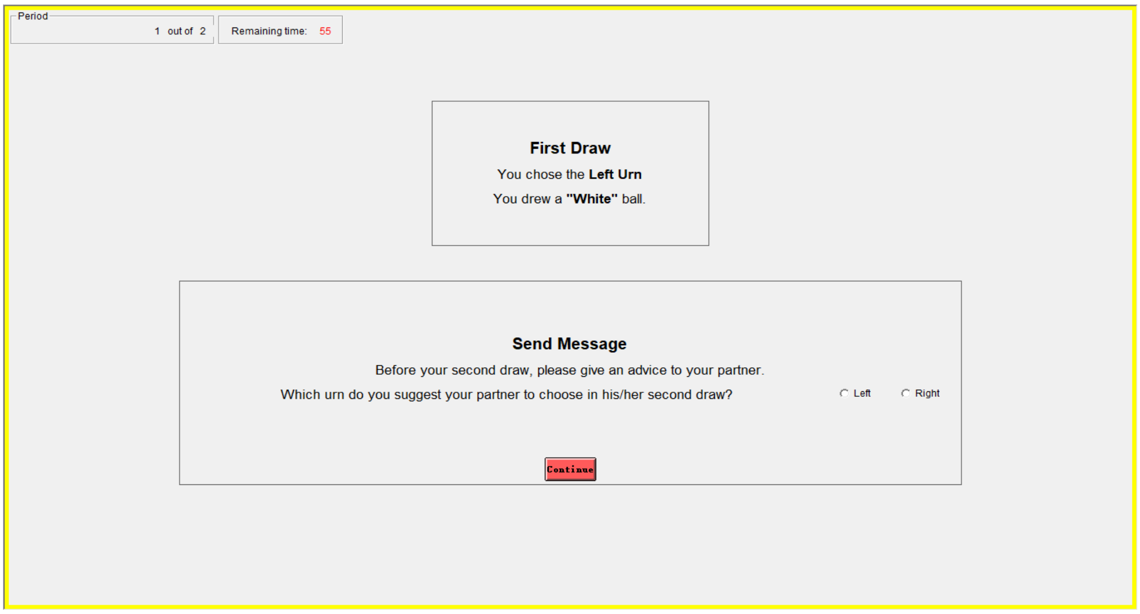

Figure 5: The screen Player A faced when giving advice to their partner in the AD treatment

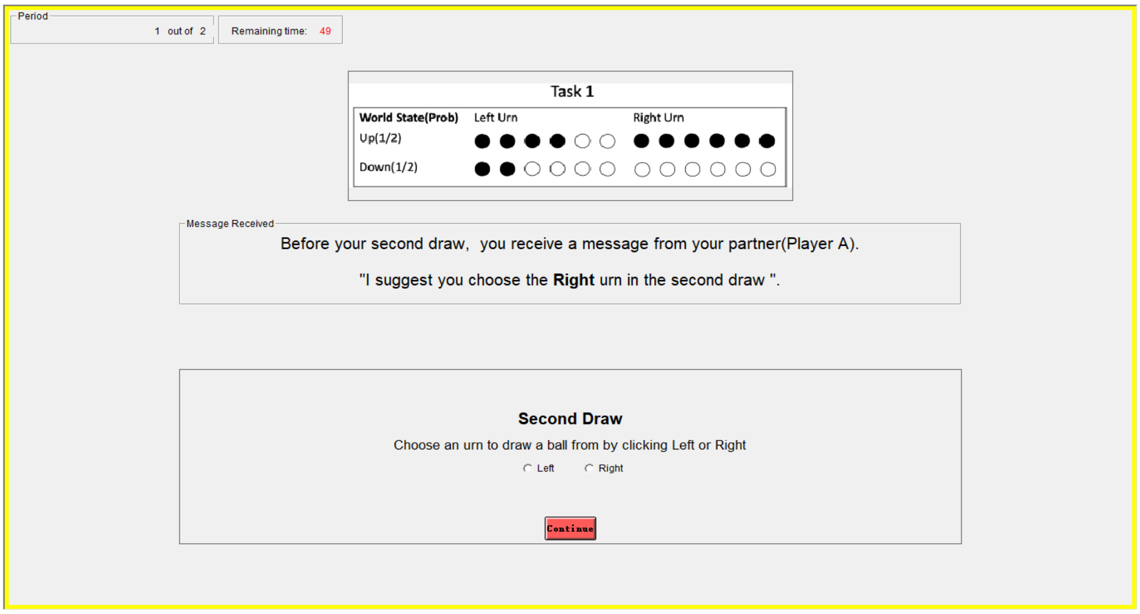

Figure 6: The screen Player B faced when making the second draw in Task 1 in the AD treatment

3.2.3 Belief with observation and advice

While the belief updating of participants following the Bayesian updating rule is mentioned above, their beliefs are different when they can observe others' choices or receive advice from others.

In the Reinforcement task of the *Observation* treatment, if Player B has the first draw from the left urn and the colour of the ball is black, then following Bayesian updating rules, it's more likely the state is Up($P = 2/3$) than Down(P=1/3). When they receive a message from their partner saying that Player A chose the right urn in the second draw, they can obtain more private information from their partners' behaviours. To be more specific, assuming all participants are Bayesian decision-makers, observing Player A choose the right urn in the second draw implies that Player A believes it's more likely that the world state is Up. Player B can deduce in reverse that there are two conditions which can make Player A believe the probability of the "Up" state is higher than that of the "Down" state. One is that Player A chose the left urn in the first draw and got the black ball, which implies that the probability of the "Up" state is 2/3. Another one is that Player A chose the right urn in the first draw and got the black ball, implying that the probability of the "Up" state is 1. Since both of these two conditions can lead to Player A choosing the right urn in the second draw and Player B does not know which urn Player A chose in the first draw, it's reasonable for Player B to assume that his/her partner choose left urn or right urn in the first draw randomly(50%-50%). Thus, for Player B, the message from Player A choosing the right urn in the second draw implies that the probability of the "Up" state is 1/2(2/3+1)=5/6. Since the belief updated from one's own draw(P(Up) = 2/3) and the belief deduced by the partner's choice(P(Up) = 5/6) both indicate that the world state is more likely to be "Up", this perceived belief of the world state from the message aligns with the individual belief of the world state when there is no observation.

