House of Finance

SAFE

Working Paper Series

Wenhui Li – Christian Wilde

Belief Formation and Belief Updating under Ambiguity: Evidence from Experiments

SAFE Working Paper No. 251

SAFE | Sustainable Architecture for Finance in Europe
A cooperation of the Center for Financial Studies and Goethe University Frankfurt

House of Finance | Goethe University
Theodor-W.-Adorno-Platz 3 | 60323 Frankfurt am Main

Tel. +49 69 798 30080 | Fax +49 69 798 33910
info@safe-frankfurt.de | www.safe-frankfurt.de

Belief formation and belief updating under ambiguity: evidence from experiments*

Wenhui Li and Christian Wilde†

Preliminary version

June 2, 2019

Abstract

Decisions under ambiguity depend on both the belief regarding possible scenarios and the attitude towards ambiguity. This paper exclusively investigates the belief formation and belief updating process under ambiguity, using laboratory experiments. The results show that half of the subjects tend to adopt a simple heuristic strategy when updating beliefs, while the other half seems to partially adopt the Bayesian updates. We recover beliefs, represented by distributions of the priors/posteriors. The recoverable initial priors mostly follow a uniform distribution. We also find that subjects on average demonstrate slight pessimism in an ambiguous environment.

Keywords: ambiguity, learning strategy, belief updates, non-Bayesian updates, pessimism, laboratory experiments

1 Introduction

One central question in the decision theory is how the presence of ambiguity affects individuals' decision making. Ambiguity may affect decision makings through two aspects: how is the belief about the possible scenarios in the ambiguous environment formed and updated, and how ambiguity-averse (-neutral, -seeking) is the decision maker? These two aspects work very differently in shaping the final decisions. A sizable number of theoretical literature studies ambiguity by separating these two aspects (Ahn 2008; Brennan 1998; Cao et al. 2005; Chen and Epstein 2002; Epstein and Schneider 2007; Galaabaatar and Karni 2013; Karni 2018, to name a few). However, empirically it is difficult to distinguish beliefs from attitudes, and especially to elicit beliefs. Most of the empirical literature fails to do so. Our paper contributes to the literature in several ways. First, this paper isolates beliefs from attitudes in an ambiguous

*We thank Peter Ockenfels for valuable discussions. We also thank seminar participants at the FLEX 10-year anniversary conference in 2019 for helpful comments. We gratefully acknowledge research support from the Research Center SAFE, funded by the State of Hessen initiative for research LOEWE.

†Correspondence: Center of Excellence SAFE, Goethe University Frankfurt, Grüneburgplatz 1, 60323 Frankfurt am Main, Germany. Emails: w.li@safe.uni-frankfurt.de; wilde@finance.uni-frankfurt.de.

environment by adopting a special laboratory experiment design. Second, we recover beliefs, represented by the entire distributions of initial priors, in situations involving ambiguity. Third, we investigate how subjects update their beliefs in response to new information. We consider five possible models of belief updating dynamics, and determine which models are applied by the individual subjects for belief updating. Fourth, we investigate whether beliefs are biased towards good or bad outcomes, i.e. whether there exists pessimism or optimism among subjects.

The concept of ambiguity is usually defined in comparison with the concept of risk. Risk is defined as a situation in which a decision maker faces an event, whose outcome contingencies have clear and objective probability measurements, while ambiguity is mostly defined as a situation in which a decision maker cannot obtain the full information of the probability measurement, due to the scarcity or imprecision of the information. (Becker et al. 1964; Epstein 1999; Knight 1921, to name a few). The belief under ambiguity can be defined as the subjective perception of the unknown probability measurement. Previous literature argues that the probability measurement under ambiguity can be degenerated to a certain subjective probability evaluation (Gilboa et al. 2008a; Savage 1954), or expressed by a non-additive probability system theorized by the Choquet integral (Choquet 1954; Schmeidler 1989), or understood as a two-order probability measurement (Ghirardato et al. 2004; Gilboa and Schmeidler 1989; Klibanoff et al. 2005). This paper adopts the theory which interprets the ambiguous environment as two-layer uncertainty. The ambiguous environment, in our experiment, is operationalized by an ambiguous lottery, whose winning probability is unknown to the players (for the detailed design, see the experiment design section). The ambiguous lottery can be translated into a package of multiple single lotteries. Each single lottery in the package has a singular and known winning probability, defined as the first-order probability. But the occurrence probability of each single lottery is unknown. It is how ambiguity arises. This occurrence probability is defined as the second-order probability. In this paper, we investigate beliefs by tracking down the subjective evaluations of the second-order probability at a given point in time. We design the ambiguous lottery to be completely ambiguous. From the subjects' point of view, it means the winning probability of the lottery lies between $[0, 1]$, incremented by 0.01. Thus, our experiment generates a large number of possible scenarios including extreme scenarios, thereby offering a rich setting with a high degree of ambiguity. It supplements the existing literature with binary set-ups (Buser et al. 2018; Dominitz and Hung 2009; Filippis et al. 2017) and rather narrow interval set-ups (Chew et al. 2017).

One objective of this paper is to investigate the initial subjective evaluation of the second-order probability at the very beginning of the experiment, known as the initial prior. Our interest is to recover the initial prior distribution, deduce its functional form, and capture several main characteristics of the distribution, for instance, the shape of the distribution, mode, and variance. We first characterize the initial prior distribution by the *beta* distribution. The *beta* distribution is widely applied by conjugate prior literature to model the initial prior (Diaconis and Ylvisaker 1979; Gelman et al. 2004; Schlaifer and Raiffa 1961). The *beta* distribution is a family of various distributions, including but not limited to uniform distribution, bell-shaped distribution and distributions with monotonic PDF. Its shape is governed by only two shape parameters. The *beta* distribution ranges between zero and one, naturally fitting the winning probability in our experiment design that ranges from zero to one. Due to its attractive features, diversity of distribution, simplicity of construction and range, we first assume that one's initial prior follows the *beta* distribution and subsequently recover the shape of the distribution by estimating the shape parameters. We develop a method to recover the characteristics of the initial prior of each subject, also acknowledging the possibility that subjects may think simple and do not have the specific prior distribution in mind.

Another objective of this paper is to understand the subjective belief updating process, namely, how does a subject update her belief responding to the new information provided to her. This process is also called *learning*, and we use these two terms, belief updating and learning, interchangeably in the paper. More specifically, learning can be understood as a process that subjects update initial priors to posteriors using the new information. The experiment permits learning in the way that: a subject always faces the same and the only ambiguous lottery through the entire experiment. The lottery is repeatedly played in front of her and reset to the original status after each play. The result of each play, winning or losing, is observable and traceable to her (for the detailed design, see the experiment design section). Therefore, she can update her belief about the winning probability of the lottery, referring to the provided new information in whatever way she wishes. Previous literature concerning learning strategies can be roughly categorized into two streams: Bayesian updates (Branger et al. 2013; Gilboa and Schmeidler 1993; Hanany and Klibanoff 2007; Peijnenburg 2014; Pires 2002), where subjects learn the new information and update their beliefs employing the Bayes rule, and non-Bayesian updates (Epstein 2006; Epstein et al. 2008; Marinacci 2002). This paper considers both possibilities. We propose five learning strategies (henceforth: LS1, LS2 $\cdots$ LS5) to accommodate both Bayesian and non-Bayesian learning. LS1 and LS3 are categorized as Bayesian updating rules, where LS1 assumes a uniformly distributed initial prior while LS3 assumes a *beta*-distributed initial prior. LS2 and LS4 are categorized as non-Bayesian updating rules. They both update beliefs by weighting between the previous beliefs and the Bayesian reference. What distinguishes them is that LS2 refers to the Bayesian updates with a uniformly distributed initial prior while LS3 refers to the Bayesian updates with a *beta*-distributed initial prior. This paper intends to explicitly recover the shape of the belief distributions after each update. Most of the existing empirical studies fail to recover the entire belief distribution, rather they are limited to the discussion of belief range (Campanale 2011; Gilboa et al. 2008b), belief as single value (Buser et al. 2018; Dominitz and Hung 2009; Filippis et al. 2017), or different specifications (Campanale 2011; Epstein and Schneider 2007). Our experiment design and the constructions of LS1-LS4 help us to fill this gap. LS5 is rather unique. It is categorized as non-Bayesian updating rule and does not relate to the Bayes rule in any visible way. LS5 leaves the shape of the initial prior distribution unspecified and is a model for a heuristic learning strategy. For each given subject and LS, we use experimental data of the belief updates to estimate the parameters which govern the LS features, and later recover the belief update dynamics. The proposed five LSs incorporate abundant choice possibilities along the learning process. Such wide coverage seems to be diverse enough, since the factual belief updates of a given subject, in most of the cases, are well captured by at least one of the five LSs. Therefore we can determine, for each given subject, which learning strategy she adopts. Based on these results, we can make statements on their initial prior distributions.

The third objective of this paper is to analyze the tendency that subjects display towards pessimism/optimism during the belief formation and the belief updating process. The concept of pessimism/optimism is widely applied in different disciplines. In the psychology literature, pessimism is usually defined as a negative mental and emotional disposition in which an undesirable outcome is anticipated. In philosophy works, it is defined as the worldview which perceives life has no intrinsic meaning or value. In decision theory, however, the concept of pessimism is not well-defined nor thoroughly discussed. Especially in the ambiguity literature, the concept of pessimism/optimism is usually mingled with other concepts such as ambiguity attitude. In some cases, authors use the word pessimism/optimism interchangeably with the word ambiguity aversion/seeking (Giraud and Thomas 2017). We argue that pessimism/optimism can be conceptualized as the perceptive tendency to bias toward good/bad scenarios in assigning probabilities to scenarios. We clarify that pessimism/optimism is, in essence, a characteristic of belief, rather than an attitude. It bases on the shape of belief, prior to and independent from the decision making based on this belief. Attitude only enters and takes effect during the phase of decision making. Hence pessimism/optimism belongs to the discussion of belief characterization and should be clearly differentiated from the discussion of attitude. This paper manages to do so. The experiment design facilitates a clear cut between belief and attitude, and thus pessimism/optimism is partitioned out as a clean independent measurement. In addition, the term pessimism/optimism is embedded with judgmental features. Pessimism reflects the perceptive tendency towards bad situations, whereas optimism reflects the perceptive tendency towards good situations. The experiment design also accordingly operationalizes this structure: the higher the winning probability of the lottery is, the better the situation is considered to be. The degree of pessimism can be measured by the negative bias towards the worst possible situation of the factual belief, deviating from some chosen benchmark. The benchmark is a theoretically constructed belief profile that is neutral between good and bad outcomes. For sure, the choice of the neutral benchmark is an arguable question. This paper proposes six benchmark candidates, covering a range of possible initial prior distributions. Under the experiment set-up, pessimism is measured by the difference between the subjective conceived winning probability of the ambiguous lottery and the winning probability delivered by the belief formation/updating strategy according to a neutral benchmark.

This paper reaches the following findings: (1) Subjects tend to prefer simple initial prior formation: more than one-third of the subjects are inclined to form a uniformly distributed initial prior, and nearly half of the subjects keep silence about the initial prior distribution and constrain their attention to the mode of the prior distribution. (2) As for the learning strategy, nearly half of the subjects employ the Bayes rule to update their beliefs, but all of them choose to, to some degree, deviate from the perfect Bayesian updates path. The other half tends to adopt a heuristic strategy, characterized by a simple update pattern. (3) Subjects on average demonstrate slight pessimism in situations with higher degrees of ambiguity.

The rest of the paper is organized as follows: Section 2 introduces the experiment design. Section 3 presents the descriptive analysis of the belief updates data. Section 4 introduces the set-up of the five LSs. Section 5 discusses the characteristics of LSs. Section 6 discusses the subjective choices of the LS. Section 7 discuses the pessimism/optimism of subjects. Section 8 presents robustness checks. Section 9 is the conclusion.

2 Experiment design

The laboratory experiment is computerized by Z-tree (Fischbacher 2007). We operationalize the ambiguous environment by the design of a completely ambiguous lottery urn. Subjects of the experiment are told that the urn contains in total 100 balls. The color of a ball is either white or black. Neither the true proportion of the white ball(s) nor the true proportion of the black ball(s) is known to any subject. It is also explained to the subjects that the number of the white ball(s) in the urn can be any integer between zero and 100 (both ends inclusive). A so-called *guess game* is designed to track down the subjective belief. In a *guess game*, a subject needs to answer the following question: standing at this point, how many white balls do you think are in the urn? A subject enters her own belief about the proportion of the white ball(s) in the guess game. Figure 1a presents the screen display of the first *guess game*.

In order to track down the learning and belief updating process, new information about the urn is provided to the subjects. New information is generated by implementing draws from the ambiguous urn. In each draw, one ball is drawn out from the urn and its color, either white or black, is displayed to the relevant subjects. Then the ball is immediately put back to the urn. Therefore it is the draw with replacement. *Guess games* and draws are played/implemented for multiple times and the sequence is designed as follows: we first define 15 periods, indexed by $t = 1, 2 \cdots 15$. In each period t, subjects first play a *guess game*, followed by a draw implementation. Therefore the $1 + 1$ pack (one *guess game* plus one draw) is repeated for 15 times. Table 1 displays the experiment procedure.

The belief reported at time t can be seen as an updated belief based on the learning of $t - 1$ times of draw history. The past draw history, if any, is displayed on screen for subjects' reference. Figure 1b presents the *guess game* screen display after some draws are implemented (as an example, $t = 6$, five draws are done and the draw results are displayed on the screen.). The true proportion of the white balls in the ambiguous urn is fixed at 40 for good. Subjects are incentivized such that: every time one enters the correct white ball proportion (=40) in the *guess game*, she is rewarded with two Euro, otherwise zero. In other words, subjects are incentivized to insert the mode value of the prior/posterior distribution (if any) in each *guess game*, since the mode corresponds with the highest second-order probability, which indicates that the urn is most likely to contain as many white balls as the mode value. By inserting the mode value, one maximizes her chance to win the reward. The earning is only announced to the subjects at the very end of the experiment, hence the ambiguous feature of the urn sustains through the entire experiment.

It is worthwhile to mention that the *guess game* design guarantees that the data obtained from the games are purely related to the belief about the ambiguous environment, independent from the attitude towards ambiguity. Attitude plays a role in decision and becomes observable only when another alternative exists for comparison, for instance, another ambiguous or risky lottery, or a riskless alternative. Since the *guess game* does not imply any preference-related decisions between two alternatives, and it merely requires the subjects report their beliefs about the sole designed ambiguous environment, such design ensures the reported data in the *guess games* are only belief relevant.

Three sessions of this experiment have been conducted. Subjects are all randomly selected from the subject pool of the Frankfurt Laboratory for Experimental Economic Research (FLEX), Goethe University Frankfurt. Most of the selected subjects are students from Goethe University. 13 subjects attended Session I on 27.03.2018, 14 subjects attended Session II on 16.05.2018 and 14 subjects attended Session III on 23.05.2018. In Session II and III, one extra *guess game*, denoted as the *guess game* at $t = 16$, is played after the 15^{th} draw is implemented. All subjects are grouped into markets. Each market implements its draws independently (the ambiguous urns used in all markets are identical though), hence the draw history is market-specific. A subject only knows the draw information of her own market. There are in total five markets: one in Session I, two in Session II, and two in Session III. Thus five independent draw history paths are generated. Table 2 displays this design information.

Apart from the *guess game*, the experiment also includes other parts, such as the choice games and the asset trading sections. Since they are not of interest of this paper, we suppress the detailed introduction of these parts, except for one point worthwhile to be mentioned: In the choice games, the ambiguous urn is indeed served as a lottery involving high and low payouts. In the asset trading section, the draws from the urn are also used to determine the asset dividend: each time a white ball is drawn, the asset pays out a positive dividend; each time a black ball is drawn, the asset pays out zero dividend. This pattern sustains through the entire experiment. Therefore white draws are always interpreted as *good news*, while black draws correspond to *bad news*. This information is useful when we analyze the pessimism/optimism of the subjects. In addition, some quiz and a demonstration of a physical ambiguous lottery urn (in fact, a big box containing 100 cards with “white” or “black” written on) are implemented, so as to help the subjects fully understand the concept of lottery and the *guess game*. The content of the big box is not observable to any subjects, and no draws are implemented during the demonstration. Therefore when the first *guess game* is played, the lottery urn is as completely ambiguous as intended. The complete experiment design is also reported in Table 1. A complete experiment session lasts 2 hours 15 minutes on average. The average earning per subject is 27.8 Euro (total earning from all parts).

3 Descriptive analysis

The guess game entry of subject i at the beginning of period t is denoted by $white_{it}$. Each subject has 15 (in Session I) or 16 (in Session II and III) entries of $white_{it}$ at the end of the experiment. We compute the by-period market mean value: $white_{mt} = \frac{1}{N_m}\sum_i white_{it}$, where N_m denotes the number of subjects in market m. The results are reported in Table 5 and illustrated in Figure 2. In Figure 2, the shaded area at t indicates that a white draw is observed at the end of period $t-1$. Analogously, the non-shaded area at t indicates that a black draw is observed at the end of period $t-1$. Learning can be understood as the belief adjustment responding to the draw information. As can been seen, the $white_{mt}$ value mostly fluctuates around the 50-white-ball line in all markets except Market 3. Market 1, 2 and 5 experience seven-time white draw and Market 4 experiences six-time out of 15 times. In contrast, Market 3 only experiences three times of white draws out of 15. This may be the reason why beliefs are on average lower in Market 3 than in other markets.

The across-subject standard deviation reported in Table 5 sd.columns represent the heterogeneity regarding the responses in the *guess game* in a given period. The standard deviation tends to decrease when time evolves in all markets except in Market 3. It supports the argument that learning may help mitigate the belief heterogeneity. However, since the frequency of learning is finite, only taking place for 15 times, the belief heterogeneity always sustains.

In theory, a rational belief update should steer such that each time a white draw is observed, current belief is upwards adjusted or maintains at this current level; Each time a black draw is observed, current belief is downwards adjusted or maintains at the current level. It can be seen as weakly rational. Following this definition, we compute the rate of rational updates for each subject and report the results in Table 6. As can be seen that subjects mostly update their beliefs rationally. The market average rates of rational updates are all above 0.7, and full sample average is nearly 0.8. There also exists some cases where subjects act more irrationally, for example, Subject 1, 3, 15, 25, 39, whose rational updates rate is equal or below 0.5. We discuss how irrational updates are accommodated by different learning strategies in the later section.

4 Learning strategy

The main goal of this paper is to analyze the learning strategy one may adopt to form and update her belief in an ambiguous environment. The guess game records the dynamic mechanism how a subject forms her initial belief about the proportion of the white balls in the given ambiguous urn, and how she updates her belief through learning from the dividend draws. We first model five learning strategies that a subject may adopt for belief updating. The five learning strategies include both Bayesian and non-Bayesian updates, with various initial prior distributions.

4.1 LS1: Bayesian updates with uniformly distributed initial prior

First we propose a strategy which starts with a uniformly distributed initial prior and employs the Bayes rule to update the priors into posteriors. The lottery urn is designed to be completely ambiguous. Hence there are 101 possible scenarios regarding the true composition of the urn. These 101 scenarios can be indexed by the possible proportion of the white ball(s) in the urn, denoted by W, where $W = 0, 1, 2 \cdots\ 100$. The belief of a subject at any time point can be expressed in the way how she assigns probabilities to each of the 101 scenarios. Assume there is a representative subject who adopts the following belief updating strategy: At the beginning of Period 1 when no draws occur, she assigns equal probability to each of the 101 scenarios, forming a uniformly distributed initial prior:

$$Prior(W) = \frac{1}{101} \approx 0.99\%; \quad W \in \{0, 1, 2 \cdots 100\} \tag{1}$$

Draws are implemented within market. At the beginning of Period t, Market m sees $n \equiv t-1$ balls are drawn out with replacement. Suppose $k_{n,m}$ out of n are white draws, the representative subject of market m constructs her posteriors employing Bayes rule:

$$Posterior(W|n, k_{n,m}) = \frac{Prob(n, k_{n,m}|W) \times Prior(W)}{\sum_{j=0}^{100} Prob(n, k_{n,m}|j) \times Prior(j)} \tag{2}$$

$$W, j \in \{0, 1, 2 \cdots 100\}; \quad 0 \leq k_{n,m} \leq n \equiv t-1 \leq 15, \quad k_{n,m}, n \subseteq \mathbb{Z}; \quad m \in \{1, 2 \cdots 5\}$$

where $\mathbb{Z}$ denotes the set of all integers. $Prob(n, k_{n,m}|W)$ denotes the probability of observing $k_{n,m}$ units of white balls in n draws with replacement from an urn containing W units of white balls and $100 - W$ units of black balls. And $Prior(j) = 1/101$ for each j. We can write:

$$Prob(n, k_{n,m}|W) = \binom{n}{k_{n,m}} (W/100)^{k_{n,m}} (1 - W/100)^{n-k_{n,m}} \tag{3}$$

$$\binom{n}{k_{n,m}} = \frac{n!}{(n - k_{n,m})!} \tag{4}$$

Plugging into equation (2) yields:

$$Posterior(W|n, k_{n,m}) = \frac{(W/100)^{k_{n,m}} (1 - W/100)^{n-k_{n,m}}}{\sum_{j=0}^{100} (j/100)^{k_{n,m}} (1 - j/100)^{n-k_{n,m}}} \tag{5}$$

$$W, j \in \{0, 1, 2 \cdots 100\}; \quad 0 \leq k_{n,m} \leq n \equiv t-1 \leq 15, \quad k_{n,m}, n \subseteq \mathbb{Z}; \quad m \in \{1, 2 \cdots 5\}$$

Therefore at the beginning of period t, the representative subject updates the probability of scenario W from $Prior(W)$ to $Posterior(W|n, k_{n,m})$. The guess game rewards a subject each time when she correctly guesses the true proportion of the white balls. In other words, from Period 2 on, subjects are incentivized to insert the number which equalizes with the scenario index W, to which the highest probability is attached at this time. Therefore LS1, denoted by BU_{mt} (Bayesian Updates), reads:

$$BU_{mt} = \underset{W}{\text{argmax}} Posterior(W|n, k_{n,m}) \tag{6}$$

The priors and posteriors at each time t can be seen as the probability mass function (PMF) for the 101 scenarios, with the mode being BU_{mt}. Since the draw history is market-specific, BU_{mt} is also market-specific.

LS1': Maximum likelihood updates. Another way to construct a learning strategy for the representative subject is to apply the Maximum Likelihood (ML) approach. Suppose at the beginning of Period t, market m sees $n \equiv t-1$ balls are drawn out with replacement and $k_{n,m}$ of them are white balls, the ML update reads:

$$ML_{mt} = \frac{100k_{n,m}}{n} \tag{7}$$

$$0 \leq k_{n,m} \leq n \equiv t-1 \leq 15, \quad k_{n,m}, n \subseteq \mathbb{Z}; \quad m \in \{1, 2 \cdots 5\}$$

Now we show that LS1 and LS1' are equivalent.

Proposition 1. $BU_{mt} = ML_{mt}$

Proof.

$$\begin{aligned} BU_{mt} &= \underset{W}{\text{argmax}} Posterior(W|n, k_{n,m}) \\ &= \underset{W}{\text{argmax}}[(W/100)^{k_{n,m}}(1-W/100)^{n-k_{n,m}}] \end{aligned} \tag{8}$$

First consider an internal solution such that $W^* \neq 0$ and $W^* \neq 100$. The maximization problem can be rewritten as:

$$\max_{W}[k_{n,m}\ln(W/100) + (n-k_{n,m})\ln(1-W/100)] \tag{9}$$

FOC leads to

$$\frac{k_{n,m}}{W^*} = \frac{(n-k_{n,m})}{100-W^*}; \quad \text{if } W^* \neq 0, \text{ and } W^* \neq 100 \tag{10}$$

Hence

$$BU_{mt} = W^* = \frac{100k_{n,m}}{n} = ML_{mt}; \quad \text{if } W^* \neq 0, \text{ and } W^* \neq 100 \tag{11}$$

Now consider if $W = 0$ is the solution to (8). If $k_{n,m} \neq 0$, the objective function is always positive with any $W \in (0, 100)$, therefore $W = 0$ cannot be the solution. However if $k_{n,m} = 0$, the objective function reduces to $(1-W/100)^n$. Given the domains of the variables, $W^* = 0$ is the solution. In other words, only if $k_{n,m} = 0$, $W^* = 0$ is the solution to (8). Hence

$$BU_{mt} = W^* = 0 = ML_{mt}; \quad \text{if } k_{n,m} = 0. \tag{12}$$

Analogously, $W^* = 100$ is the solution to (8) only if $k_{n,m} = n$, therefore

$$BU_{mt} = W^* = 100 = ML_{mt}; \quad \text{if } k_{n,m} = n. \tag{13}$$

Combining equation (11)-(13) yields:

$$BU_{mt} = ML_{mt} \tag{14}$$

$$m \in \{1, 2 \cdots 5\}; \quad t \in \{1, 2 \cdots 16\}$$

□

Therefore LS1 (Bayesian updates with uniformly distributed initial prior) and LS1' (Maximum likelihood updates) are equivalent. From now on we do not distinguish between LS1 and LS1', and simply use LS1 to refer to the Bayesian updates with uniformly distributed initial prior. Bernstein von Mises theorem (Doob 1949) states that Bayesian updates asymptotically converge to the true parameter as long as initial prior obeys the Cromwell's rule. The Cromwell's rule (Lindley 1991) requires that no scenarios are assigned with probability one or zero. LS1 satisfied this condition. The construction of the LS1 is only dependent of the objectively-generated draw information, independent from any subjective behavior. Given these good features, we choose LS1 as the benchmark of belief formation and updates. The following four LSs vary the construction of LS1 in various ways.

4.2 LS2: Imperfect Bayesian updates with uniformly distributed initial prior

A subject may adjust his belief in period t on the basis of her belief in the previous period $t-1$, using the Bayesian updates with uniformly distributed initial prior as a reference (Epstein et al. 2010). The Bayesian updates reference is identical to BU_{mt} in LS1. LS2, to a certain extent, applies the Bayes rule to derive the posteriors, but it allows deviations from the Bayesian updates. Hence we denote it as imperfect Bayesian updates with uniformly distributed initial prior. The learning process reads:

$$white_{it} = (1-\gamma_i)white_{i,t-1} + \gamma_i BU_{mt} \tag{15}$$

$$i \in m; \quad t = 2 \cdots 16$$

where $white_{it}$ denotes the self-reported guess game entry of subject i at the beginning of Period t. $white_{i,t-1}$ is the lag variable of $white_{it}$; $t=16$ corresponds with the time point when the last *guess game* is played at the end of period 15. This *guess game* is only played in Session II and III. γ_i denotes the weight that subject i assigns to the Bayesian reference. It varies across subjects but time-invariant within subject. γ_i represents subject i's degree of receptivity to the new information. The higher the γ_i is, the more susceptibly subject i reacts to the newly learned information. On the other hand, $1-\gamma_i$ represents the stickiness of belief. The lower the γ_i is, hence the higher the $1-\gamma_i$ is, the more stubborn subject i sticks to her previous belief.

We estimate γ_i for each i based on equation (15). The by-subject results are reported in Table 7 Column LS2 and summarized in Table 8 Panel A. All $\hat{\gamma}_i$s lie between zero and one and the mean value is 0.352. More than 80% of the $\hat{\gamma}_i$s are less than 0.6, which implies a large proportion of subjects who react to new information with reservation. Epstein et al. (2010) calls it an underreacting to information. On the other hand, $\hat{\gamma}_i > 1$ suggests that subjects attach too much weight to new observations, implying an overreacting to new information. But no cases of $\hat{\gamma}_i > 1$ are observed in our samples. T-tests are also conducted to check whether coefficients are significantly different from one, in other words, whether the coefficients significantly deviate from the Bayesian updates. The conclusion does not change much. Still 34 out of 41 subjects significantly underreact to Bayesian updates. It may be concluded that in general subjects react partially to the new information represented by the Bayesian reference. Epstein et al. 2010 argue that underreaction to the Bayesian updates ($\gamma_i < 1$) eventually converges to the true parameter value. It is rather straightforward, since $\gamma_i < 1$ results in that belief keeps adjusting towards the Bayesian update. Asymptotically belief converges to the Bayesian update, which converges to the true parameter value. However, overreaction to the Bayesian updates ($\gamma_i > 1$) may converge to incorrect parameter value, since it may swing around the Bayesian update but not converge. The estimation results of γ_i imply that most subjects are on the converging track.

4.3 LS3: Bayesian updates with *beta*-distributed initial prior

LS3 follows the thought of the conjugate prior in the Bayesian analyses (Diaconis and Ylvisaker 1979; Gelman et al. 2004; Schlaifer and Raiffa 1961). It assumes that a subject forms her initial prior characterized by a *beta*-distribution, and employs the Bayes rule to update the prior using new information. According to Bernstein von Mises Theorem, LS3 generates belief updates which asymptotically converge to the true parameter value. In comparison with LS1, LS3 eases the assumption inflicted on initial prior: from the only-possible uniformly distributed initial prior to various initial prior shapes. Uniform distribution is still included as a special case. The *beta*-distribution is chosen to simulate the initial prior since it characterizes a wide range of distributions defined in the interval $[0, 1]$, parametrized by only two shape parameters α and β. In case of $\alpha = \beta = 1$, the *beta*-distribution becomes the uniform distribution. Other shapes of *beta*-distribution can be obtained by varying the values of (α, β). Figure 3 illustrates some examples. In case of $\alpha = 1, \beta > 1$, the PDF is strictly decreasing within the domain. In case of $\beta = 1, \alpha > 1$, the PDF is strictly increasing within the domain. In case of α and β are both larger than one, the PDF has a bell-like shape.

Although the *beta*-distribution mostly applies to continuous distributions, the 101 discrete scenarios of white-ball proportion is to some extent dense enough. Later we discretize the *beta*-distribution, translating the PDF of the prior/posterior distribution into PMF to match the 101 discrete scenarios. Since our main interest is the mode of the distribution, the maximum discrepancy between the mode read from the distribution described by continuous PDF and the mode read from the distribution described by discrete PMF is only 0.5 (out of 100). Therefore starting with a continuous prior distribution does little harm. To satisfy the domain criteria of the *beta*-distribution, we first translate the 101 scenarios, indexed by the number of the white balls $W = 0, 1, 2 \cdots 100$, into a corresponding winning probability interval $\theta \in [0, 1]$. Additionally, we restrict $\alpha \geq 1$ and $\beta \geq 1$ to guarantee the uniqueness of mode value, if mode exists.

Under LS3, the initial prior of subject i (in PDF) reads:

$$Prior(\theta|\alpha_i, \beta_i) = \frac{\Gamma(\alpha_i + \beta_i)}{\Gamma(\alpha_i)\Gamma(\beta_i)} \theta^{\alpha_i - 1}(1 - \theta)^{\beta_i - 1} \tag{16}$$

$$\theta \in [0, 1]; \quad i \in \{1, 2, \cdots N\}; \quad \alpha, \beta \geq 1$$

The prior of subject i follows the a *beta*-distribution (PDF) , which is parametrized by α_i and β_i. $\Gamma(\cdot)$ denotes the gamma function. The posteriors are updated employing the Bayes rule. Suppose until the beginning of Period $t \geq 2$, Subject i in Market m witness $n \equiv t-1$ draws with replacement, $k_{n,m}$ of which are white draw(s). The draw is implemented by market, therefore subjects in the same market share an identical draw history. For the time being, we suppress the subscript of $k_{n,m}$ for the simplicity. The posterior of subject i employing the Bayes rule reads:

$$Posterior(\theta|n,k;\alpha_i,\beta_i) = \frac{Prior(\theta|\alpha_i,\beta_i) \times Prob(k|n,\theta)}{\int_{\theta'} Prior(\theta'|\alpha_i,\beta_i) \times Prob(k|n,\theta')d\theta'} \tag{17}$$

where

$$Prob(k|n,\theta) = \binom{n}{k} \theta^k (1-\theta)^{n-k} \tag{18}$$

$$\binom{n}{k} = \frac{n!}{(n-k)!} \tag{19}$$

Plugging equation (16)(18) into (17) yields

$$\begin{aligned} Posterior(\theta|n,k;\alpha_i,\beta_i) &= \frac{\theta^{\alpha_i-1}(1-\theta)^{\beta_i-1}\theta^k(1-\theta)^{n-k}}{\int_{\theta'} \theta'^{\alpha_i-1}(1-\theta')^{\beta_i-1}\theta'^k(1-\theta')^{n-k}d\theta'} \\ &= \frac{\theta^{\alpha_i+k-1}(1-\theta)^{\beta_i+(n-k)-1}}{\int_{\theta'} \theta'^{\alpha_i+k-1}(1-\theta')^{\beta_i+(n-k)-1}d\theta'} \end{aligned} \tag{20}$$

Since

$$\int_\theta Prior(\theta|\alpha_i,\beta_i)d\theta = \frac{\Gamma(\alpha_i+\beta_i)}{\Gamma(\alpha_i)\Gamma(\beta_i)} \int_\theta \theta^{\alpha_i-1}(1-\theta)^{\beta_i-1}d\theta = 1 \tag{21}$$

Hence

$$\int_\theta \theta^{\alpha_i-1}(1-\theta)^{\beta_i-1}d\theta = \frac{\Gamma(\alpha_i)\Gamma(\beta_i)}{\Gamma(\alpha_i+\beta_i)} \tag{22}$$

Replacing α_i with $\alpha_i + k$ and β_i with $\beta_i + (n-k)$ in equation (22) yields:

$$\int_\theta \theta^{\alpha_i+k-1}(1-\theta)^{\beta_i+(n-k)-1}d\theta = \frac{\Gamma(\alpha_i+k)\Gamma(\beta_i+n-k)}{\Gamma(\alpha_i+k+\beta_i+n-k)} \tag{23}$$

Plugging (23) into (20) yields:

$$Posterior(\theta|n,k_{n,m};\alpha_i,\beta_i) = \frac{\Gamma(\alpha_i+k_{n,m}+\beta_i+n-k_{n,m})}{\Gamma(\alpha_i+k_{n,m})\Gamma(\beta_i+n-k_{n,m})}\theta^{\alpha_i+k_{n,m}-1}(1-\theta)^{\beta_i+(n-k_{n,m})-1} \tag{24}$$

$$\theta \in [0,1]; \quad i \in m; \quad \alpha_i, \beta_i \geq 1; \quad k_{n,m} \leq n \equiv t-1 \leq 15, \quad k_{n,m}, n \subseteq \mathbb{Z}$$

Therefore the posterior inherits the *beta*-distribution PDF fashion, with the updated parameters $\alpha_i + k_{n,m}$ and $\beta_i + (n - k_{n,m})$. The intuition is rather obvious: $k_{n,m}$ denotes the number of white draws out of n draws implemented in market m (ie. the frequency of white draw). Whereas $n - k_{n,m}$ denotes the frequency of black draw. If the white draw is more frequently observed than the black draw, the α-parameter increases relatively more quickly than the β-parameter, reshaping the posterior *beta*-distribution more left-skewed and resulting in a higher value of mode. On the other hand, if more black draw is observed, the β-parameter increases relatively more quickly than the α-parameter, causing the posterior *beta*-distribution more right-skewed and resulting in a lower value of mode. This property of *beta*-distribution updates assumes that subjects upwards adjust their beliefs if they witness proportionally more white balls, and downwards adjust their beliefs if they witness proportionally more black draws, which is intuitive. In addition, when time evolves, more draws are implemented. n increases for sure and $k_{n,m}$ is on the track of increase. Hence both $\alpha_i + k_{n,m}$ and $\beta_i + (n - k_{n,m})$ are on the track of increase. The PDF of *beta*-distribution with larger parameter bundle concentrates more densely on some narrower bandwidth. The bell-shape PDF curve looks thinner and taller. It is also intuitive, since when learning evolves, subjects tend to feel more confident about their beliefs, and assign more probabilities to a limited number of scenarios.