However, if Player B has the first draw from the left urn and the colour of the ball is black(P(Up) = 2/3) while observing Player A choose the left urn in the second draw, the perceived belief of the world state from message conflicts individual belief of the world state. Choosing the left urn in the second draw implies that Player A believes it's more likely that the world state is Down. Following the same process of belief updating, the probability of the "Down" state deduced from this message is also 1/2(2/3+1)=5/6, which is inconsistent with the individual belief P(Up) = 2/3.

Similarly, in the Inertia task, if Player B had the first draw from the left urn and the colour of the ball is black, then following bayesian updating rules, it's more likely the state is Down(P = 2/3) than Up(P=1/3). When they receive a message saying that Player A chose the left urn in the second draw, it implies that Player A believes it's more likely that the world state is Down than Up. Following the same process mentioned above, two conditions lead to Player A believing the probability of "Down" is higher than "Up". One is that Player A chooses the left urn and gets the black ball, which updates the belief of Player A that P(Down) = 2/3. Another is that Player A chooses the right urn and gets the white ball, which also updates the belief of Player A that P(Down) = 2/3. Assume that Player A chooses the left urn or the right urn in the first draw randomly(50%-50%). For Player B, the message from Player A choosing the left urn in the second draw implies that the probability of the "Down" state is 1/2(2/3+2/3)=2/3. This perceived belief of the world state from the message aligns with the individual belief of the world state when there is no message. However, when Player B observe Player A choosing the right urn in the second draw, the perceived belief of the world state from this message is P(Up) = 2/3, which is conflicting with the individual belief.

In *Advice* treatment, the advice that Player B receive from Player A should be as informative as the message in *Observation* treatment. To be more specific, the advice that Player A send should be equal to what Player A chose in the second draw if assume Player A is a Bayesian decision maker. Thus, the perceived belief of the world state from the advice is the same as that of Player B who observes the choices of Player A.

The presence of observations and advice may lead to consistency or contrary between an individual's updated beliefs and the individual's perceived updated beliefs of others. Thus, in this paper, we also tend to investigate whether observation and advice have different effects on participants' decision-making behaviour when these two beliefs are contradictory and congruent.

4 Results

In this section, we use three types of empirical analyses to test the above-mentioned hypotheses. First, we present non-parametric results by simply plotting the dependent variables such as the mean error rate. Second, we use non-parametric statistics such as the Mann-Whitney-Wilcoxon (MWW) test and two-tailed Wilcoxon Signed-Rank tests (WSR) to show the simple and straightforward tests of the hypotheses. Finally, we present reduced-form regressions to further test the validity of results from non-parametric statistics.

4.1 Error Rate

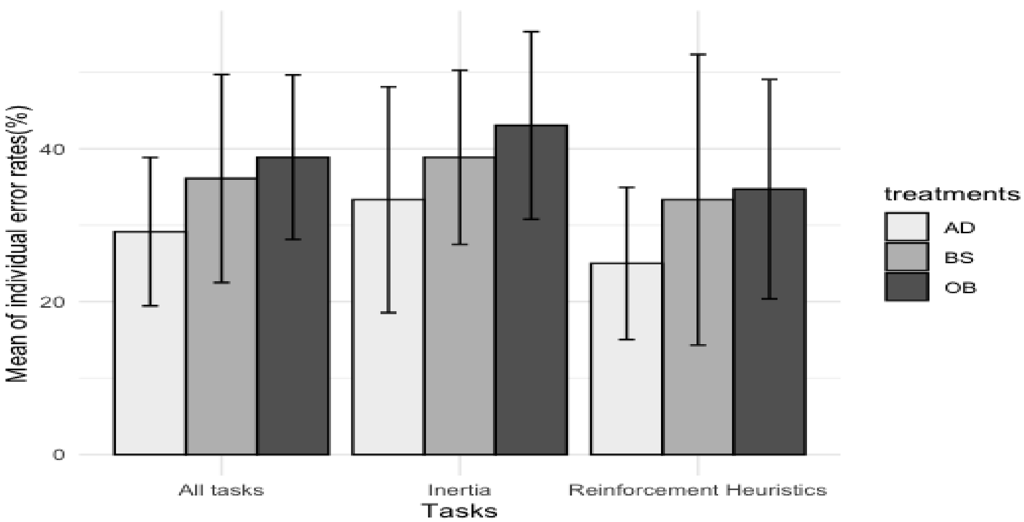

Figure 7: Mean individual error rates by treatments and tasks

The individual error rate is measured as the percentage of the non-bayesian optimal choices selected by each participant. The mean of individual error rates by treatments and tasks are depicted in Figure 7. The total mean error rate including both tasks(Inertia task and Reinforcement task) in the *Advice*(AD) treatment is 29.17%(SD = 9.71%) , vs 38.89%(SD = 10.76%) in the *Observation*(OB) treatment and 36.11%(SD = 13.61%) in the *Baseline*(BS) treatment. In both the Inertia task and Reinforcement task, mean error rates in AD treatment are also lower than that of OD treatment and BS treatment. To test for differences in the distribution of individual-level error rates in different treatments, here we rely on the non-parametric Kruskal–Wallis test which is a multisample generalization of the two-sample Wilcoxon (Mann–Whitney) rank-sum test and used for the analysis of inter-group independent observations. The difference in mean error rates is insignificant when considering both tasks together(N = 20, $\chi^2(2) = 0.381$, p = 0.8266). When we split the test into two tasks separately, the result holds both for the Inertia task(N = 20, $\chi^2(2) = 1.942$, $p = 0.3787$) and the Reinforcement task(N = 20, $\chi^2(2) = 3.010$, is $p = 0.2221$).