The estimation of the shape parameter bundle (α_i, β_i) of each subject i is done in the following procedure. Like in LS1, a subject translates her posterior distribution into her *guess game* entry by reporting the mode of the posterior distribution, since they are financially incentivized to do so. The mode function of a *beta*-distribution with parameters $\alpha > 1$ and $\beta > 1$ is written as $\frac{\alpha-1}{\alpha+\beta-2}$, and thus we can write:

$$\frac{white_{it}}{100} = \frac{\alpha_i + k_{n,m} - 1}{\alpha_i + \beta_i + n - 2} \tag{25}$$

$$i = 1, 2 \cdots N; \quad \alpha_i, \beta_i > 1; \quad k_{n,m} \leq n \equiv t - 1 \leq 15, \quad k_{n,m}, n \subseteq \mathbb{Z}, \quad k_{0,m} = 0$$

Re-arrange to achieve:

$$100k_{n,m} - white_{it}(n-2) - 100 = (white_{it} - 100)\alpha_i + white_{it}\beta_i \tag{26}$$

In order to restrict $\alpha_i, \beta_i > 1$, we first transfer α_i and β_i into the reciprocals of the standard logistic function. Since a standard logistic function ($logistic(x) = /(1 + e^{-x})$) has a co-domain $(0, 1)$, the co-domain of its reciprocal is $(1, \infty)$. The regression equation reads:

$$100k_{n,m} - white_{it}(n-2) - 100 = \frac{white_{it} - 100}{logistic(a_i)} + \frac{white_{it}}{logistic(b_i)} \tag{27}$$

$$\alpha_i = [logistic(a_i)]^{-1} \tag{28}$$

$$\beta_i = [logistic(b_i)]^{-1} \tag{29}$$

$$i \in m; \quad k_{n,m} \leq n \equiv t - 1 \leq 15, \quad k_{n,m}, n \subseteq \mathbb{Z}, \quad k_{0,m} = 0$$

First we run the nonlinear regression based on equation (27) to estimate $(a_i,\ b_i)$ for each subject i. Then we recover $(\alpha_i,\ \beta_i)$ using (28) and (29), respectively. If $\hat{a}_i$ is large enough, α_i is estimated by unity. Analogously, if $\hat{b}_i$ is large enough, β_i is estimated by unity. The parameter estimation results are reported in Table 7 column LS3 and Table 8 Panel B and Figure 4. As is shown that most of the subjects formulate their initial priors characterized by low $\hat{\alpha}_i$, $\hat{\beta}_i$ values, which correspond with a rather flat PDF curve. More than one third of the subjects assign the probability perfectly evenly ($\hat{\alpha}_i = \hat{\beta}_i = 1$, uniformly distributed initial prior, equivalent to LS1).

Furthermore, we discretize the continuous prior and posterior distributions (in PDF form) to the PMF of the 101 discrete scenarios ($W = 0, 1, \cdots 100$) in the following way:

$$Prior(W) = \frac{Prior(\theta)}{\sum_{\theta'} Prior(\theta')} \tag{30}$$

$$Posterior(W) = \frac{Posterior(\theta)}{\sum_{\theta'} Posterior(\theta')} \tag{31}$$

$$\text{for each } W \equiv 100\theta \in \{0, 1, \cdots 100\}; \quad \theta' \in \{0, 0.01, 0.02 \cdots 1\}$$

$Prior(W)/Posterior(W)$ is the PMF for urn scenario with W units of white balls, while $Prior(\theta)/Posterior(\theta)$ is the PDF value for winning probability $\theta = W/100$. The translation is done for each subject in each period. The $Prior(W)$ (PMF of initial prior) is illustrated in Figure 4b. As $(\hat{\alpha}_i, \hat{\beta}_i)$ suggests, there exists heterogeneity in the formation of initial prior. $(\hat{\alpha}_i = 1, \hat{\beta}_i = 1)$, $(\hat{\alpha}_i > 1, \hat{\beta}_i = 1)$, $(\hat{\alpha}_i = 1, \hat{\beta}_i > 1)$, $(\hat{\alpha}_i > 1, \hat{\beta}_i > 1)$ implying uniformly distributed, increasing, decreasing, bell-shaped initial priors, respectively, are all seen in the samples.

4.4 LS4: Imperfect Bayesian updates with *beta*-distributed initial prior

LS4 assumes that a subject adjusts her belief in Period t on the basis of her belief in the previous period $t-1$, using Bayesian updates with *beta*-distributed initial prior as a reference. The Bayesian updates reference is constructed analogously to LS3. In comparison to LS3, LS4 allows belief to deviate from the Bayesian updates. This flexibility is facilitated by the choice of γ_i^c. In comparison to LS2, LS4 shares the "imperfect" feature, but the Bayesian updates start with a *beta*-distributed initial prior, not restricted to the uniformly distributed initial prior as in LS2. Hence, LS4 is so far the most flexible learning strategy compared with the other three. It is modeled as follow:

$$white_{it} = (1 - \gamma_i^c) white_{i,t-1} + \gamma_i^c \frac{100(\alpha_i^c + k_{n,m} - 1)}{\alpha_i^c + \beta_i^c + n - 2} \tag{32}$$

$$i \in m; \quad k_{n,m} \leq n \equiv t - 1 \leq 15, \quad k_{n,m}, n \subseteq \mathbb{Z}, \quad k_{0,m} = 0$$

To avoid confusion, we upper-script the parameters with a c to distinguish those in LS2 and LS3. For each subject i, we estimate three parameters, γ_i^c, α_i^c and β_i^c, using the data from

$t = 1, \cdots 16$ based on equation (32). Again we restrict $\alpha_i^c, \beta_i^c \geq 1$ by first estimating a_i^c and b_i^c analogously as in equation (28)(29) and later recovering α_i^c and β_i^c. In addition, we restrict $\gamma_i^c > 0$ by estimating r_i^c first, where $\gamma_i^c = exp(r_i^c)$. Later we recover γ_i^c. A negative enough $\hat{r}_i^c$ is estimated by $\hat{\gamma}_i^c = 0$. Underreaction to the Bayesian reference leads to convergence to the true parameter value, while overreaction may converge to incorrect values. The reasoning is analogous to that in LS2.

The by-subject results are reported in Table 7 column LS4 and summarized in Table 8 Panel C. The mean value of $\hat{\gamma}_i^c$ is 0.732, indicating that on average subjects reluctantly deviate from their current beliefs, only partially responding to the new information represented by the Bayesian reference. If significance is taken into account, still about half of the subjects (21 out of 40) present clear tendency to underreact to the Bayesian reference. As for the initial prior, Figure 5a presents the distribution of estimated bundle $(\hat{\alpha}_i^c, \hat{\beta}_i^c)$ (same as in Table 8 Panel C). Initial prior is most likely to be formed as a uniform distribution (24 out of 40 subjects). Figure 5b recovers the initial prior distribution based on the estimation results. It is shown that the initial prior formation in LS4 is comparable to that in LS3: most of the subjects tend to come up with a uniformly distributed prior.

4.5 LS5: Heuristic updates

LS3 models the belief updating process such that a subject upwards adjusts her current belief each time a white draw is observed, proportional to the maximum possible adjustment in the upward direction. Analogously, she downwards adjusts her current belief each time a black draw is observed, proportional to the maximum possible adjustment in the downward direction. Compared to LS1-LS4, LS5 seems to be more straightforward and heuristic. LS1-LS4 model beliefs by specifying the probability distribution across all possible scenarios. LS5 simply suppresses this specification, directly attending to the mode value of the distribution (ie. the response in the *guess game*). As for the updating rule, LS5 adopts a rather heuristic strategy which has nothing to do with the Bayes rule in any visible way. Such heuristic feature pays the price in the way that belief updates by LS5 probably do not converge to the true parameter value. The learning process reads:

$$\begin{aligned} white_{it} - white_{i,t-1} = &\delta_{i,w} 1[white]_{m,t-1}(100 - white_{i,t-1}) \\ &+ \delta_{i,b} 1[black]_{m,t-1}(0 - white_{i,t-1}) \end{aligned} \tag{33}$$

$$i \in m; \quad t = 2 \cdots 16$$

$1[white]_{m,t-1}$ is a dummy variable, equal to one if at the end of period $t-1$ market m observes a white draw, otherwise zero. Analogously, $1[black]_{m,t-1}$ is equal to one if at the end of period $t-1$ market m observes a black draw, otherwise zero. $\delta_{i,w}$ captures the belief adjustment of subject i when she observes a white draw. $\delta_{i,b}$ captures the adjustment when a black draw is observed. $\delta_{i,w}$ and $\delta_{i,b}$ are heterogeneous across subjects but time-invariant for a given subject. A positive $\delta_{i,w}$ implies rational belief updating: upward adjustment if a white draw is observed. Similarly, a positive $\delta_{i,b}$ implies downward adjustment if a black draw is observed. The estimation of $\delta_{i,w}$ and $\delta_{i,b}$ for each i are reported in Table 7 column LS5 and summarized in Table 8 Panel D. The results imply adjustment rationality as we expect: $\hat{\delta}_{i,w}$ and $\hat{\delta}_{i,b}$ tend to be positive. Significance test for whether coefficients are different from zero reduces the number of positive pairs, but still leaves 24 out 41 subjects who have at least one significantly positive coefficient. Negative coefficients imply that such subjects update their beliefs probably irrationally. It is justified by the low rates of rational updates of these subjects, shown in Table 6. In addition, the mean values of $\hat{\delta}_{i,w}$ and $\hat{\delta}_{i,b}$ are not significantly different (p-value=0.181), both around $0.11 - 0, 12$. It implies that subjects tend to react to white and black draws symmetrically, around $11\% - 12\%$ proportionally to the maximum possible adjustment.

5 Characteristic of learning strategy

In this section, we summarize the characteristics of the five learning strategies, using the fitted values obtained from the estimations. We already estimate the subject-specific parameters which characterizes LS2-LS5 using the *guess game* data in the previous section. Consequently we can recover the fitted values of the estimations. Now we use the fitted values to summarize three characteristics of each learning strategy: the initial prior, the goodness of fit to the factual data, the speed of convergence.

5.1 Initial prior

Based on the assumptions, LS1 and LS2 assume deterministic uniformly distributed initial priors. LS3 and LS4 instead assume that initial priors follow a *beta*-distribution characterized by some parameter bundle (α_i, β_i). The parameter bundle is subject-specific and therefore a subject is endowed with the flexibility to form her own initial prior. For LS3 and LS4, we make great efforts to recover the specific shape of the subjective initial priors. LS5, however, acts indifferent to the shape of the distribution but directly focuses on the final interest, the mode of the distribution. It is simply recorded as $white_{i1}$ in the data. Therefore, in this part, we only focus on the characteristics of the initial priors in LS3 and LS4.

As a reminder, the estimated parameter bundles ($\hat{\alpha}_i$, $\hat{\beta}_i$) of LS3 for each subject are reported in Table 7 column LS3, Table 8 Panel B and illustrated in Figure 4. It shows that there exists some heterogeneity in terms of initial prior formation. As is shown in Figure 4a, The big bubble circling coordinate (1,1) represents 37% (15 out of 41) of the subjects who formulates a uniformly distributed prior, parameterized by $\hat{\alpha}_i = \hat{\beta}_i = 1$. The other circles lie on the vertical line $\alpha = 1$ corresponds with a relatively pessimistic initial prior, parameterized by $\hat{\alpha}_i = 1$, $\hat{\beta}_i > 1$ and thus the mode of the initial prior distribution rests at 0. The initial prior shapes decreasingly from the peak at $\theta = 0$ to the trough at $\theta = 1$. This initial prior choice accounts for 24% (10 out of 41) of the subjects. Analogously, the circles lie on the horizontal line $\beta = 1$ corresponds with a relatively optimistic initial prior, parameterized by $\hat{\beta}_i = 1$, $\hat{\alpha}_i > 1$, and thus the mode of the initial prior distribution finds itself at 100. The initial prior shapes increasingly from the trough at $\theta = 0$ to the peak at $\theta = 1$. 20% (8 out of

41) of the subjects formulates their initial priors in such fashion according to the estimation. The rest of the subjects, accounting for 20% (8 out of 41), start with a bell-shape initial prior, parameterized by $\hat{\alpha}_i > 1$, $\hat{\beta}_i > 1$, and thus the mode of the initial prior finds itself at an internal singular value. Although the initial prior formations are heterogeneous across subjects, most of the initial priors tend to shape rather flat. It is not a surprise that subjects tend to assign the probability rather evenly to the 101 scenarios when the situation is completely ambiguous. For those initial priors which appear to be non-flat, Figure 4b shows that, as a matter of fact, no large PMF is assigned to some specific scenarios. The maximum PMF is merely as large as 0.02. It turns out that most subjects have a rather flat-shaped initial prior.

The estimated initial prior parameter bundles ($\hat{\alpha}_i^c$, $\hat{\beta}_i^c$) of LS4 are reported in Table 7 column LS4 and Table 8 Panel C, and illustrated in Figure 5 . LS4 recovers more flat-shaped initial priors than LS3 does: 24 out of 41 subjects see $\hat{\alpha}_i^c = \hat{\beta}_i^c = 1$, implying a uniformly distributed initial prior. Four subjects come up with an optimistic initial prior ($\hat{\alpha}_i > 1$, $\hat{\beta}_i = 1$), less than the eight in LS3. Eight subjects formulate a pessimistic initial prior ($\hat{\alpha}_i = 1$, $\hat{\beta}_i > 1$), less than the ten in LS3. Only four subjects formulate an initial prior which has an internal mode value ($\hat{\alpha}_i, \hat{\beta}_i > 1$), less than the eight in LS3.

In conclusion, if we account for possible different initial priors, although there exists heterogeneity regarding the initial prior formation, subjects are most likely to be found to adopt uniformly distributed initial priors. It seems that simplicity, symmetry, and high degree of ambiguity prevail regarding the initial prior formation.

5.2 Goodness of fit

In this part, we analyze how well the belief update dynamic following one learning strategy fits the factual belief update dynamic, and which learning strategy of the five fits it best. The guess game entries of subject i ($white_{it}$; $t \in \{0, 1, \cdots 16\}$) records her belief updating dynamic. Therefore the problem comes down to analyzing the difference between the belief updates simulated by one of the five learning strategies and the factual belief updates recorded by the guess game. We first compute the average difference over all subject for each learning strategy. The difference is expressed by its absolute value to avoid the cancellation between positive and negative differences during summation. The average differences between the factual belief updates and the simulated belief updates by LS1-LS5 are written, respectively, as follows:

$$\text{diff(guess, LS1)}_t = \frac{1}{N}\sum_{i=1}^{N}|white_{it} - BU_{mt}| \tag{34}$$

$$\text{diff(guess, LS2)}_t = \frac{1}{N}\sum_{i=1}^{N}|white_{it} - (1-\hat{\gamma}_i)white_{i,t-1} - \hat{\gamma}_i BU_{mt}| \tag{35}$$

$$\text{diff(guess, LS3)}_t = \frac{1}{N}\sum_{i=1}^{N}|white_{it} - \frac{100(\hat{\alpha}_i + k_{n,m} - 1)}{\hat{\alpha}_i + \hat{\beta}_i + t - 3}| \tag{36}$$

$$\text{diff(guess, LS4)}_t = \frac{1}{N}\sum_{i=1}^{N}|white_{it} - (1-\hat{\gamma}_i^c)white_{i,t-1} - \hat{\gamma}_i^c \frac{100(\hat{\alpha}_i^c + k_{n,m} - 1)}{\hat{\alpha}_i^c + \hat{\beta}_i^c + t - 3}| \tag{37}$$

$$\begin{aligned}\text{diff(guess, LS5)}_t = \frac{1}{N}\sum_{i=1}^{N}|white_{it} - white_{i,t-1} - \hat{\delta}_{i,w}1[white]_{m,t-1}(100 - white_{i,t-1}) \\ - \hat{\delta}_{i,b}1[black]_{m,t-1}(0 - white_{i,t-1})|\end{aligned} \tag{38}$$

$$i \in m; \quad n \equiv t-1; \quad t \in \{1, \cdots 16\}$$

where N denotes the number of the subjects. The results are reported in Table 9 and Figure 6. BU_{mt} in LS1 and LS2 has no unique value at $t = 1$, since a uniformly distributed initial prior has no unique mode value. And LS2, LS4, LS5 includes the lagged variable $white_{i,t-1}$. Therefore the difference at $t = 1$ is only applicable to LS3. Only subjects in session II and III play the guess game at $t = 16$, so N reduces from 41 to 28 at $t = 16$.

The results are reported in Table 9 and illustrated in Figure 6. As is shown in Figure 6, all differences tend to shrink when learning evolves. In periods $t = 3 \cdots 16$, LS4 tends to fit the factual belief updates best. Flexibility seems to be the source of the improvement of goodness of fit: LS4 proposes initial prior not to be restricted to uniform distribution and updating rule to allow deviation from Bayesian reference. This argument is also supported by the fact that LS3 fits the factual beliefs better than LS1 in all periods. The heuristic strategy LS5 performs more or less averagely, dominating LS1 and LS2 for most periods. The goodness of fit of each LS tends to converge to a similar level as time evolves, around 5-10 units. This remaining discrepancy may be explained by the irrational proportion of the factual updates.

5.3 Speed of convergence

The speed of convergence of one learning strategy describes its speed of convergence to a constant number, or a narrow-windowed interval (for instance, fluctuates only ± 3 units out of the maximum width 101 units). The analysis does not restrict the constant number to be the true proportion of the white balls in the ambiguous urn, since the experiment only permits 15 times of learning, a finite case. In practice it is hard to see the belief update converge to the true parameter value. The speed of convergence of one learning strategy focus on how fast the fluctuation of the belief adjustment following one learning strategy undermines as learning evolves. The inter-temporal belief adjustment of each learning strategy computed by the estimated parameters of subject i is written as:

$$\text{Diff}(j)_{it} = \widehat{white}(j)_{i,t} - \widehat{white}(j)_{i,t-1} \tag{39}$$

$$j \in \{LS1 \cdots LS5\}; \quad i \in \{1, 2 \cdots N\}; \quad t \in \{2, \cdots 16\}$$

where $\widehat{white}(j)_{i,t}$ denotes the simulated belief update of subject i in period t, simulated by learning strategy j. $\widehat{white}(j)_{i,t-1}$ is its lagged variable. The inter-temporal belief adjustment of each learning strategy is then averaged over all subjects.

$$\overline{\text{Diff}}(j)_t = \frac{1}{N}\sum_{i=1}^{N} \text{Diff}(j)_{it} \tag{40}$$

$$j \in \{LS1 \cdots LS5\}; \quad t \in \{2, \cdots 16\}$$

The results of $\overline{\text{Diff}}(j)_t$ (by j and t) are reported in Table 10 and Figure 7. As is shown that LS5 seems to converge most slowly. Among LS1-LS4, LS2 seems to outperforms the other three in the early periods, but the other three catch up after period 8. The speed of convergence differs slightly among LS1-LS4 after period 10. LS1/LS3 seems to have a smaller standard deviation, meaning that the speed of convergence of LS1/LS3 is more robust across subjects. It can be concluded that a learning strategy initializing with a well-described prior distribution (like in LS1-LS4) can more easily converge to some constant number than a learning strategy initializing with a heuristic prior (like in LS5). Among the learning strategies initializing with a well-described prior distribution, those that fully follow the Bayes rule (LS1/LS3) see its fast convergence more robust across subjects than those which partially refer to the Bayesian reference (LS2/LS4).

As a reference, we derive the speed of convergence of the factual guess game entries. Analogously the Difference between the guess game entry and its lagged term of each subject is computed by:

$$\text{Diff}(guess)_{it} = white_{it} - white_{i,t-1} \tag{41}$$

$$i \in \{1, 2 \cdots N\}; \quad t \in \{2, \cdots 16\}$$

Then averaging over all subjects for given t:

$$\overline{\text{Diff}}(guess)_t = \frac{1}{N}\sum_{i=1}^{N} \text{Diff}(guess)_{it}; \qquad t \in \{2, \cdots 16\} \tag{42}$$

The results are reported in Table 10 and Figure 7. In comparison with the convergence speed of the belief updates generated by LS1-LS5, the factual belief updates seem to converge more slowly. The temporal factual belief adjustment is usually around ten units (out of 100) between two neighboring periods. Although the adjustment in the early periods tend to be smaller than LS1-LS5, the later periods still sustain rather large adjustment. There exists large heterogeneity across subjects in terms of convergence speed of the factual belief updates, supported by large standard deviation shown in Table 10. In practice, the belief updates are more unstable than the five learning strategies suggest.

6 Heterogeneity

Based on the parameter estimation, we observe large heterogeneity across the 41 subjects in terms of the initial prior formation and the choice of belief updating. In this section, we try to answer the question: who chooses which LS. In order to answer this question, we define that the initial prior and learning strategy actually formed/chosen by subject i produce a belief update dynamic which best fits the factual belief update dynamic observed from subject i in the *guess game*. In addition, since learning strategy is a dynamic concept, the comparison of the goodness of fit between two learning strategies must base on their performances over a multi-period duration, rather than a point-wise comparison in some specific period. The criteria are thus set as follows: For a given subject i, a given learning strategy $j \in \{\text{LS1} \cdots \text{LS5}\}$, we compute the root mean squared error (RMSE). The error arises from the difference between i's factual belief and the belief generated by learning strategy j in each period, and the mean refers to averaging the errors over period. Therefore the RMSE of subject i for a given learning strategy j reads:

$$RMSE_{i,j} = \sqrt{\frac{1}{T}\sum_{t}[white_{it} - \widehat{white}(j)_{i,t}]^2} \tag{43}$$

$$j \in \{LS1 \cdots LS5\}; \quad t \in \{2, \cdots 16\}$$

where $white_{it}$ denotes the guess game entry of subject i in period t; $\widehat{white}(j)_{i,t}$ denotes the simulated belief of subject i in period t simulated by the Learning strategy j. For subjects from Session I, only 15 guess games are played, hence the final period $T = 15$. LS1, LS2, LS4 and LS5 cannot recover the mode value of initial priors, since LS1 assumes a uniform distribution which has no unique mode and the other three include a one-period lagged variable in the regression. For the comparability, $RMSE$ is computed only including $t \geq 2$, and $T = 15$ for Session I while $T = 16$ for Session II and III. The learning strategy j which achieves the smallest $RMSE_{i,j}$ value among the five learning strategies is defined as the chosen learning strategy of subject i. Table 11 and Figure 8 report the ranking of LS using the RMSE criteria. Table 12 summarizes the chosen learning strategy by subject, and accordingly recovers the corresponding initial prior formation and the choice of belief updating rule.

As is shown in Figure 8, most of the subjects see their best fitted LS has a rather low RMSE, below 10 (out of 100). It implies that for most of the subjects in our samples, there exists at least one LS among the five, which explains the factual belief updates very well. Outliers are Subject 1, 3, 15, 25, 26, and 39, whose updates cannot be explained by any of the five LSs well. It may be due to the fact that these subjects update their beliefs irrationally for a substantial proportion. This argument is supported by Table 6, which shows that the rate of rational updates of these subjects are relatively low. In fact, the highest is merely 0.6, lower than the full sample average 0.79. It is not a surprise to see that subjects with substantial irrational updates are poorly explained by any of the five LSs.

As for "who chooses which LS", Table 11 shows that LS5 is mostly chosen, followed by LS2 and LS4. Regarding the initial prior formation, as is shown in Table 12 Panel B, nearly half of the subjects (20 out of 41) form an initial prior which can be specifically characterized by either a uniform distribution (17 out of 20) or a non-uniform distribution (6 out of 20). The non-uniform distributions, however, are mostly flat-shaped: the *beta*-distribution shape parameters are mostly close to one, indicating flat PDF curves. The other half (20 out of 41) tends to keep the initial prior distribution unspecific and directly attend to the mode of the distribution.

In regard of the belief updating rule, the subjects, whose initial prior distribution is specific, mostly underreact to the Bayesian updates. Table 12 Panel B shows that 17 out of 20 subjects have a $\hat{\gamma}$ ($\hat{\gamma}^c$) which is less that one. Even if the statistical significance (significance level=95%) is taken into account, still 11 out of 20 subjects underreact to the Bayesian updates, one overreacts, and eight subjects seem to perfectly follow the Bayesian updates. Among the eight alleged Bayesian players, only four of them (subject No. 14, 18, 35, 38, see Table 8) can probably be seen as the real Bayesian player, justified by their low RMSE values generating by the Bayesian updates (by LS3). The other four (subject No. 1, 15, 25, 39, see Table 8), as discussed above, should be categorized into irrational players instead of Bayesian players, since their RMSEs for each LS are relatively high. For those who do not reveal a specific initial prior distribution, the belief updating follows a rather heuristic strategy. They tend to simply upwards adjust their beliefs if a white draw is observed, and downwards adjust their beliefs if a black draw is observed. The belief adjustment turns out to be mostly rational and proportional to the maximum possible adjustment along the rational direction. It is justified by the mean value of $\hat{\delta}_{i,w}$ which is 0.127, and the mean value of $\hat{\delta}_{i,b}$ which is 0.114. Such value corresponds with 5-10 units of white balls. In regard of the choice of the belief updating rule, it may be concluded that subjects are mostly non-Bayesian players: either deviating from the Bayesian updates path (12 out of 40), or adopting heuristic updating strategy in no relation to the Bayes rule (20 out of 40).

7 Pessimism/optimism

Pessimism/optimism in general describes the perceptive tendency that one biases towards the good/bad outcomes of an event. In an ambiguity setting, pessimism/optimism can be defined as bias towards the good/bad scenarios when forming the second-order probabilities for each possible scenarios in an ambiguous environment. If such a bias exists in an ambiguity setting, it should be expected to weaken or disappear when learning makes the environment less ambiguous.

In order to investigate the bias, we quantify the pessimism/optimism for each subject in each period, in the way that we measure the difference between the mode of one's factual belief and the mode of a chosen benchmark. It is clear that the discussion of pessimism/optimism relies on the choice of a benchmark. Hence the paper proceeds as follow: we first propose some choices for benchmark and discuss their quality. Subsequently, using the chosen benchmarks, we measure the average pessimism/optimism across subjects.

7.1 Choice of benchmarks

As discussed, pessimism/optimism can be diagnosed by the belief bias in comparison with a chosen benchmark. Question arises for how to choose a benchmark. Prior to learning (at the beginning of Period 1), it is intuitive to choose the mid-point, 50, as the benchmark. Without any draw information, all 101 scenarios (indexed by the possible number of the white balls: 0, 1, $\cdots$ 100) are equally possible. For a rational subject, the probability of inserting one integer (inside the designed domain) in the *guess game* at $t = 1$ should be equal to the probability of inserting any other integer (inside the designed domain). Asymptotically speaking, if infinite many subjects play the *guess game* at $t = 1$, the inserted answers should form a uniform distribution between $[0, 1]$. Therefore, we naturally choose the mean value of this distribution, 50, as the benchmark for the initial belief at $t = 1$.

From $t = 2$ onwards, how should one update her belief responding to the draw information is a open question. 50, or any fixed number, is not any more a good candidate since it ignores the effect of new information. The benchmark is supposed to be a dynamic process, neutrally updating the belief responding to the new information. It is challenging to argue one belief update process is better than another. Not to mention to define the best way of updating. However, some basic criteria can be discussed for a benchmark: (a) It should start with neutrality, defined as neither pessimistic nor optimistic, in the initial prior formation. It implies that the initial prior distribution should be symmetric about the mid-point 50. It results in that no preexisting bias exists, and it also seamlessly fit in the chosen benchmark in period $t = 1$ (b) It should demonstrate clear and unique mode value, at least for most of the time. (c) It should rationally respond to the draw information: when good news comes, update the belief towards the good direction; when bad news comes, bad direction. (d) It should asymptotically converge to the true parameter value: in our case, the mode value of the updated posterior should converge to 40, the true proportion of white balls in the urn. Taken these criteria into account, a *beta*-distribution with $\alpha = \beta > 1$ seems to be a straightforward candidate, since by construction it satisfies criteria (a)(b). (c)(d) can be satisfied by the application of the Bayes rule. Therefore, we propose six belief updating benchmarks, all of which incorporate the basic criteria.

The six benchmarks are identical such that they all employ Bayes rule and perfectly follow the Bayesian updates. They are differentiated in terms of the initial prior distribution:

B1: uniform distribution; (identical to LS1)
B2: a symmetric triangle distribution: linear, increasing (decreasing) between [0, 0.5] ([0.5, 1]), and kink at 0.5 (see Figure 9)
B3: *beta*-distribution with shape parameters $\alpha = \beta = 2$;
B4: *beta*-distribution with shape parameters $\alpha = \beta = 3$;
B5: *beta*-distribution with shape parameters $\alpha = \beta = 10$;
B6: *beta*-distribution with shape parameters $\alpha = \beta = 50$;

The initial prior distributions (in PMF) of all six benchmarks are illustrated in Figure 9. We propose them as the benchmarks since they all satisfy the criteria mentioned above at least from Period 2 onward. The belief updating, employing the Bayes rule, can be written as:

$$Posterior^B(W|n, k_{n,m}) = \frac{(W/100)^{k_{n,m}}(1 - W/100)^{n-k_{n,m}} \times Prior^B(W)}{\sum_{j=0}^{100}(j/100)^{k_{n,m}}(1 - j/100)^{n-k_{n,m}} \times Prior^B(j)} \tag{44}$$

$$W, j \in \{0, 1, 2 \cdots 100\}; \quad 0 \leq k_{n,m} \leq n \equiv t - 1 \leq 15,$$
$$k_{n,m}, n \subseteq \mathbb{Z}; \quad m \in \{1, 2 \cdots 5\}; \quad B \in \{B1 \cdots B6\}$$

$$BU^B_{mt} = \underset{W}{\text{argmax}} Posterior^B(W|n, k_{n,m}); \quad B \in \{B1 \cdots B6\}; \tag{45}$$

B1, as defined, is identical to LS1. B2-B6 are analogous to LS1, with the different initial priors, $Prior^B(W)$ for $B \in \{B2 \cdots B6\}$, respectively. We suppress the explicit functional forms for $Prior^B(W)$, but it is easy to construct the linear PMF function for B2 and recover the *beta*-distribution for B3-B6. BU^B_{mt} recovers the Bayesian updates of market m in period t, starting with the initial priors whose shape is characterized by $B \in \{B1 \cdots B6\}$.

As a further quality check, goodness of fit of each benchmark is computed to check whether one benchmark captures the factual belief updates well. A bad goodness of fit implies that one updating process is rarely chosen by subjects in practice. The benchmark belief updates ought to also consider the pragmatism and feasibility. Hence, a bad goodness fit may reduce the candidature power of one benchmark. The goodness of fit of B1-B6 are computed as follow:

$$\text{diff(guess, v)}_t = \frac{1}{N}\sum_{i=1}^{N} |white_{it} - BU^v_{mt}|; \tag{46}$$

$$i \in m, \quad t = 2 \cdots 16 \quad v \in \{V1 \cdots V5\}$$

The results are reported in Table 13 Panel A and illustrated in Figure 10a. As can be seen that, B1 present the best goodness of fit from Period 5 onwards, but performs worst in early periods. It may be due to fact that B1, by construction, always generates the mode equal to either extreme value (0 or 100) in Period 2. B2-B4, similar to each other, tend to be the top three among the five. B5-B6 present the worse goodness of fit in later periods. It may be explained by the construction of the initial prior distribution which concentrate too much probability to scenario $W = 50$.

7.2 Analysis of pessimism/optimism

The pessimism/optimism in Period 1 is rather straightforward to diagnose. The mode value of the initial priors (the response inserted in *guess game* at $t = 1$) is on average equal to 48.56 (N=41, sd.=10.28, min=23, max=70). It appears to be lower than the chosen benchmark 50, but the t-test shows that it is not significantly different from 50 ($H_0 : white_{t=1} = 50$, p-value=0.375). It means that subjects tend to be neither pessimistic nor optimistic at first.

In period $t = 2 \cdots 16$, We compute the exact differences between the factual belief updates and each of the benchmark updates by subject by period. The average exact difference across subjects (in a given period, using one particular benchmark) presents the average pessimism/optimism (in that period, using that particular benchmark). The equation reads:

$$\text{diff(guess, B)}_t^{exact} = \frac{1}{N}\sum_{i=1}^{N}(white_{it} - BU_{mt}); \quad i \in m, \quad t = 2 \cdots 16 \tag{47}$$

$$i \in m, \quad t = 2 \cdots 16 \quad B \in \{B1 \cdots B6\}$$

diff(guess, B)$_t^{exact}$ describes the average pessimism (if diff< 1) or optimism (if diff> 1) of subjects in period t, using benchmark $B \in \{B1 \cdots B6\}$. The results are reported in Table 13 Panel B. Each column represents the analysis using one specific benchmark. Figure 10b illustrates the same results. As can be seen in Table 13, significantly negative difference values are more frequently seen in early periods than in later periods. It may imply that subjects tend to demonstrate slight pessimism in situations with higher degree of ambiguity, and that the pessimism disappears when the degree of ambiguity decreases as more information about the urn is available.