Table 2 shows more detailed mean error rates in the case of conflict between decision bias and Bayesian updating(lose-stay errors in Inertia task and errors following right first-draw in Reinforcement Heuristics task) and in the case of alignment(win-switch errors in Inertia and errors following left first-draw in Reinforcement Heuristics task). We then rely on two-tailed Wilcoxon Signed-Rank tests (WSR) to test the difference in distributions of mean error rates here because the Inertia task and Reinforcement heuristics tasks are within-subject designs and observations from these two tasks are dependent. The difference of mean error rates between conflict case and alignment case is significant when consider both tasks together(N = 20, $z = 3.809$, $p < 0.01$), but the result doesn't hold in inertia task(N = 20, If $z = -1.44$, $p = 0.1498$) or Reinforcement task separately(N = 20, $z = 1.382$, $p = 0.1669$).

	win-switch(Inertia)	lose-stay(Inertia)	RH(Left draw)	RH(Right draw)
BS	30%(9/30)	45.2%(19/42)	100%(12/12)	26.7%(16/60)
OB	34.2%(13/38)	52.9%(18/34)	52.2%(12/23)	26.5%(13/49)
AD	31.7%(13/41)	34.5%(19/55)	65%(13/20)	14.5%(11/76)

Table 2: Mean error rates of different types of errors in two tasks

Table 3 depicts the mean error rates of participants assigned to different roles(Player A or Player B) in each treatment. No matter for all treatments or each separate treatment, the difference between the mean individual error rates of Player A and Player B is very small. With independent observations in group 'Player A' and group 'Player B', we again use Mann-Whitney-Wilcoxon (MWW) test to further investigate the difference in the distribution of mean error rates. The results from the MWW test are consistent with the data shown in Table 3. There is no significant difference in mean error rates between two different roles, both when consider all observations($N = 20$, $z = 0.498$, $p = 0.6187$) and consider observations in each treatment individually.

	All treatments	Baseline	Observation	Advice
Player A	35%(84/240)	38.89%(28/72)	37.5%(27/72)	30.21%(29/96)
Player B	33.33%(80/240)	33.33%(24/72)	40.28%(29/72)	28.13%(27/96)

Table 3: Mean error rates of different roles in each treatment

To confirm the stability of results from non-parametric tests and gain further insights on the relationship between decision errors and other factors, we build a regression model on the second-draw decisions(eq(1)) and conduct random-effects probit regression on the strongly balanced panel data collected from the experiments.

$$\begin{aligned} Errors = Conflict + First - draw + Round + Task + OB \\ + AD + Role + Gender + Math + Knowledge \end{aligned} \tag{1}$$

A dummy dependent variable(*Errors*) denotes the existence of second-draw errors, where 0 = Correct and 1 = Error. *Conflict* dummy = 1 implies the case when inertia or reinforcement heuristics is conflicting with Bayesian updating rules, which is win-stay in inertia task and right first draw in reinforcement task. *Conflict* = 0 represents the case that two decision biases align with Bayesian updating rules, lose-stay in inertia task and left first draw in reinforcement task. *First-draw* denotes the color of the ball drawn from the first draw where 0 = white(lose) and 1 = black(win). *Round* is a continuous variable that stands for different rounds of the experiments, ranging from 1 to 12. The *Task* is a dummy variable where 0 = Inertia task and 1 = Reinforcement Heuristics task. *OB* and *AD* both are dummy variables for the *observation* treatment and the *advice* treatment. *Role* = 0 denotes Player A and Role = 1 denotes Player B. Dummy variable *Gender* = 0 stands for female and 1 stand for male. *Math* is the dummy variable measuring the math ability of participants, where 0 = Poor, 1 = Fair and 2 = Good. *Knowledge* is also a dummy variable capturing participants' knowledge of Bayesian updating rules. *Knowledge* = 0 implies no knowledge of Bayesian updating rules. *Knowledge* = 1 implies participants having an understanding of Bayesian updating rules.

Table 4 reports the results from the random-effects probit regression in controlling for personal characteristics(gender, math ability and knowledge about the Bayesian updating rule). The findings are consistent with the results of the above non-parametric tests. First, there shows no significant difference in the likelihood of second-draw errors between BS treatment with OB and AD treatments(both $p > 0.1$). The distribution of second draw errors in two different tasks(Inertia and Reinforcement heuristics) also exhibits no significant difference($p > 0.1$). Second, the conflict between decision biases and Bayesian updating rules shows a large and highly significant effect on increasing the likelihood of second-draw errors for participants in all three treatments($\beta_{Conflict} = 0.792$, $p < 0.01$) and each treatment separately(all with $p < 0.01$). Furthermore, we control for repeated rounds in the experiment and find a small but highly significant dropping in second-draw error rates over the periods($\beta_{Round} = -0.098$, $p < 0.01$), which indicates a possible existence of learning effect. The math ability of participants has an effect on reducing errors. Higher math ability shows less likelihood of errors when considering all three treatments together($\beta_{Math} = -0.338$, $p < 0.01$).

4.2 Learning Effect

The above regression on second-draw errors shows a possible existence of the learning effect over periods. In this section, we tend to gain more about the learning effect. Figure 8 depicts the trends of mean error rates in each treatment, with each observation for every three rounds. It's clear that mean error rates gradually decrease over time in BS treatment and OB treatment. But in AD treatment, mean error rates decrease largely in the first 9 rounds and increase again in the last 3 rounds. To further test the existence of the learning effect, we split observations from the first 4 rounds and the last 4 rounds into two groups. Here, we again use two-tailed Wilcoxon Signed-Rank tests (WSR) to investigate the difference in individual error rates between the first and the last 4 rounds. The WSR test result shows a highly significant difference in mean error rate when consider all participants (N = 20, $z = 3.158, p < 0.01$). The results hold for *Observation* treatment(N = 6, $z = 2.003, p = 0.0452$) and *Advice* treatment(N = 8, $z = 2.345, p = 0.0190$). When all 12-round observations are split into two groups, the first 6 rounds and the last 6 rounds, the WSR test result also implies a significant difference in mean error rate when considering all three treatments(N = 20, $z = 2.804, p < 0.01$),

Variables	All	BS	OB	AD
OB treatment	-0.075			
	(0.204)			
AD treatment	-0.262			
	(0.192)			
Conflict	0.792^{***}	0.756^{***}	0.778^{***}	0.861^{***}
	(0.044)	(0.086)	(0.085)	(0.064)
First-draw	0.132	0.043	0.127	0.346
	(0.161)	(0.28)	(0.292)	(0.284)
Round	-0.098^{***}	-0.121^{***}	-0.067^{*}	-0.107^{***}
	(0.026)	(0.042)	(0.04)	(0.042)
Task	-0.022	-0.294	0.074	0.233
	(0.155)	(0.285)	(0.276)	(0.267)
Role	-0.036	-0.19	0.223	-0.355
	(0.167)	(0.465)	(0.33)	(0.272)
Gender	0.317	0.454	0.254	-0.056
	(0.219)	(0.481)	(0.608)	(0.336)
Math	-0.338^{***}	-0.113	-0.363	-0.052^{*}
	(0.125)	(0.307)	(0.227)	(0.316)
Knowledge	-0.038	0.325	-0.51	0.566
	(0.185)	(0.39)	(0.323)	(0.41)
Number of observations	480	144	144	192

Significance level: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$
Notes: The table shows the marginal effect and standard error(in parentheses) of random -effect probit regression. The dependent variable is a dummy variable indicating second -draw error and the independent variables are shown in the first column named Variables. The second column(All) shows the regression results for all observations in three treatments. The regression results for observations in each treatment separately are shown in the BS, OB and AD columns.