8 Robustness check

8.1 LS2 variation

We consider a variation of LS2 with the homogeneous and time-variant γ_t:

$$\begin{aligned} white_{it} =& (1-\gamma_t) white_{i,t-1} + \gamma_t BU_{mt} \\ & i \in m; \quad t = 2 \cdots 16 \end{aligned} \tag{48}$$

We report the estimation results for both full sample and the rational sub-sample. As defined in descriptive analysis, rational belief updates are those $white_{it}$ such that $white_{it} \geq white_{i,t-1}$ if a white draw is observed at the end of period $t-1$, and $white_{it} \leq white_{i,t-1}$ if a black draw is observed at the end of period $t-1$. The results of $\hat{\gamma}_t$ are reported in Table 14. All $\hat{\gamma}$s are between zero and one. It implies that on average subjects underreact to the new information. Intertemporally, on average subjects tend to more actively react to the new information from period 6 to 11, less actively at the beginning and at the end. It is intuitive to react less actively when learning evolves for many periods. The marginal effect of the new information on reshaping the belief is diminishing. The underreaction in early periods may be explained by the wait-and-see attitude. The scarcity of new information may keep subjects sticky to their initial beliefs. The estimation based on the rational sample only slightly differs from that based on the full sample. The relatively large differences appear at $t = 1, 14, 15$, and the results are more robust at $t \in [2, 13]$.

8.2 LS5 variation

First we consider a variation of LS5 with the homogeneous and time-variant $\delta_{w,t}$ and $\delta_{b,t}$. We denote this variation as *LS5 Variation1*:

$$\begin{aligned} white_{it} - white_{i,t-1} =& \delta_{w,t} 1[white]_{m,t-1}(100 - white_{i,t-1}) \\ &+ \delta_{b,t} 1[black]_{m,t-1}(0 - white_{i,t-1}) \\ &i \in m; \quad t = 2 \cdots 16 \end{aligned} \tag{49}$$

Like LS2 variation, we again report the estimation results for both full sample and the rational sub-sample in Table 14. All $\hat{\delta}_{w,t}$ and $\hat{\delta}_{b,t}$ present intuitive positive signs. It implies that even in full sample case, on average, subjects upwards adjust if observing a white draw, downwards if a black draw. Naturally by design, the rational sub-sample case consistently presents stronger effects than the full sample case, since sub-sample cases exclude the counter-directional behaviors. On average, subjects seem to respond equally to white draws as to black draws, consistent with the finding in the original LS5. The average value of $\hat{\delta}_{w,t}$ over time, as well as the average value of $\hat{\delta}_{b,t}$ over time, is around 0.19. The values of $\hat{\delta}_{w,t}$ and $\hat{\delta}_{b,t}$ seem to be higher during the mid-period of learning. The difference between $\hat{\delta}_{w,t}$ and $\hat{\delta}_{b,t}$ are more obvious in the mid period of learning but milder at the beginning and at the end of learning. It also consolidates the finding that subjects react more actively to draw information during the mid-period, and less actively at the beginning and at the end of learning.

In addition,we consider another variation of LS5 with homogeneous and time-invariant δ_w and δ_b, denoted as *LS5 Variation2*. The belief updating dynamic then becomes:

$$\begin{aligned} white_{it} - white_{i,t-1} =& \delta_w 1[white]_{m,t-1}(100 - white_{i,t-1}) \\ &+ \delta_b 1[black]_{m,t-1}(0 - white_{i,t-1}) \end{aligned} \tag{50}$$

$$i \in m; \quad m \in \{1, \cdots 5\}; \quad t \in \{2 \cdots 16\}$$

The estimation of δ_w and δ_b can be done by simply pooling all subject-level data. The results of $\hat{\delta}_w$ and $\hat{\delta}_b$ are shown in Table 15 Column (1) and (2), without and with market fixed effects in the regression, respectively.

At last, we test whether peer effect exists. The peer effect can be described as how the belief updates of one subject affect the belief updates of other subjects in the same market. Since there is no direct interaction across subjects when the *guess games* are played, we proxy the beliefs of others within the market as follows: one part of the experiment is called *asset trading*, where subjects trade assets in a double auction within market. The dividend of the assets are ambiguous: the draws from the ambiguous urn, designed to permit learning in the *guess games*, is also used to determine the dividend of the trading assets: each time a white draw is observed, each unit of asset pays out a positive dividend; Each time a black draw, zero dividend. The asset trading takes place for 15 periods (1+1+1 pack: one *guess games* + one *asset trading* + one draw implementation, repeats 15 times, see Table 1). The traded price reflects the belief of the traders: if one believes that the proportion of the white balls is high, she is likely to trade at a higher price. Hence we proxy beliefs of other subjects in market m in Period t by the last traded price of market m in Period t, denoted by $LTP_{m,t}$. Therefore, the last variation of LS5, denoted as *LS5 Variation3*, reads:

$$\begin{aligned} white_{it} - white_{i,t-1} = & \delta_w 1[white]_{m,t-1}(100 - white_{i,t-1}) \\ & + \delta_b 1[black]_{m,t-1}(0 - white_{i,t-1}) \\ & + \delta_p LTP_{m,t-1} \end{aligned} \tag{51}$$

$$i \in m; \quad m \in \{1, \cdots 5\}; \quad t \in \{2 \cdots 16\}$$

$LTP_{m,t-1}$ denotes the last traded price in market m in period $t-1$. The estimation results of δ_w and δ_b are shown in Table 15 Column (3) and (4), without and with market fixed effects, respectively. The parameters of interest in Table 15 Column (1)-(4) are all positively significant, consistent with the finding in the original LS5. The adjustment is around 0.1 to 0.2 proportional to the distance reference. This value is slightly higher than the results from original LS5 estimation reported in Table 7 LS5 columns and Table 8 Panel D. Given the fact that the heterogeneous $\hat{\hat{\delta}}_{i,w}$ and $\hat{\hat{\delta}}_{i,b}$ are more sensitive to the irrational individual behaviors than homogeneous $\hat{\hat{\delta}}_w$ and $\hat{\hat{\delta}}_b$ (which pool all individual data, and the effect of the irrational data is averaged out), the slight difference is understandable.

9 Conclusion

This paper manages to distinguish beliefs from attitudes in situations involving ambiguity. The initial belief formation and belief updating process are directly tracked down by a simple and clear-cut experiment design. The results show that, regarding the initial prior formation, more than one third of the subjects tend to form a uniformly distributed initial prior, and nearly half of the subjects tend to constrain their attention to the mode of the distribution, rather than base their decisions on an entire prior distribution. Among the rest of the subjects, one subject forms a strictly increasing initial prior distribution, three form strictly decreasing initial prior distributions, and two form bell-shaped initial prior distributions. However, all these six distributions tend to avoid allocating large probability to some specific scenarios, and thus they all result in rather flat PDF curves. As for the belief updating strategy, subjects are mostly non-Bayesian players: more than one third (12 out of 41) of the subjects significantly deviate from the perfect Bayesian updates. Among them, 11 subjects underreact to the Bayesian updates and one subject overreacts. Half the subjects (20 out of 41) employ the heuristic updating strategy, irrelevant to the Bayes rule. Among the rest of the subjects, four tend to perfectly follow the Bayesian updates, four subjects adopt some learning strategies which are not well captured by any of the five proposed strategies, and one subject never updates the belief. As for pessimism/optimism, subjects on average demonstrate slight pessimism in a highly ambiguous environment. When the environment becomes less ambiguous due to learning, such pessimism disappears.

References

Ahn, D. S. (2008). Ambiguity without a state space. *The Review of Economic Studies*, 75(1):3–28.

Becker, G., DeGroot, M., and Marschak, J. (1964). Measuring utility by a single-response sequential method. *Behavioral Science*, 9:226–232.

Branger, N., Larsen, L. S., and Munk, C. (2013). Robust portfolio choice with ambiguity and learning about return predictability. *Journal of Banking & Finance*, 37(5):1397–1411.

Brennan, M. J. (1998). The Role of Learning in Dynamic Portfolio Decisions *. *Review of Finance*, 1(3):295–306.

Buser, T., Gerhards, L., and van der Weele, J. (2018). Responsiveness to feedback as a personal trait. *Journal of Risk and Uncertainty*, 56(2):165–192.

Campanale, C. (2011). Learning, ambiguity and life-cycle portfolio allocation. *Review of Economic Dynamics*, 14(2):339 – 367.

Cao, H. H., Wang, T., and Zhang, H. H. (2005). Model Uncertainty, Limited Market Participation, and Asset Prices. *The Review of Financial Studies*, 18(4):1219–1251.

Chen, Z. and Epstein, L. (2002). Ambiguity, risk, and asset returns in continuous time. *Econometrica*, 70(4):1403–1443.

Chew, S. H., Miao, B., and Zhong, S. (2017). Partial ambiguity. *Econometrica*, 85(4):1239–1260.

Choquet, G. (1954). Theory of capacities. *Ann. Inst. Fourier*, 5:131–295.

Diaconis, P. and Ylvisaker, D. (1979). Conjugate priors for exponential families. *The Annals of Statistics*, 7(2):269–281.

Dominitz, J. and Hung, A. A. (2009). Empirical models of discrete choice and belief updating in observational learning experiments. *Journal of Economic Behavior & Organization*, 69(2):94 – 109.

Doob, J. L. (1949). *Application of the theory of martingales*, chapter 13, pages 23–27. Colloq. Intern. du C.N.R.S (Paris).

Epstein, L. (2006). An axiomatic model of non-bayesian updating. *Review of Economic Studies*, 73(2):413–436.

Epstein, L., Noor, J., and Sandroni, A. (2008). Non-bayesian updating: A theoretical framework. *Theoretical Economics*, 3(2).

Epstein, L. G. (1999). A definition of uncertainty aversion. *The Review of Economic Studies*, 66(3):579–608.

Epstein, L. G., Noor, J., and Sandroni, A. (2010). Non-Bayesian Learning. *The B.E. Journal of Theoretical Economics*, 10(1):1–20.

Epstein, L. G. and Schneider, M. (2007). Learning under ambiguity. *The Review of Economic Studies*, 74(4):1275–1303.

Filippis, R. D., Guarino, A., Jehiel, P., and Kitagawa, T. (2017). Updating ambiguous beliefs in a social learning experiment. CeMMAP working papers CWP13/17, Centre for Microdata Methods and Practice, Institute for Fiscal Studies.

Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2):171–178.

Galaabaatar, T. and Karni, E. (2013). Subjective expected utility with incomplete preferences. *Econometrica*, 81(1):255–284.

Gelman, A., Carlin, J. B., Stern, H. S., and Rubin, D. B. (2004). *Bayesian Data Analysis*. Chapman and Hall/CRC, 2nd ed. edition.

Ghirardato, P., Maccheroni, F., and Marinacci, M. (2004). Differentiating ambiguity and ambiguity attitude. *Journal of Economic Theory*, 118(2):133–173.

Gilboa, I., Maccheroni, F., Marinacci, M., and Schmeidler, D. (2008a). Objective and Subjective Rationality in a Multiple Prior Model. Carlo Alberto Notebooks 73, Collegio Carlo Alberto.

Gilboa, I., Postlewaite, A. W., and Schmeidler, D. (2008b). Probability and Uncertainty in Economic Modeling. *Journal of Economic Perspectives*, 22(3):173–188.

Gilboa, I. and Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153.

Gilboa, I. and Schmeidler, D. (1993). Updating ambiguous beliefs. *Journal of Economic Theory*, 59(1):33–49.

Giraud, R. and Thomas, L. (2017). Ambiguity, optimism, and pessimism in adverse selection models. *Journal of Economic Theory*, 171:64 – 100.

Hanany, E. and Klibanoff, P. (2007). Updating preferences with multiple priors. *Theoretical Economics*, 2(3).

Karni, E. (2018). A mechanism for eliciting second-order beliefs and the inclination to choose. *American Economic Journal: Microeconomics*, 10(2):275–85.

Klibanoff, P., Marinacci, M., and Mukerji, S. (2005). A smooth model of decision making under ambiguity. *Econometrica*, 73(6):1849–1892.

Knight, F. H. (1921). *Risk, Uncertainty and Profit*. Houghton Mifflin Co, Boston, MA.

Lindley, D. (1991). *Making Decisions*, page 104. Wiley.

Marinacci, M. (2002). Learning from ambiguous urns. *Statistical Papers*, 43(1):143–151.

Peijnenburg, K. (2014). Life-Cycle Asset Allocation with Ambiguity Aversion and Learning. 2014 Meeting Papers 967, Society for Economic Dynamics.

Pires, C. (2002). A rule for updating ambiguous beliefs. *Theory and Decision*, 53(2):137–152.

Savage, L. J. (1954). *The foundations of statistics.* Wiley Company, New York.

Schlaifer, R. and Raiffa, H. (1961). Applied statistical decision theory. *Journal of the American Statistical Association*, 57.

Schmeidler, D. (1989). Subjective Probability and Expected Utility without Additivity. *Econometrica*, 57(3):571–587.

Table 1: Procedure of the experiment

	The 15 periods				
$t = 0$	$t = 1$	$t = 2$	$\cdots$	$t = 15$	$t = 16$
choice games	**Guess game** ↓ Asset trading ↓ **Draw implementation**	**Guess game** ↓ Asset trading ↓ **Draw implementation**	$\cdots$ $\cdots$ $\cdots$	**Guess game** ↓ Asset trading ↓ **Draw implementation**	**Guess game** ↓ choice games Questionnaire Payment
	Draw implementation: Each market implemets its own draw. In a draw, one ball is drawn out from the completely ambiguous urn with replacement. The color is displayed to the relevant subjects and then the ball is put back to the urn. The whole process is computerized				

Table 2: Subject group and *guess game* played

Session	Market	No. of the subjects	Guess games played
I	1	13	t=1-15
II	2	7	t=1-16
II	3	7	t=1-16
III	4	7	t=1-16
III	5	7	t=1-16

Table 3: Short summary of the five learning strategies

This table briefly summarizes the characteristics of the five learning strategies, in terms of the assumptions on initial prior and updating rule.

	Initial prior	**Updating rule**	**Bayesian or not**
LS1: Bayesian updates with uniformly distributed initial prior	uniform distr.	Bayes rule	Bayesian
LS2: Imperfect Bayesian updates with uniformly distributed initial prior	The Bayesian reference term: uniform distri. (as in LS1)	weight between current belief and Bayesian reference (in LS1)	non-Bayesian
LS3: Bayesian updates with *beta*-distributed initial prior	*beta*-distri.	Bayes rule	Bayesian
LS4: Imperfect Bayesian updates with *beta*-distributed initial prior	The Bayesian reference term: *beta*-distri. (as in LS3)	weight between current belief and Bayesian reference (in LS3)	non-Bayesian
LS5: Heuristic updates	Mode of the distribution	heuristic adjustment responding to new info	non-Bayesian

Table 4: Summary of the five learning strategies and their variations

Learning strategy (LS)	Initial Prior	Updating rule	Bayesian or not	Remarks
LS1: Bayesian updates with uniformly distributed initial prior				
LS1	Uniformly distribution prior	Bayes rule	Bayesian	Belief (prior/posterior) is a distribution
LS1': Maximum likelihood (ML) updates				
LS1'		Belief updates strictly follow the ML method	Non-Bayesian	Belief is single-valued, equal to the mode value of the belief distribution in LS1.
LS2: Imperfect Bayesian updates with uniformly distributed initial prior				
LS2	Bayesian updates component assumes uniformly distributed prior	Current belief is updated partially/fully/overly to the Bayesian reference (in LS1).	Non-Bayesian	The Bayesian reference refers to the mode value of one posterior distribution (as in LS1).
LS2 Variation	same as above	same as above	Non-Bayesian	The weight attached to the Bayesian reference (partially/fully/overly) is time-variant but homogeneous.
LS3: Bayesian updates with *beta*-distributed initial prior				
LS3	Prior is characterized by a *beta*-distribution with shape parameter (α_i, β_i)	Belief updates strictly follow the Bayes rule. Posterior sustains the *beta*-distribution fashion	Bayesian	Differing from LS1 only in prior assumption. The shape parameter is subject-specific.

LS4: Imperfect Bayesian updates with *beta*-distributed initial prior

LS4	Prior is characterized by a *beta*-distribution with shape parameter bundle (α_i^c, β_i^c)	Current belief is updated partially/fully/overly to the Bayesian reference with a *beta*-distributed prior (in LS3). In practice, it means that belief is updated by weighing between current belief and the Bayesian reference.	Non-Bayesian	The Bayesian reference refers to the unique peak point of one posterior distribution. The weight parameter is subject-specific but time-invariant.

LS5: Heuristic updates

LS5		Belief is upwards (downwards) adjusted if a white (black) draw is observed.	Non-Bayesian	Belief is single-valued. Adjustment scale is heterogeneous and time-invariant.
LS5 Variation1		same as above	Non-Bayesian	Adjustment scale is homogeneous and time-variant.
LS5 Variation2		same as above	Non-Bayesian	Adjustment scale is homogeneous and time-invariant.
LS5 Variation3		same as above *plus* controlling market-wide belief	Non-Bayesian	Adjustment scale is homogeneous and time-invariant.

Table 5: Market-average *guess game* responses

This table reports the market-average guess game responses of subjects for each market m: $white_{mt} = \frac{1}{N_m}\sum_i white_{it}$, where $i \in m$ and N_m denotes the number of the subjects in market m. Each market implements its own draw, and thus the draw history is market-specific. The results of $white_{mt}$ are reported in mean columns. sd. stands for the corresponding standard deviation. The number of white draws (out of 15 draws) are presented for each market. The fraction of white draws is computed by dividing the number of white draws by 15.

	Market 1		**Market 2**		**Market 3**		**Market 4**		**Market 5**	
Period	mean	sd.	mean	sd.	mean	sd.	mean	sd.	mean	sd.
1	47.54	8.53	46.86	9.36	53.86	10.59	49.71	9.21	45.71	11.94
2	45.92	13.67	44.57	14.26	48.14	18.38	49.57	10.98	48.29	15.10
3	42.69	13.37	50.86	10.83	58.29	16.18	45.86	6.56	63.29	9.28
4	42.38	13.73	57.00	9.35	49.86	13.89	41.43	11.11	69.14	17.32
5	46.92	10.67	48.43	5.21	46.00	21.39	34.57	11.39	52.00	29.94
6	42.38	11.12	64.43	16.88	52.86	18.13	21.00	10.77	71.29	14.49
7	47.46	9.52	54.86	23.55	42.43	19.14	33.43	16.03	48.71	21.00
8	52.08	6.39	55.43	12.07	40.86	26.54	38.14	10.29	62.29	13.32
9	53.77	6.83	52.00	17.84	25.86	20.57	30.00	13.89	57.29	9.59
10	54.31	7.03	45.71	13.57	37.43	31.54	40.71	8.07	64.43	8.35
11	49.92	8.15	56.71	11.41	20.86	21.20	45.14	8.29	54.29	12.26
12	50.38	4.75	51.43	7.98	39.43	18.74	40.14	9.16	51.14	5.14
13	53.85	4.65	44.57	9.59	39.86	28.63	43.29	8.03	52.14	4.52
14	52.85	4.38	49.14	12.86	33.86	23.84	48.71	4.30	49.57	11.07
15	51.23	2.08	53.71	14.80	38.57	30.43	43.43	6.88	44.14	18.53
16	-	-	46.71	4.33	32.71	23.96	42.86	6.20	49.43	7.87
No. of white draws	7		7		3		6		7	
Fraction of white draws	46.7%		46.7%		20%		40%		46.7%	

Table 6: Rationality of subjects' belief updates

This table reports the rate of rational belief updates by subject. A rational belief update is defined as such: if a white draw is observed, a subject upwards adjusts her belief or maintains her current belief; if a black draw is observed, a subject downwards adjusts her belief or maintains her current belief. Subjects' beliefs are those reported in the guess games ($white_{it}$). The rational updates rate of a subject is computed by dividing the number of her rational belief updates by the total number of her belief updates. The market average value and the full sample average value are reported in the last two rows, respectively.

Market 1		**Market 2**		**Market 3**		**Market 4**		**Market 5**	
Subject No.	Rational rate	Subject No.	Rational rate	Subject No.	Rational rate	Subject No.	Rational rate	Subject No.	Rational rate
1	0.50	14	1.00	21	0.67	28	0.93	35	1.00
2	0.93	15	0.47	22	1.00	29	0.93	36	0.87
3	0.50	16	1.00	23	0.67	30	1.00	37	0.67
4	0.64	17	0.73	24	0.93	31	0.60	38	0.87
5	0.79	18	1.00	25	0.47	32	0.87	39	0.27
6	0.64	19	0.93	26	0.60	33	0.87	40	0.80
7	0.64	20	0.93	27	0.80	34	0.93	41	0.73
8	0.79								
9	1.00								
10	1.00								
11	0.86								
12	0.71								
13	0.86								
Market avg.	0.76		0.87		0.73		0.88		0.74
Total avg. ($N = 41$)									0.79

Table 7: Parameterization of LS2-LS5

This table displays the parameter estimation of LS2-LS5. LS2 assumes that a subject updates her belief by weighting between the Bayesian reference (Bayesian updates with uniformly distributed initial prior) and her current belief. The weight assigned to Bayesian reference $\hat{\gamma}_i$ is estimated based on equation (15). LS3 assumes that a subject starts with a *beta*-distributed prior and employs the Bayes rule to update her belief. $\hat{\alpha}_i$ and $\hat{\beta}_i$ is the estimated parameter which characterizes the shape of the initial prior. The estimation is based on equation (27)-(29). LS4 assumes that a subject update her belief by weighting between the Bayesian reference (Bayesian updates with *beta*-distributed initial prior) and her current belief. The estimation is based on equation (32). $\hat{\gamma}_i^c$ is the weight assigned to the Bayesian reference. $\hat{\alpha}_i^c$ and $\hat{\beta}_i^c$ are the shape parameters characterizing the initial prior. LS5 assumes that a subject upwards (downwards) adjusts her belief if a white (black) draw is observed, proportional to the maximum possible adjustment along the upward (downward) direction. The estimation is based on equation (33). $\delta_{i,w}$ governs the adjustment reacting to white draws, while $\delta_{i,b}$ to black draws. Mean value in the last row is the mean value of the corresponding column (mean value across subjects). The parameters of Subject 22 cannot be estimated for LS4 since he/she never updates the belief. $\hat{\gamma}_i$ of LS2 and $\hat{\gamma}_i^c$ of LS4 are tested for whether significantly different from one; $\hat{\delta}_{i,w}$ and $\hat{\delta}_{i,b}$ of LS5 are tested for whether significantly different from zero

			LS2	LS3		LS4			LS5	
Session	**Market** (m)	**Subject** (i)	$\hat{\gamma}_i$	$\hat{\alpha}_i$	$\hat{\beta}_i$	$\hat{\gamma}_i^c$	$\hat{\alpha}_i^c$	$\hat{\beta}_i^c$	$\hat{\delta}_{i,w}$	$\hat{\delta}_{i,b}$
0	1	1	0.509	1.000	1.088	0.962	1.000	1.000	-0.060	-0.103
		2	0.037***	1.000	1.000	0.190***	1.000	1.000	0.005	0.028
		3	0.391**	1.000	1.000	0.595*	1.000	1.000	0.030	-0.044
		4	0.188***	1.000	1.000	1.138	1.000	1.000	-0.053	-0.095
		5	0.320***	1.000	2.140	0.269***	1.000	1.000	0.055	0.071
		6	0.083***	1.000	1.000	0.165***	1.000	1.000	0.007	0.003
		7	0.251***	1.242	1.000	0.907	1.245	1.000	0.088	0.050
		8	0.545***	1.000	1.000	0.863	1.000	1.000	0.173***	0.127**
		9	0.277***	1.077	1.000	0.946	1.078	1.000	0.135***	0.123***
		10	0.605***	1.075	1.131	0.941	1.031	1.081	0.165***	0.167***
		11	0.354***	1.000	1.000	0.718**	1.000	1.000	0.120**	0.138***
		12	0.147***	1.839	2.330	0.462***	1.000	1.272	0.040*	0.031
		13	0.373***	1.603	1.000	0.647**	1.606	1.000	0.171**	0.090
1	2	14	0.326***	1.942	1.962	0.964	1.960	1.965	0.127***	0.101***
		15	0.764	1.000	1.000	0.933	1.000	1.000	0.107	0.069
		16	0.354***	1.772	2.144	0.118***	1.000	1.000	0.062	0.050
		17	0.310***	1.000	1.068	0.592**	1.000	1.089	0.069	0.072*
		18	0.320***	1.719	2.163	0.982	1.954	2.390	0.106**	0.056
		19	0.275***	1.000	1.113	0.690**	1.000	1.123	0.105***	0.081**
		20	0.692	1.000	1.000	1.274	1.000	1.000	0.514***	0.347***
	3	21	0.555***	1.000	1.508	0.338***	1.000	1.000	0.184*	0.185*
		22	0.000***	11.639	1.006	-	-	-	0.000	0.000
		23	0.095***	1.000	1.000	1.261	1.017	1.000	-0.017	0.094
		24	0.122***	1.000	1.000	0.364***	1.000	1.000	0.044	0.113***
		25	0.729	1.110	1.000	0.766	1.000	1.000	0.194	0.415
		26	0.611	1.000	1.000	0.560**	1.000	1.000	0.293	0.194
		27	0.316***	1.000	1.000	0.573**	1.000	1.000	0.126	0.204*
2	4	28	0.730*	1.000	1.000	1.153	1.000	1.000	0.203***	0.282***
		29	0.104***	1.410	2.291	0.699***	1.000	1.430	0.085***	0.097**
		30	0.128***	1.101	1.471	1.008	1.106	1.502	0.101***	0.116***
		31	0.133***	2.214	1.000	0.425**	1.000	1.000	0.058	0.072
		32	0.200***	1.981	1.000	0.557**	1.000	1.000	0.125*	0.116
		33	0.283***	1.169	1.000	0.581***	1.000	1.000	0.137**	0.126*
		34	0.301***	1.410	1.000	0.334***	1.000	1.000	0.000	0.100
	5	35	0.570**	1.000	1.465	1.159	1.000	1.515	0.296**	0.254***
		36	0.323***	1.000	1.195	0.888	1.000	1.217	0.212***	0.110***
		37	0.046***	1.000	1.237	0.202***	1.000	1.000	0.077	0.025
		38	0.268***	1.000	1.820	1.066	1.000	1.879	0.163**	0.154***
		39	0.568	1.000	1.000	0.924	1.000	1.000	0.051	-0.124
		40	0.877	1.000	1.641	1.389**	1.000	2.401	0.586***	0.510***
		41	0.346***	1.000	1.000	0.660**	1.000	1.000	0.295**	0.125
Mean value			0.352	1.446	1.263	0.732	1.075	1.172	0.126	0.110

For $\hat{\gamma}_i$ of LS2 ($\hat{\gamma}_i^c$ of LS4), t-test: $H_0 : \hat{\gamma}_i(\hat{\gamma}_i^c) = 1$. For $\hat{\delta}_{i,w}$ ($\hat{\delta}_{i,b}$) of LS5, t-test: $H_0 : \hat{\delta}_{i,w}(\hat{\delta}_{i,b}) = 0$. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 8: Statistical summary of estimated parameters in LS2-LS5

This table displays the across-subject statistical summary of each estimated parameters in Table 7. sd. denotes standard deviation.

Panel A: $\hat{\gamma}_i$ in Table 7 column LS2			
mean	variance	sd.	N
0.352	0.049	0.222	41
Percentiles		Distribution	
25%	0.188	$[0, 0.3)$	41.5%
50%	0.319	$[0.3, 0.6)$	41.5%
75%	0.545	$[0.6, 1]$	17.0%
90%	0.692		

Panel B: $\hat{\alpha}_i$ and $\hat{\beta}_i$ in Table 7 column LS3		
	count	percent
$\hat{\alpha}_i = \hat{\beta}_i = 1$	15	37%
$\hat{\alpha}_i = 1$, $\hat{\beta}_i > 1$	10	24%
$\hat{\alpha}_i > 1$, $\hat{\beta}_i = 1$	8	20%
$\hat{\alpha}_i > \hat{\beta}_i > 1$	1	2%
$\hat{\beta}_i > \hat{\alpha}_i > 1$	7	17%
Total	41	100%

Panel C: $\hat{\gamma}_i^c$, $\hat{\alpha}_i^c$ and $\hat{\beta}_i^c$ in Table 7 column LS4				
$\hat{\gamma}_i^c$:	mean	variance	sd.	N
	0.732	0.113	0.336	40
	Percentiles		Distribution	
	25%	0.510	$[0, 0.3)$	12.5%
	50%	0.709	$[0.3, 0.6)$	27.5%
	75%	0.963	$[0.6, 1]$	40.0%
	90%	1.156	> 1	20.0%
		count	percent	
	$\hat{\alpha}_i^c = \hat{\beta}_i^c = 1$	24	60%	
	$\hat{\alpha}_i^c = 1$, $\hat{\beta}_i^c > 1$	8	20%	
	$\hat{\alpha}_i^c > 1$, $\hat{\beta}_i^c = 1$	4	10%	
	$\hat{\alpha}_i^c > \hat{\beta}_i^c > 1$	0	0%	
	$\hat{\beta}_i^c > \hat{\alpha}_i^c > 1$	4	10%	
	Total	40	100%	

Panel D: $\hat{\delta}_{i,w}$ and $\hat{\delta}_{i,b}$ in Table 7 column LS5		
	$\hat{\delta}_{i,w}$	$\hat{\delta}_{i,b}$
mean	0.126	0.110
sd.	0.130	0.124
Percentiles		
25%	0.051	0.050
50%	0.106	0.100
75%	0.171	0.138
90%	0.293	0.254
N	41	41

Table 9: Goodness of fit: LS1-LS5

This table reports the differences (differences are expressed by the absolute values) between the factual belief updates (the guess game entry $white_{it}$) and the belief updates simulated by one of the five learning strategies (LS1-LS5). It checks the goodness of fit for each learning strategy. Five learning strategies are developed to capture subjects' belief update dynamics. LS1 assumes that a subject starts with a uniformly distributed prior and employs the Bayes rule to update her belief. LS2 assumes that she weights between LS1 and her previous belief. LS3 assumes that she starts with a *beta*-distributed prior and employs the Bayes rule. LS4 assumes that she weights between LS3 update and her previous belief. LS5 assumes that she responds to the white (black) draw by upwards (downwards) adjusting her belief proportional to the maximum possible adjustment in the upward (downward) direction. The diff columns report the average difference over all subjects by LS by period, computed based on equation (34)-(38), respectively. sd. denotes the corresponding standard deviation.

	diff(guess, LS1)$_t$		diff(guess, LS2)$_t$		diff(guess, LS3)$_t$		diff(guess, LS4)$_t$		diff(guess, LS5)$_t$		
Period	diff	sd.	diff	sd.	diff	sd.	diff	sd.	diff	sd.	N
1					37.34	18.15					41
2	48.46	14.84	17.74	17.44	38.85	22.65	29.70	22.23	11.06	11.47	41
3	13.49	14.95	8.97	12.56	12.03	11.72	8.37	10.73	11.38	16.42	41
4	16.56	13.25	10.50	10.94	14.91	12.72	11.66	11.74	10.70	11.75	41
5	12.71	15.64	11.82	15.57	11.45	13.90	9.29	12.02	10.35	11.82	41
6	11.37	10.81	11.50	12.16	9.87	10.95	8.59	11.28	11.47	16.31	41
7	12.90	15.43	12.40	14.87	10.91	14.59	9.43	13.02	10.64	17.33	41
8	10.17	12.74	8.02	10.56	8.18	11.41	8.36	11.95	9.03	13.34	41
9	9.10	10.98	9.21	9.87	7.71	8.61	6.73	8.66	9.00	11.33	41
10	9.46	14.27	8.35	12.78	7.67	12.47	7.52	13.19	7.34	15.19	41
11	8.00	10.55	8.12	9.06	6.71	7.73	6.31	8.27	7.48	10.31	41
12	8.00	8.58	6.01	7.75	5.91	6.51	5.34	7.93	7.73	8.98	41
13	8.24	12.58	5.64	9.55	7.29	10.87	5.45	9.61	6.20	9.28	41
14	7.88	11.02	6.63	8.07	6.54	9.31	6.15	8.18	6.10	7.75	41
15	8.95	15.65	8.44	15.19	8.34	14.11	7.14	14.02	10.77	18.79	41
16	8.07	12.33	5.84	7.12	6.95	9.30	5.82	7.55	9.97	10.14	28

Table 10: Speed of convergence of LS1-LS5

This table displays how fast one belief update process converges to a constant number. The five belief updates generated by five different learning strategies are considered, respectively. For a given learning strategy (LS1-LS5), the belief adjustments between every two neighboring periods of subject i are first computed based on equation (39), and then averaged over all subjects. The results are reported in $\overline{\text{Diff}}(LS1)_t$-$\overline{\text{Diff}}(LS5)_t$ column as a reference for speed of convergence of LS1-LS5, respectively. The corresponding standard deviations are reported in sd. columns. As a reference, the belief adjustment based on factual guess game data is computed by equation (41)-(42) and shown in Factual belief columns.

	LS1		**LS2**		**LS3**		**LS4**		**LS5**		**Factual belief**		
Period	$\overline{\text{Diff}}(LS1)_t$	sd.	$\overline{\text{Diff}}(LS2)_t$	sd.	$\overline{\text{Diff}}(LS3)_t$	sd.	$\overline{\text{Diff}}(LS4)_t$	sd.	$\overline{\text{Diff}}(LS5)_t$	sd.	$\overline{\text{Diff}}(guess)_t$	sd.	N
1													
2					31.45	32.73					11.41	12.27	41
3	41.46	18.81	14.92	12.68	33.67	16.80	25.74	17.56	13.08	12.33	13.32	15.61	41
4	14.10	6.40	9.27	9.44	12.86	5.52	10.89	7.20	12.95	16.09	12.49	11.23	41
5	16.83	4.97	8.13	8.04	14.70	4.65	12.10	7.81	12.30	14.14	14.44	17.74	41
6	8.29	2.37	7.16	7.01	7.96	2.12	7.97	7.21	9.55	9.59	17.34	20.81	41
7	10.00	2.48	7.38	7.37	9.10	2.42	8.65	8.40	10.63	11.74	15.54	19.64	41
8	5.98	2.26	8.30	6.34	6.05	1.53	7.06	6.38	11.97	17.24	11.05	13.34	41
9	6.32	1.81	5.65	6.19	5.88	1.83	6.68	6.40	9.62	10.84	12.46	16.28	41
10	5.66	1.51	5.95	5.71	5.05	1.45	5.23	4.18	11.71	12.30	9.54	14.60	41
11	5.66	1.72	4.75	5.72	5.07	1.37	4.73	3.82	11.10	12.39	10.39	16.47	41
12	5.17	0.91	5.24	5.36	4.72	1.22	3.74	3.24	9.70	8.84	8.49	11.71	41
13	4.02	1.42	4.47	4.05	3.86	0.84	4.35	3.59	8.49	8.45	5.63	7.04	41
14	3.66	0.75	2.91	2.46	3.68	0.85	3.24	2.46	6.56	5.25	7.05	12.23	41
15	3.49	0.77	3.91	4.12	3.30	0.77	3.83	3.26	8.89	11.09	10.22	18.39	41
16	2.50	0.87	6.37	7.73	2.66	0.70	3.83	3.87	14.41	19.88	8.46	11.58	28

Table 11: Ranking of LS by subject

This table displays the by-subject ranking of the five learning strategies, measured by the root mean square error (RMSE). RMSE of a given LS is used to quantify, for a given subject, the explanatory power of this LS to her factual belief updates. The RMSEs are computed based on equation (43). The lower the value of RMSE is, the better the corresponding LS explains the factual belief updates, and hence the higher the rank of the LS is. Subject 22 has no entry since he/she never updates the belief.