Table 4: Regression on second-draw errors

Observation treatment(N = 6, $z = 1.992, p = 0.0464$) and *Advice* treatment separately (N = 8, $z = 1.829, p = 0.0673$).

To further study the learning effect, we build a regression model(eq(2)) including a dummy variable *Last* where 0 implies the observations from the first 4 rounds and *Last* = 1 captures observations from the last 4 periods, an interaction term(*Role* & *Last*) capturing the relationship between role and learning effect. Interaction variables(*OB* & *Last* and *AD* & *Last*) indicate the interaction effect between different treatments and learning.

$$\begin{aligned} Errors = Conflict + First - draw + Round + Study + OB + AD + OB\&Last \\ + AD\&Last + Role + Role\&Last + Gender + Math + Knowledge \end{aligned} \tag{2}$$

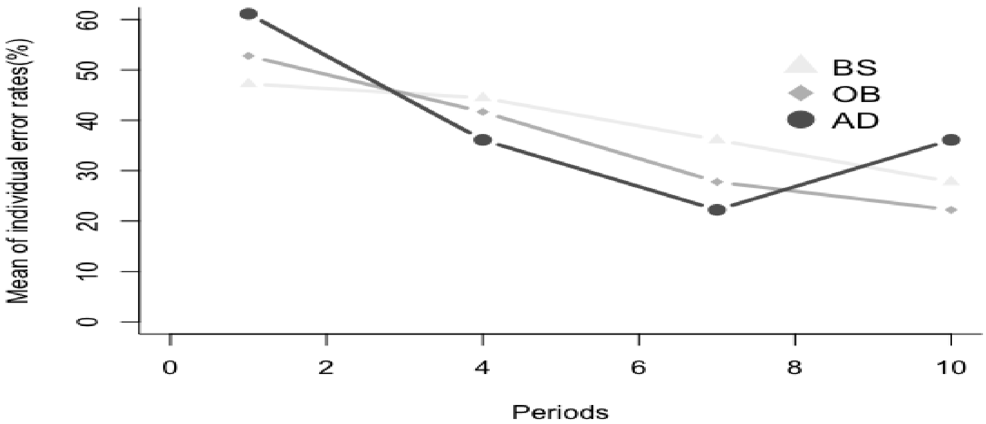

Figure 8: Mean error rates over periods by treatments

Table 5 presents the results of the random-effect probit regression on second-draw errors. The results are incongruent with the WSR test for the learning effect. The regression shows a significant difference in the likelihood of errors between the first 4 rounds and the last 4 rounds considering all participants($\beta_{Last} = -1.071$, $p < 0.01$). The results also hold for observations in *OB* treatment($\beta_{Last} = -0.907$, $p = 0.08$) and *AD* treatment ($\beta_{Last} = -0.961$, $p = 0.024$). There is no significant effect of the interaction between different role assignments($\beta_{RoleLast} = -0.190$, $p > 0.1$) and the two periods on the probability of errors. However, when the dummy variable *Last*=0 implies the observations from the first 6 rounds and 1 denotes the observations from the last 6 rounds, the regression results are not consistent with non-parametric tests, which show no existence of a learning effect($\beta_{Last} = -0.789$, $p > 0.1$).

4.3 Peer Effect

In the OB treatment, player B is able to see his partner's choices in the second draw. In the AD treatment, player B receives the advice given by player A about the decisions for the second draw. This section aims to investigate whether the quality of the advice or the observed choices affects decision errors. *Quality* here is a dummy variable(0 = Bad, 1 = Good) measured as whether Player A in OB treatment makes optimal decisions following Bayesian updating rule in the second draw or whether Player A in AD treatment gives optimal advice following Bayesian updating rule. For example, if player A chooses the left urn in the first draw in the Reinforcement task and draws a black ball, then the optimal

Variables	All	BS	OB	AD
OB treatment	-0.399			
	(0.333)			
AD treatment	-0.223			
	(0.298)			
OB treatment & Last	0.393			
	(0.475)			
AD treatment & Last	0.085			
	(0.463)			
Conflict	0.796^{***}	0.767^{***}	0.775^{***}	0.865^{***}
	(0.443)	(0.088)	(0.086)	(0.064)
First-draw	0.043	-0.028	-0.053	0.283
	(0.199)	(0.357)	(0.391)	(0.329)
Last	-1.071^{***}	-0.732	-0.907^{*}	-0.961^{**}
	(0.391)	(0.492)	(0.517)	(0.426)
Task	0.279	0.064	0.504	0.418
	(0.191)	(0.350)	(0.359)	(0.327)
Role	-0.190	0.014	0.008	-0.612
	(0.255)	(0.5)	(0.531)	(0.407)
Role & Last	0.248	-0.399	0.718	0.199
	(0.379)	(0.557)	(0.687)	(0.645)
Gender	0.326	0.847	-0.126	-0.139
	(0.222)	(0.609)	(0.713)	(0.341)
Math	-0.458^{***}	-0.414	-0.376	-0.707^{*}
	(0.153)	(0.371)	(0.228)	(0.378)
Knowledge	0.008	0.029	-0.213	0.581
	(0.229)	(0.489)	(0.394)	(0.41)
Number of observations	254	76	76	102

Significance level: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Notes: The table shows the marginal effect and standard error(in parentheses) of random-effect probit regression. The dependent variable is a dummy variable indicating second-draw error and the independent variables are shown in the first column named Variables. The second column(All) shows the regression results for all observations in three treatments. The regression results for observations in each treatment separately are shown in the BS, OB and AD columns.