			Learning strategy No.				
Session	Market	Subject	rank 1^{st}	rank 2^{nd}	rank 3^{rd}	rank 4^{th}	rank 5^{th}
I	1	1	2	3	5	4	1
		2	5	2	4	3	1
		3	2	4	5	3	1
		4	5	2	3	1	4
		5	2	4	5	3	1
		6	2	4	5	3	1
		7	5	2	4	3	1
		8	5	2	4	1	3
		9	5	2	4	3	1
		10	2	5	4	3	1
		11	5	2	4	1	3
		12	2	5	3	4	1
		13	4	5	3	2	1
II	2	14	4	3	5	2	1
		15	2	4	3	1	5
		16	2	3	4	5	1
		17	2	5	4	3	1
		18	4	3	5	2	1
		19	5	2	4	3	1
		20	5	2	1	3	4
	3	21	2	4	5	3	1
		22	-	-	-	-	-
		23	5	2	1	3	4
		24	5	2	4	1	3
		25	4	2	3	1	5
		26	4	2	3	1	5
		27	5	2	4	3	1
III	4	28	5	2	3	1	4
		29	5	3	2	4	1
		30	5	2	4	3	1
		31	5	2	4	3	1
		32	5	2	4	3	1
		33	5	2	4	3	1
		34	2	4	5	1	3
	5	35	4	3	5	2	1
		36	5	2	4	3	1
		37	5	2	4	3	1
		38	4	3	5	2	1
		39	2	4	3	1	5
		40	4	3	5	2	1
		41	5	2	4	1	3

Count: LS ranked 1^{st}	LS5	LS2	LS4
	20	12	8

Table 12: Heterogeneity: initial prior and learning strategy

This table summarize the chosen learning strategy and initial prior by subject. The chosen learning strategy of subject i is defined as the learning strategy (of LS1-LS5) which produces the smallest root mean squared error (RMSE) over all belief updating periods of subject i. The error is defined as the difference between the factual belief (guess game entry) and the simulated belief generated by a given learning strategy. The RMSE is computed based on equation (43) by subject by learning strategy. Panel A reports the initial prior choice (uniformly distributed, *beta*-distributed, or a singular value) and the learning strategy choice (responsiveness to the Bayesian reference; or the scaled parameters in directional adjustment) for each subject. Panel B displays the statistical summary of the choices. Subject 22 has no entry since he/she never updates the belief.

Panel A: choice of initial prior and learning strategy

			Initial prior: uniform or *beta* distribution learning strategy: referring to the Bayesian reference					Initial prior: distribution unspecified learning strategy: heuristic strategy		
Session	**Market** (m)	**Subject** (i)	Initial prior	$\hat{\alpha}_i^c$	$\hat{\beta}_i^c$	Responsiveness to Bayesian reference	$\hat{\gamma}_i^c$	Initial prior	$\hat{\delta}_{i,w}$	$\hat{\delta}_{i,b}$
I	1	1	uniform			partially	0.509			
		2						60	0.005	0.028
		3	uniform			partially	0.391			
		4						40	-0.053	-0.095
		5	uniform			partially	0.320			
		6	uniform			partially	0.083			
		7						45	0.088	0.050
		8						40	0.173	0.127
		9						50	0.135	0.123
		10	uniform			partially	0.605			
		11						65	0.120	0.138
		12	uniform			partially	0.148			
		13	beta	1.606	1.000	partially	0.647			
II	2	14	beta	1.960	1.965	partially	0.964			
		15	uniform			partially	0.764			
		16	uniform			partially	0.354			
		17	uniform			partially	0.310			
		18	beta	1.954	2.390	partially	0.982			
		19						42	0.105	0.081
		20						63	0.514	0.347
	3	21	uniform			partially	0.555			
		22	-	-	-	-	-	-	-	-
		23						50	-0.017	0.094
		24						50	0.044	0.113
		25	uniform			partially	0.766			
		26	uniform			partially	0.560			
		27						42	0.126	0.204
III	4	28						63	0.203	0.282
		29						40	0.085	0.097
		30						35	0.101	0.116
		31						50	0.058	0.072
		32						60	0.125	0.116
		33						50	0.137	0.126
		34	uniform			partially	0.301			
	5	35	beta	1.000	1.515	overly	1.159			
		36						37	0.212	0.110
		37						50	0.076	0.025
		38	beta	1.000	1.879	overly	1.066			
		39	uniform			partially	0.568			
		40	beta	1.000	2.401	overly	1.389			
		41						42	0.295	0.125

Panel B: Summary

The best fitted LS is LS5

N=20	Initial prior $white_{i1}$ (mean)	$\hat{\delta}_{i,w}$ (mean)	$\hat{\delta}_{i,b}$ (mean)
	48.7	0.127	0.114

The best fitted LS is LS2 or LS4

N=20	Initial prior		Responsiveness to the Bayesian reference	
	$\hat{\alpha} = \hat{\beta} = 1$	14	partially ($\hat{\gamma}, \hat{\gamma}^c < 1$)	17
	$\hat{\alpha} > 1$, $\hat{\beta} = 1$	1	fully ($\hat{\gamma}, \hat{\gamma}^c = 1$)	0
	$\hat{\alpha} = 1$, $\hat{\beta} > 1$	3	overly ($\hat{\gamma}, \hat{\gamma}^c > 1$)	3
	$\hat{\alpha} > 1$, $\hat{\beta} > 1$	2	mean=0.622	

Table 13: Pessimism/optimism of subjects: in comparison with benchmark B1-B6

Panel A of this table reports the absolute difference between the factual guess game entries and six benchmark strategies (B1-B6), respectively. All benchmark strategies employ Bayes rule to update beliefs. They differ from each other only in terms of the initial prior. B1 is constructed with a uniform distribution. B2 assumes a symmetric triangle-shaped initial prior distribution (see Figure 9); B3-B6 assumes the initial priors characterized by the *beta*-distribution with $\alpha = \beta = 2, 3, 10, 50$, respectively. The belief updates of B1-B6 are computed by equation (44)-(45). The per period average absolute difference diff(guess, B1)$_t$-diff(guess, B6)$_t$ are computed by equation (46). sd.columns report the corresponding standard deviation. The absolute difference can be interpreted as the goodness of fit of B1($\cdots$ B6) when regressing the guess game responses.
Panel B of this table reports the exact differences between the factual guess game entries and the six benchmark strategies (B1-B6), respectively, computed by equation (47). sd.columns report the corresponding standard deviation. The exact difference can be interpreted as the pessimism/optimism of subjects, comparing with the chosen benchmark (B1 $\cdots$ B6). A negative value indicates, on average, a lower value of the updates (towards zero) than the benchmark value, hence implying pessimism. In contrast, a positive value indicates optimism. For each exact difference, a one-sided t-test is conducted for whether it is significantly positive or negative.

Panel A: Goodness of fit: **absolute** difference between guess games responses and benchmark B1-B6 updates

	diff(guess, B1)$_t$		diff(guess, B2)$_t$		diff(guess, B3)$_t$		diff(guess, B4)$_t$		diff(guess, B5)$_t$		diff(guess, B6)$_t$		
Period	diff	sd.	diff	sd.	diff	sd.	diff	sd.	diff	sd.	diff	sd.	N
1			7.24	7.26	7.24	7.26	7.24	7.26	7.24	7.26	7.24	7.26	41
2	48.46	14.84	12.02	8.83	17.66	12.14	13.63	10.29	12.29	8.44	12.05	8.68	41
3	13.49	14.95	8.59	10.19	9.41	10.49	8.59	10.19	9.02	10.27	9.37	10.64	41
4	16.56	13.25	11.46	9.68	11.63	9.27	10.54	8.87	11.76	9.45	12.68	10.70	41
5	12.71	15.64	11.73	13.36	11.98	13.99	11.73	13.48	11.98	13.30	12.27	13.62	41
6	11.37	10.81	12.83	10.86	11.10	10.65	11.71	10.80	14.78	11.79	17.07	12.94	41
7	12.90	15.43	12.76	13.14	12.46	14.06	12.46	13.66	13.10	13.03	13.54	13.27	41
8	10.17	12.74	10.68	10.00	9.83	11.56	9.90	10.79	11.61	9.38	12.90	9.98	41
9	9.10	10.98	9.83	10.45	9.34	10.59	9.78	10.48	11.68	11.39	13.07	13.04	41
10	9.46	14.27	10.54	12.28	9.59	13.07	9.98	12.24	11.54	11.21	13.12	11.76	41
11	8.00	10.55	8.85	9.65	8.56	9.94	9.05	9.83	10.41	10.63	11.68	12.23	41
12	8.00	8.58	6.66	7.79	7.68	7.57	7.37	7.26	7.32	6.98	7.20	8.46	41
13	8.24	12.58	8.24	11.38	8.24	11.76	8.24	11.39	8.51	10.63	8.93	10.96	41
14	7.88	11.02	7.73	10.13	7.73	10.39	7.78	10.15	8.22	9.72	8.85	10.47	41
15	8.95	15.65	9.20	14.59	9.07	14.88	9.15	14.41	9.44	13.63	9.71	13.88	41
16	8.07	12.33	8.54	10.88	8.18	11.30	8.43	10.73	9.61	9.26	10.96	10.46	28

Panel B: Pessimism/Optimism: **exact** difference between guess games responses and benchmark B1-B6 updates

	diff(guess, B1)$_t^{exact}$		diff(guess, B2)$_t^{exact}$		diff(guess, B3)$_t^{exact}$		diff(guess, B4)$_t^{exact}$		diff(guess, B5)$_t^{exact}$		diff(guess, B6)$_t^{exact}$		
Period	diff	sd.	diff	sd.	diff	sd.	diff	sd.	diff	sd.	diff	sd.	N
1			-1.44	10.16	-1.44	10.16	-1.44	10.16	-1.44	10.16	-1.44	10.16	41
2	-4.12	50.52	-2.90	14.63	-3.32	21.17	-3.15	16.79	-2.98	14.61	-2.93	14.56	41
3	-7.73***	18.59	-2.10	13.16	-3.46*	13.66	-2.10	13.16	-0.05	13.67	0.63	14.16	41
4	-5.49**	20.49	-3.71*	14.54	-3.10*	14.55	-2.00	13.63	-0.34	15.08	0.39	16.59	41
5	-4.22*	19.71	-4.22*	17.27	-4.22*	17.92	-4.22*	17.36	-4.22*	17.39	-4.22*	17.84	41
6	2.39	15.50	-0.78	16.79	1.44	15.32	1.12	15.89	-0.15	18.91	-0.78	21.41	41
7	-4.32*	19.65	-4.32*	17.80	-4.32*	18.29	-4.32*	17.98	-4.32*	17.96	-4.32*	18.46	41
8	0.27	16.30	2.15	14.48	0.07	15.17	0.05	14.64	0.15	14.92	0.27	16.31	41
9	1.73	14.16	0.02	14.34	0.37	14.12	-0.32	14.33	-2.71	16.09	-4.24*	17.97	41
10	3.27	16.81	2.78	15.94	2.56	16.00	2.02	15.66	0.61	16.07	-0.29	17.62	41
11	1.17	13.19	-0.02	13.10	0.32	13.11	-0.37	13.35	-2.07	14.73	-3.44*	16.56	41
12	0.10	11.73	1.34	10.16	-0.27	10.78	-0.29	10.34	-1.85	9.94	-2.41*	10.84	41
13	0.88	15.02	0.44	14.05	0.34	14.35	0.15	14.06	-0.95	13.58	-1.85	14.01	41
14	1.73	13.43	1.29	12.68	1.20	12.89	0.85	12.76	-0.71	12.71	-1.78	13.59	41
15	3.10	17.76	1.39	17.19	2.24	17.28	1.73	16.98	-0.32	16.58	-2.20	16.80	41
16	4.43*	14.06	0.68	13.81	3.18	13.58	2.43	13.43	-1.32	13.28	-5.82**	13.99	28

t-test: H_0 : diff(guess, .)$_t^{exact} = 0$. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 14: Robustness check: LS2 variation and LS5 variation

This table reports the estimation results for intertemporal LS2 and intertemporal LS35. LS2 assumes that a subject updates her belief by weighting between the Bayesian reference and the one-period lagged belief. The intertemporal version of LS2 is modeled such that the weight parameter γ_t is homogeneous across subjects but time-variant. The OLS estimation is based on equation (48). LS5 assumes that a subject upwards (downwards) adjusts her belief if a white (black) draw is observed. Such adjustment is proportional to the distance between her current belief and extreme scenario. The intertemporal version of LS5 is modeled such that the adjustment parameters $\delta_{w,t}$ and $\delta_{b,t}$ are homogeneous across subjects but time-variant. The OLS estimation is based on equation (49). Two regressions are run based on full sample and rational sub-sample, respectively. The rational sub-sample only include the rational belief updates. Rational belief updates, as defined, are those guess game entries $white_{it}$ such that $white_{it} \geq white_{i,t-1}$ if a white draw is observed at the end of period $t-1$, and $white_{it} \leq white_{i,t-1}$ if a black draw is observed at the end of period $t-1$. The results of $\hat{\gamma}_t$ ($\hat{\delta}_{w,t}$ and $\hat{\delta}_{b,t}$) over t are reported in LS2 columns (LS5 column).

	LS2		**LS5**			
	$\hat{\gamma}_t$	$\hat{\gamma}_t$	$\hat{\delta}_{w,t}$	$\hat{\delta}_{b,t}$	$\hat{\delta}_{w,t}$	$\hat{\delta}_{b,t}$
t=2	0.040	0.177***	0.029	0.053	0.188***	0.165***
	(0.048)	(0.043)	(0.075)	(0.083)	(0.059)	(0.061)
t=3	0.512***	0.429***	0.136**	0.0192	0.196***	0.123*
	(0.092)	(0.077)	(0.061)	(0.093)	(0.042)	(0.066)
t=4	0.328***	0.375***	0.085	0.0892	0.181***	0.213***
	(0.108)	(0.091)	(0.087)	(0.088)	(0.061)	(0.067)
t=5	0.726***	0.614***	0.117	0.220***	0.193***	0.324***
	(0.149)	(0.128)	(0.083)	(0.059)	(0.061)	(0.043)
t=6	0.980***	0.892***	0.508***	0.157**	0.509***	0.217***
	(0.106)	(0.089)	(0.088)	(0.074)	(0.058)	(0.052)
t=7	0.854***	0.617***	0.137**	0.292***	0.159***	0.351***
	(0.135)	(0.113)	(0.056)	(0.073)	(0.038)	(0.052)
t=8	0.538***	0.538***	0.184**	0.074	0.214***	0.137**
	(0.112)	(0.102)	(0.073)	(0.080)	(0.051)	(0.056)
t=9	0.880***	0.836***		0.145***		0.245***
	(0.142)	(0.121)		(0.053)		(0.039)
t=10	0.524***	0.389***	0.120*	0.092	0.157***	0.147**
	(0.160)	(0.133)	(0.064)	(0.105)	(0.045)	(0.073)
t=11	0.835***	0.828***	0.157*	0.213***	0.189***	0.258***
	(0.137)	(0.109)	(0.082)	(0.061)	(0.059)	(0.043)
t=12	0.692***	0.632***	0.129**	0.101	0.177***	0.140**
	(0.175)	(0.154)	(0.058)	(0.093)	(0.042)	(0.064)
t=13	0.023	0.221	0.066	0.028	0.078	0.115*
	(0.194)	(0.172)	(0.074)	(0.079)	(0.050)	(0.061)
t=14	0.532***	0.764***	0.098	0.095	0.141**	0.173***
	(0.152)	(0.136)	(0.084)	(0.065)	(0.063)	(0.050)
t=15	0.807***	0.360**	0.007	0.056	0.144	0.106**
	(0.173)	(0.152)	(0.130)	(0.062)	(0.098)	(0.043)
t=16						
Sample choice	full sample	rational sample	full sample		rational sample	
N	574	456	574		456	
R^2	0.384	0.486	0.192		0.52	

Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 15: Robustness check: LS5 variation

This table summaries the estimated parameters built on two variations of LS5. LS5 assumes that a subject upwards (downwards) adjusts her belief if a white (black) draw is observed. Such adjustment is proportional to the distance between one's current belief and the extreme scenario, governed by δ_w and δ_b, respectively. δ_w and δ_b are assumed to be time-invariant and homogeneous across subjects. POLS is applied to equation (50). The regression results are reported in column (1) (without market FE) and column (2) (with market FE). Column (3)(4) further control for the one-period lagged market traded price. POLS is applied to equation (51). The regression results are reported in column (3) (without market FE) and column (4) (with market FE).