Table 5: Learning effect: regression on second-draw errors

choice following the Bayesian updating rule in the second draw should be right. Thus, choosing the right urn in the second draw in OB treatment or suggesting Player B choose the right urn in AD treatment are regarded as good quality.

About the consistency and inconsistency between an individual's updated beliefs and the individual's perceived updated beliefs of the partner(mentioned in section 3), they can also be represented by qualities. For Player B who follows Bayesian updating rules in decision-making, the good-quality message is always consistent with the individual's updated beliefs, while the bad-quality message always conflicts with the individual's updated beliefs. Thus, by exploring the effect of the quality of observation or advice on Player B's errors in decision-making, we can also learn how errors distribute when two different beliefs(individual's beliefs and perceived beliefs of others) coincide and contradict each other.

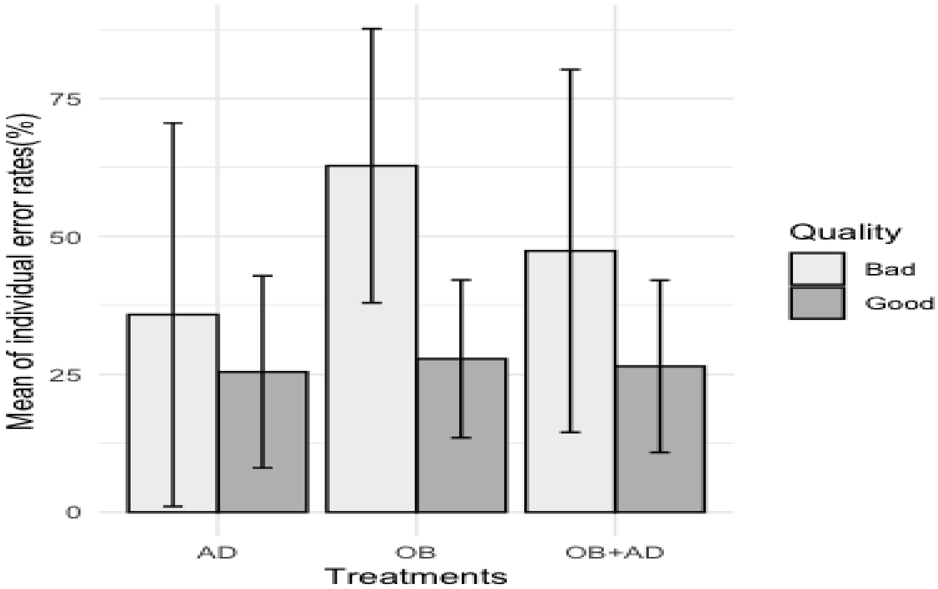

Figure 9: Mean error rates by treatments and qualities

Figure 9 depicts the mean error rate by treatments, conditional on the quality. It's obvious that the mean error rate is relatively lower when player A gives a good-quality message, especially in OB treatments(N = 6, mean = 27.82%(Good) vs 62.80%(Bad), sd = 14.30%(Good) vs 24.88%(Bad)). WSR test also confirms the results that the difference in distributions over good quality and bad quality are significant when consider both treatments(N = 14, $z = -1.665$, $p = 0.0959$) and for OB treatment separately(N = 6, $z = -1.782$, $p = 0.0747$), but insignificant in AD treatment(N = 8, $z = -0.631$, $p = 0.5281$)

Table 6 shows more detailed distributions of mean error rates for different roles. The differences in mean error rates between Player A and Player B are insignificant both for good quality and bad quality(MWW test: Good: N = 14, $z = -0.835$, $p = 0.4036$; Bad: N = 14, $z = 1.217$, $p = 0.2238$).

To further investigate the effect of quality on errors, we build a regression model(eq(3)) with a dummy variable $Quality$(0 = Bad, 1 = Good) and also interaction terms for quality

	Good(Player A)	Good(Player B)	Bad(Player A)	Bad(Player B)
OB	18.4%(9/49)	36.7 %(18/49)	78.2%(18/23)	47.8%(11/23)
AD	24.2%(16/66)	24.2%(16/66)	43.3%(13/30)	36.7%(11/30)

Table 6: Mean error rates by qualities and roles

and AD treatment and role.

$$\begin{aligned} Errors = Conflict + First - draw + Round + Task + AD + +AD\&Quality \\ + Role + Role\&Quality + Gender + Math + Knowledge \end{aligned} \tag{3}$$

Variables	All	OB	AD
AD treatment	-0.672^{**}		
	(0.341)		
Quality & ADTreatment	0.883^{**}		
	(0.409)		
Conflict	0.79^{***}	0.739^{***}	0.857^{***}
	(0.522)	(0.078)	(0.065)
First-draw	0.023	-0.087	0.179
	(0.323)	(0.522)	(0.47)
Round	-0.088^{***}	-0.072^{***}	-0.1^{**}
	(0.029)	(0.047)	(0.041)
Task	0.095	0.186	0.192
	(0.199)	(0.334)	(0.274)
Role	0.238	-0.699	-0.274
	(0.316)	(0.512)	(0.436)
Quality	-1.84^{***}	-2.846^{***}	-0.479
	(0.461)	(0.761)	(0.48)
Role & Quality	0.668^{*}	1.997^{***}	-0.133
	(0.403)	(0.708)	(0.563)
Gender	0.24	0.208	-0.023
	(0.263)	(0.646)	(0.338)
Math	-0.355^{**}	-0.311	-0.425
	(0.164)	(0.295)	(0.328)
Knowledge	-0.061	-0.536	0.447
	(0.252)	(0.376)	(0.421)
Number of observations	336	144	192

Significance level: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$
Notes: The table shows the marginal effect and standard error(in parentheses) of random -effect probit regression. The dependent variable is a dummy variable indicating second-draw error and the independent variables are shown in the first column named Variables. The second column(All) shows the regression results for all observations in both OB and AD treatments. The regression results for observations in OB and AD treatments separately are shown in the OB and AD columns.

Table 7: Peer effect: regression on second-draw errors

Table 7 presents the results of random-effect probit regression on errors. The results about the relationship between quality and errors are consistent with the above non-parametric test, while it provides more interesting information. First, *Quality* plays a highly significant role in reducing the likelihood of errors, both for all treatments($\beta_{Quality} = -1.84, p < 0.01$) and OB treatment separately($\beta_{Quality} = -2.846, p < 0.01$). These results align with the results from the MWW test. Second, there is a significant difference in errors between OB treatment and AD treatment, which is different from the above test findings($\beta_{ADtreatment} = -0.672, p < 0.05$). Also, the regression results show that quality affects the relationship between role and errors. That is, good quality has a relatively larger effect on reducing the likelihood of errors for Player A than for Player B, especially in OB treatment($\beta_{Quality} + \beta_{Role\&Quality} = -0.849, p < 0.01$).

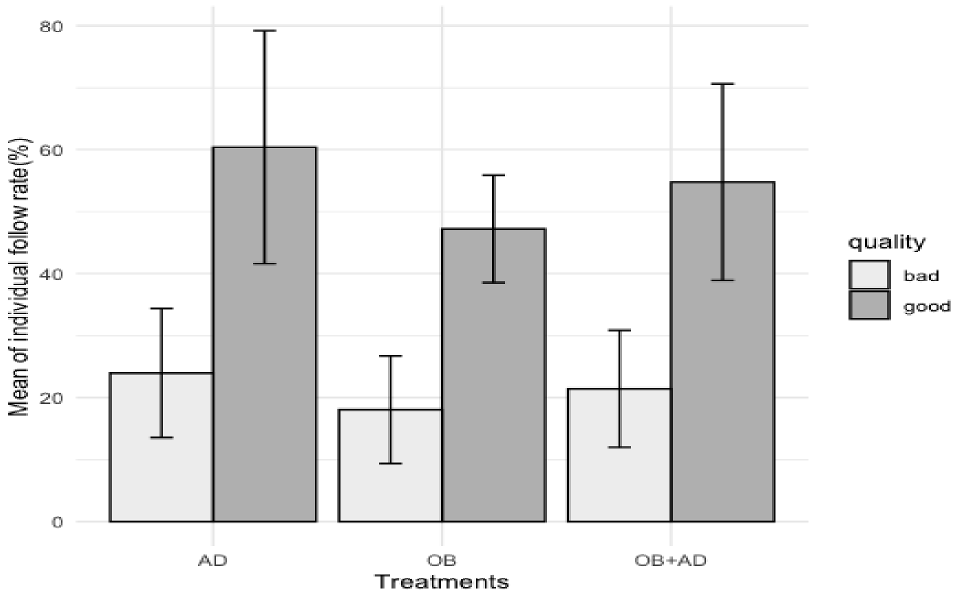

Figure 10: Mean follow rates by treatments and qualities

In addition to the relationship between quality of observation or advice with error rates, another interesting concern is whether quality affects the willingness of Player B to follow what their partners choose or suggest. To make it simple, here we define "follow" as Player B chooses the same choice they observed or Player A suggested, regardless of the first draw. Figure 10 depicts the distribution of mean follow rates conditional on the quality in different treatments. The difference in mean follow rates is obvious, where mean individual follow rates are much higher with good-quality than with bad-quality. WSR test also confirms the results that the difference in distributions between good quality and bad quality is significant when considering both treatments together(N = 7, z = 2.366, p

= 0.018) and in AD treatment(N = 4, z = 1.826, p = 0.068) individually.

$$Follow = Conflict + First - draw + Round + Study + AD + Quality \\ + Gender + Math + Knowledge \tag{4}$$

A dummy dependent variable *Follow* = 1 implies that Player B chooses the same choice they observe or Player A suggest and 0 denotes the opposite. Table 8 presents the results of random-effect probit regression on second-draw follows(shown as eq(4)). The regression results show weak evidence that good quality increases the likelihood of second-draw following for Player B($\beta_{Quality} = 0.449, p < 0.1$), especially in the *Advice* treatment($\beta_{ADtreatment} = 0.789, p < 0.05$).

Variables	All	OB	AD
AD treatment	0.789**		
	(0.389)		
Conflict	-0.357	-0.367	-0.227
	(0.286)	(0.430)	(0.446)
First-draw	-0.360	-0.187	-0.459
	(0.241)	(0.364)	(0.378)
Round	-0.299	-0.006	-0.075
	(0.033)	(0.047)	(0.062)
Task	-0.075	0.02	-0.075
	(0.225)	(0.315)	(0.367)
Quality	0.449*	0.390	0.715*
	(0.239)	(0.333)	(0.405)
Gender	0.081	0.479	0.469
	(0.272)	(0.662)	(0.491)
Math	-0.236	-0.531	0.342
	(0.232)	(0.399)	(0.549)
Knowledge	0.018	−0.536	0.447
	(0.372)	(0.376)	(0.421)
Number of observations	168	72	72

Significance level: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$
Notes: The table shows the marginal effect and standard error(in parentheses) of random-effect probit regression. The dependent variable is a dummy variable indicating second-draw follow and the independent variables are shown in the first column named Variables. The second column(All) shows the regression results for all observations in both OB and AD treatments. The regression results for observations in OB and AD treatments separately are shown in the OB and AD columns.

Table 8: Peer effect: regression on second-draw follows

5 Conclusion and Discussion

When people want to make more rational and higher-welfare choices in their decision-making, they tend to seek external information to help themselves, such as observing other people's choices or asking for their advice. In this paper, we design and trial a novel experiment to investigate how the presence of observation and advice influence people's decision-making by affecting their decision biases and make them choose more rational choices.

Given the time and funding constraints, the experiment in this paper had only 20 participants, well below the minimum number of observations required for statistical analysis. Therefore, the description of the data is not convincing enough and the conclusions of the data analysis in this paper have no statistical significance. However, this paper can be considered a pilot study prior to the formal study. Based on the performance of the experiments and data analysis, this pilot study has important implications for how to improve the formal study.

1. Increase the number of participants in each treatment.

In this pilot experiment, the number of participants in each of the three treatments is all below 10, which make the data obtained in this experiment have no statistical power and lead to inconsistent results when using different analytical methods. Therefore, in the formal experiment, a larger number of subjects is necessary in order to obtain more statistically valid conclusions from the data analysis.

2. Increase the number of rounds in the experiment.

In the analysis regarding the learning effect, the likelihood of errors decreased as the number of rounds increased. The distribution of errors also shows differences in the first 4 rounds and the last 4 rounds of the experiment. However, when the total 12 rounds of the experiment were divided into the first and second halves, the two parts did not show a significant difference in the error rates. The inconsistency in the data analysis of the learning effect may be due to the insufficient number of repetitions. Some participants also responded that the experiment was difficult to understand and was not given enough rounds to fully understand the task set. Therefore, in the formal experiment, the number of rounds should be increased to 40 or 60 per treatment, and last around 35-45 minutes, in order to give the subjects enough time to understand the task and to better study the learning effect.

With regard to further research, in our experiment, two possible incentives for participants being observed or giving advice are the payment incentive and the reputation effect. However, we didn't separate and focus on these two incentives individually in the current research. For further study, some interesting findings may come from the study on which incentive plays a more important role in influencing the observed to make choice or advice-giver to give advice. Whether the reputation of the participants affects their partners' willingness of following the observation or advice and how reputation change in repeated rounds throughout the whole experiment are also of interest and need more exploration.

References

[1]AAchtziger, A., & Alós-Ferrer, C. (2014). Fast or Rational? A Response-Times Study of Bayesian Updating. Management Science, 60(4), 923–938.https://doi.org/10.1287/mnsc.2013.1793

[2]Achtziger, A., Alós-Ferrer, C., Hügelschäfer, S., & Steinhauser, M. (2015). Higher incentives can impair performance: Neural evidence on reinforcement and rationality. Social Cognitive and Affective Neuroscience, 10(11), 1477–1483. https://doi.org/10.1093/scan/nsv036

[3]Akaishi, R., Umeda, K., Nagase, A., & Sakai, K. (2014). Autonomous Mechanism of Internal Choice Estimate Underlies Decision Inertia. Neuron, 81(1), 195–206. https://doi.org/10.1016/j.neuron.2013.10.018

[4]Alós-Ferrer, C., Buckenmaier, J., Garagnani, M., & Ritschel, A. (2022). The reinforcement paradox: Monetary incentives and Bayesian updating. Economics Letters, 211, 110200. https://doi.org/10.1016/j.econlet.2021.110200

[5]Alós-Ferrer, C., Hügelschäfer, S., & Li, J. (2016). Inertia and Decision Making. Frontiers in Psychology, 7. https://doi.org/10.3389/fpsyg.2016.00169

[6]Alós-Ferrer, C., & Ritschel, A. (2018). The reinforcement heuristic in normal form games. Journal of Economic Behavior & Organization, 152, 224–234. https://doi.org/10.1016/j.jebo.2018.06.014

[7]Anderson, L.R., Holt, C.A.(1997). Information cascades in the laboratory. Am. Econ. Rev. 847–862

[8]Banerjee, A. V. (1992). A Simple Model of Herd Behavior. The Quarterly Journal of Economics, 107(3), 797–817. https://doi.org/10.2307/2118364

[9]Barto, A. G., & Sutton, R. S. (1997). Reinforcement Learning in Artificial Intelligence. In Advances in Psychology (Vol. 121, pp. 358–386). Elsevier. https://doi.org/10.1016/S0166-4115(97)80105-7

[10]Börgers, T., & Sarin, R. (1997). Learning Through Reinforcement and Replicator Dynamics. Journal of Economic Theory, 77(1), 1–14. https://doi.org/10.1006/jeth.1997.2319

[11]Bourdieu, P. (1985). The social space and the genesis of groups. Social Science Information, 24(2), 195–220. https://doi.org/10.1177/053901885024002001

[12]Çelen, B., & Kariv, S. (2004). Distinguishing informational cascades from herd behavior in the Laboratory. American Economic Review, 94(3), 484–498. https://doi.org/10.125

7/0002828041464461
[13]Çelen, B., Kariv, S., & Schotter, A. (2010). An Experimental Test of Advice and Social Learning. Management Science, 56(10), 1687–1701. https://doi.org/10.1287/mnsc.1100.1228
[14]Çelen, B., & Kariv, S. (2005). An experimental test of observational learning under imperfect information. Economic Theory, 26(3), 677–699. https://doi.org/10.1007/s00199-004-0542-0
[15]Charness, G., Karni, E., & Levin, D. (2007). Individual and group decision making under risk: An experimental study of Bayesian updating and violations of first-order stochastic dominance. Journal of Risk and Uncertainty, 35(2), 129–148.
[16]Charness, G., & Levin, D. (2005). When Optimal Choices Feel Wrong: A Laboratory Study of Bayesian Updating, Complexity, and Affect. 11.
[17]Chaudhuri, A., & Graziano, S. (2006). Social Learning and Norms in a Public Goods Experiment with Inter-Generational Advice. REVIEW OF ECONOMIC STUDIES, 24.
[18]Chen, Y., Wang, Q., & Xie, J. (2011). Online Social Interactions: A Natural Experiment on Word of Mouth versus Observational Learning. JOUrNAL Of MArkETING RESEArCH, 17.
[19]Chevalier, J. A., & Mayzlin, D. (2006). The Effect of Word of Mouth on Sales: Online Book Reviews. Journal of Marketing Research, 43(3), 345–354.
[20]Chiang, Y.-M., Hirshleifer, D., Qian, Y., & Sherman, A. E. (2011). Do Investors Learn from Experience? Evidence from Frequent IPO Investors. Review of Financial Studies, 24(5), 1560–1589. https://doi.org/10.1093/rfs/hhq151
[21]Choi, J. J., Laibson, D., Madrian, B. C., & Metrick, A. (2009). Reinforcement Learning and Savings Behavior. The Journal of Finance, 64(6), 2515–2534. https://doi.org/10.1111/j.1540-6261.2009.01509.x
[22]Duflo, E., & Saez, E. (2002). Participation and investment decisions in a retirement plan: The influence of colleagues' choices. Journal of Public Economics, 28.
[23]Duflo, E., & Saez, E. (2003). The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence from a Randomized Experiment. The Quarterly Journal of Economics, 118(3), 815–842.
[24]Erev, I., & Haruvy, E. (2016). 10. Learning and the Economics of Small Decisions. In J. H. Kagel & A. E. Roth (Eds.), The Handbook of Experimental Economics, Volume

Two. Princeton University Press. https://doi.org/10.1515/9781400883172-011
[25]Erev, I., & Roth, A. E. (1998). Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. American Economic Review, 88(4), 848-881.
[26]Feri, F., Meléndez-Jiménez, M. A., Ponti, G., & Vega-Redondo, F. (2011). Error cascades in observational learning: An experiment on the Chinos game. Games and Economic Behavior, 73(1), 136–146. https://doi.org/10.1016/j.geb.2010.10.003
[27]Gal, D. (2006). A psychological law of inertia and the illusion of loss aversion [Data set]. American Psychological Association. https://doi.org/10.1037/e683162011-083
[28]Gale, D. (1996). What have we learned from social learning? European Economic Review, 40(3–5), 617–628. https://doi.org/10.1016/0014-2921(95)00074-7
[29]Gilchrist, D. S., & Sands, E. G. (2016). Something to Talk About: Social Spillovers in Movie Consumption. Journal of Political Economy, 124(5), 1339–1382. https://doi.org/10.1086/688177
[30]Handel, B. R. (2013). Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts. The American Economic Review, 103(7), 2643–2682.
[31]Harvey, N., Fischer, I. (1997). Taking Advice: Accepting Help, Improving Judgment, and Sharing Responsibility. Organizational Behavior and Human Decision Processes, 70(2), 117–133. https://doi.org/10.1006/obhd.1997.2697
[32]Hodgkinson, G. P. (1997). Cognitive Inertia in a Turbulent Market: The Case of UK Residential Estate Agents. Journal of Management Studies, 34(6), 921–945. https://doi.org/10.1111/1467-6486.00078
[33]Holroyd, C. B., & Coles, M. G. H. (2002). The neural basis of human error processing: Reinforcement learning, dopamine, and the error-related negativity. Psychological Review, 109(4), 679–709. https://doi.org/10.1037/0033-295X.109.4.679
[34]Hong, H., Kubik, J. D., & Stein, J. C. (2004). Social Interaction and Stock-Market Participation. The Journal of Finance, 59(1), 137–163.
[35]Iyengar, R., & Schotter, A. (2008). Learning under supervision: An experimental study. Experimental Economics, 11(2), 154–173. https://doi.org/10.1007/s10683-007-9164-2
[36]Kaustia, M., & Knüpfer, S. (2008). Do Investors Overweight Personal Experience? Evidence from IPO Subscriptions. The Journal of Finance, 63(6), 2679–2702.

[37]Li, J., & Zhang, X. (2022). Do people provide good advice? Journal of Economic Behavior & Organization, 199, 44–64. https://doi.org/10.1016/j.jebo.2022.05.009
[38]Maclin, R., & Shavlik, J. W. (1996). Creating advice-taking reinforcement learners. Machine Learning, 22(1–3), 251–281. https://doi.org/10.1007/BF00114730
[39]Madrian, B. C., & Shea, D. F. (2001). The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior. The Quarterly Journal of Economics, 116(4), 1149–1187
[40]Merlo, A., & Schotter, A. (2003). Learning by not doing: An experimental investigation of observational learning. Games and Economic Behavior, 21.
[41]Moretti, E. (2011). Social learning and peer effects in consumption: Evidence from movie sales. The Review of Economic Studies, 78(1), 356–393. https://doi.org/10.1093/restud/rdq014
[42]Schram, A., & Sonnemans, J. (1996). Why people vote: Experimental evidence. Journal of Economic Psychology, 17(4), 417–442. https://doi.org/10.1016/0167-4870(96)00022-0
[43]Schotter, A. (2003). Decision Making with Naive Advice. The American Economic Review, 93(2,), 196–201.
[44]Schotter, A., & Sopher, B. (2003). Social Learning and coordination conventions in Intergenerational games: An experimental study. Journal of Political Economy, 111(3), 498–529. https://doi.org/10.1086/374187
[45]Schotter, A., & Sopher, B. (2007). Advice and behavior in intergenerational ultimatum games: An experimental approach. Games and Economic Behavior, 58(2), 365–393. https://doi.org/10.1016/j.geb.2006.03.005
[46]Schonberg, T., Daw, N. D., Joel, D., & O'Doherty, J. P. (2007). Reinforcement learning signals in the human striatum distinguish learners from Nonlearners during reward-based decision making. Journal of Neuroscience, 27(47), 12860–12867. https://doi.org/10.1523/jneurosci.2496-07.2007
[47]Schultz, W. (1998). Predictive reward signal of dopamine neurons. Journal of Neurophysiology, 80(1), 1–27. https://doi.org/10.1152/jn.1998.80.1.1
[48]Steinel, W., Abele, A. E., & De Dreu, C. K. W. (2007). Effects of Experience and Advice on Process and Performance in Negotiations. Group Processes & Intergroup Relations, 10(4), 533–550. https://doi.org/10.1177/1368430207081541
[49]Sull, D. N. (2005). Why good companies go bad. 8. Thorndike, E. L. (1911). Animal

Intelligence; experimental studies, by Edward L. Thorndike. https://doi.org/10.5962/bhl.title.1201

[50]Tripsas, M., & Gavetti, G. (2000). Capabilities, Cognition, and Inertia: Evidence from Digital Imaging. Strategic Management Journal, 21(10/11,), 1147–1161.

[51]Tykocinski, O. E., Pittman, T. S., & Tuttle, E. E. (1995). Inaction inertia: Foregoing future benefits as a result of an initial failure to act. Journal of Personality and Social Psychology, 68(5), 793–803. https://doi.org/10.1037/0022-3514.68.5.793//

[52]Tykocinski, O. E., & Pittman, T. S. (1998). The consequences of doing nothing: Inaction inertia as avoidance of anticipated counterfactual regret. Journal of Personality and Social Psychology, 75(3), 607–616. https://doi.org/10.1037/0022-3514.75.3.607

[53]Weick, K. E., & Quinn, R. E. (n.d.). ORGANIZATIONAL CHANGE AND DEVELOPMENT. 27.

[54]Yaniv, I. (2004). Receiving other people's advice: Influence and benefit. Organizational Behavior and Human Decision Processes, 93(1), 1-13. doi:10.1016/j.obhdp.2003.08.002