Dependent variable: Belief adjustment: $white_{i,t} - white_{i,t-1}$				
	(1)	(2)	(3)	(4)
$\hat{\delta}_w$	0.136***	0.120***	0.113***	0.109**
	(0.020)	(0.026)	(0.041)	(0.048)
$\hat{\delta}_b$	0.129***	0.152***	0.231***	0.240***
	(0.019)	(0.027)	(0.046)	(0.052)
$\hat{\delta}_p$			0.083*	0.073
			(0.045)	(0.046)
market FE	No	Yes	No	Yes
N	602	602	392	392
R^2	0.135	0.138	0.169	0.171

Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure 1: Screen display of the guess game

This figure displays the computer screen a subject sees when she plays the *guess game*. A subject inserts a integer between zero and 100 (inclusive) to announce her guess about the proportion of the white balls in the ambiguous urn. Draws with replacement are implemented from the ambiguous urn as the source of new information in order to permit learning. The play goes as 1 guess game + 1 draw, and this 1+1 pack is played by 15 times. In Session II and III, one extra guess game is played after the 15^{th} draw. The previous draw history, if any, is displayed on the screen for subjects' reference. Figure (a) is the screen display of the very first guess game; Figure (b) is the screen display of the guess game when five (as an example) draws are implemented. Inactivity is not allowed. Time restriction does not apply.

(a) Guess game at t=1

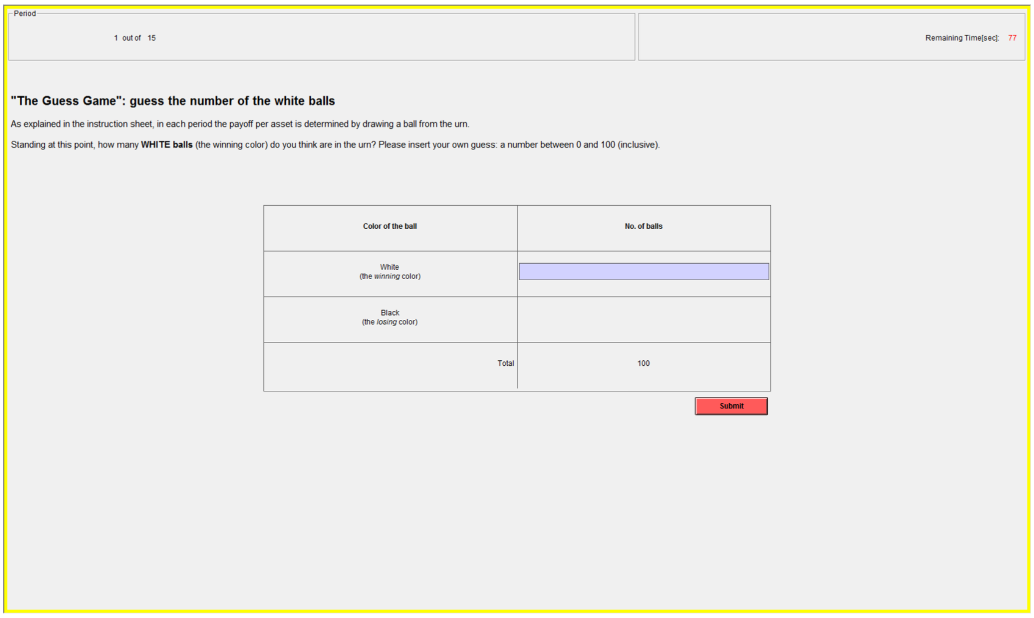

(b) Guess game at t=6

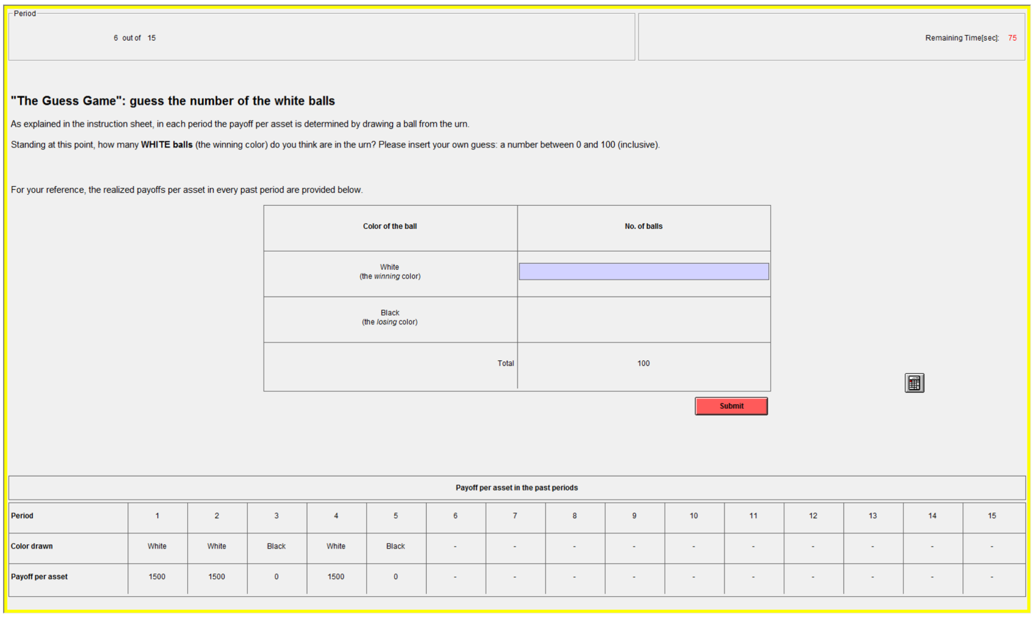

Figure 2: Guess game entry and draw history: by market

The diagram illustrates the market-average guess game entry, $white_{mt} = \frac{1}{N_m}\sum_i white_{it}$, and the market-specific draw history. $white_{it}$ is the guess game entry of subject i in period t, the subjective belief about the number of the white balls in the ambiguous urn. 15 draws are implemented in each market independently. The shaded column (located in period $t \geq 2$) represents that a white draw is observed at the end of previous period $t-1$, therefore the beliefs inside the shaded columns are reached just after the relevant subjects observe this white draw. Analogously, the non-shaded column (located in period $t \geq 2$) represents that a black draw is observed at the end of previous period $t-1$, and the beliefs inside the non-shaded columns are reached just after the relevant subjects observe this black draw. The beliefs in period 1 are reached before any draw is implemented.

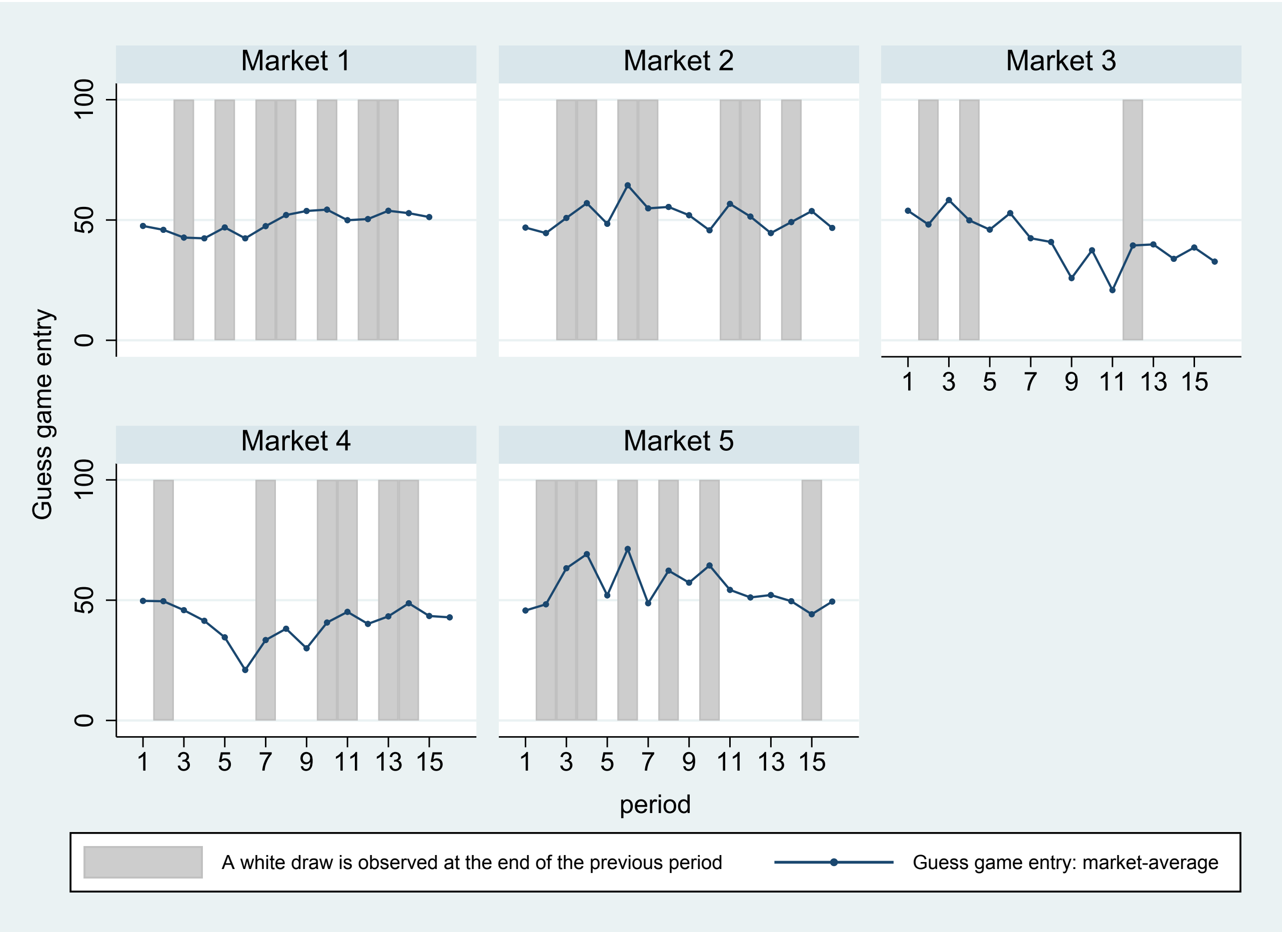

Figure 3: Examples of *beta*-distribution: PDF

This diagram illustrates the PDFs of *beta*-distributions with shape parameter bundle $(\alpha, \beta) = (2, 2); (3, 3); (1, 5); (5, 1); (3, 5); (5, 3)$, respectively. A uniform distribution is also displayed in the diagram as a reference, which is equivalent to the *beta*-distribution if $(\alpha, \beta) = (1, 1)$

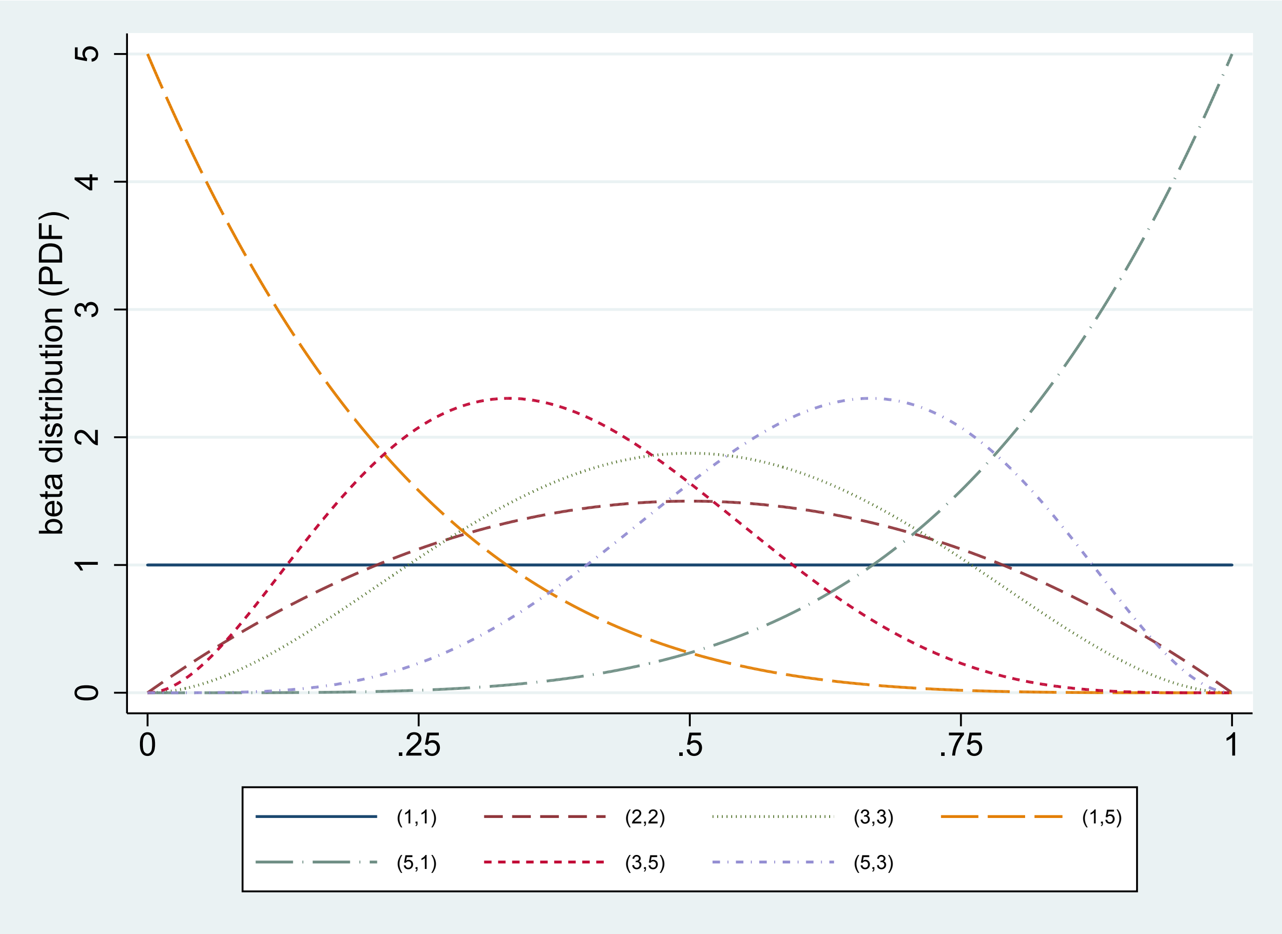

Figure 4: LS3: distribution of the initial prior

These diagrams illustrate the estimation results of the initial priors described in LS3. LS3 assumes that a subject starts with a *beta*-distributed initial prior and update it by employing the Bayes rule. The shape of the initial prior is characterized by parameter bundle (α_i, β_i). Using the guess game entries, we estimate subject-specific (α_i, β_i) bundles based on equation (27)-(29). $\hat{\alpha}_i$ and $\hat{\beta}_i$ are restricted $\hat{\alpha}_i, \hat{\beta}_i \geq 1$, . Diagram (a) illustrates the distribution of the 41 pairs of $(\hat{\alpha}_i, \hat{\beta}_i)$. The size of the bubble represents the frequency of $(\hat{\alpha}_i, \hat{\beta}_i)$ estimated at this value. Diagram (b) illustrates the PMFs of the initial prior recovered from $(\hat{\alpha}_i, \hat{\beta}_i)$ based on equation (16) and (30). Sample size $N = 41$.

(a) Distribution of $(\hat{\alpha}_i, \hat{\beta}_i)$

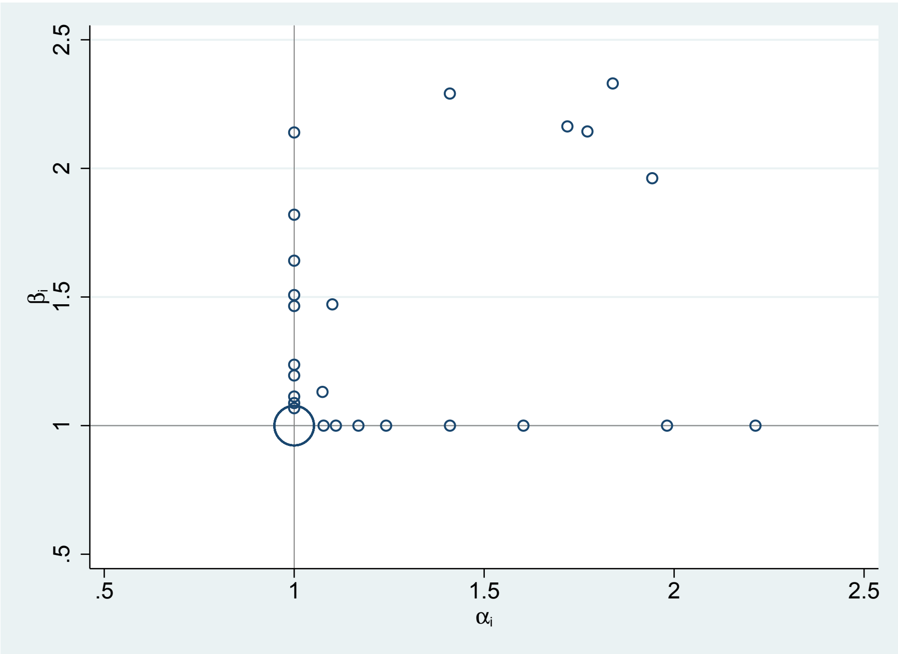

(b) PMF of the initial prior: recovered from $(\hat{\alpha}_i, \hat{\beta}_i)$

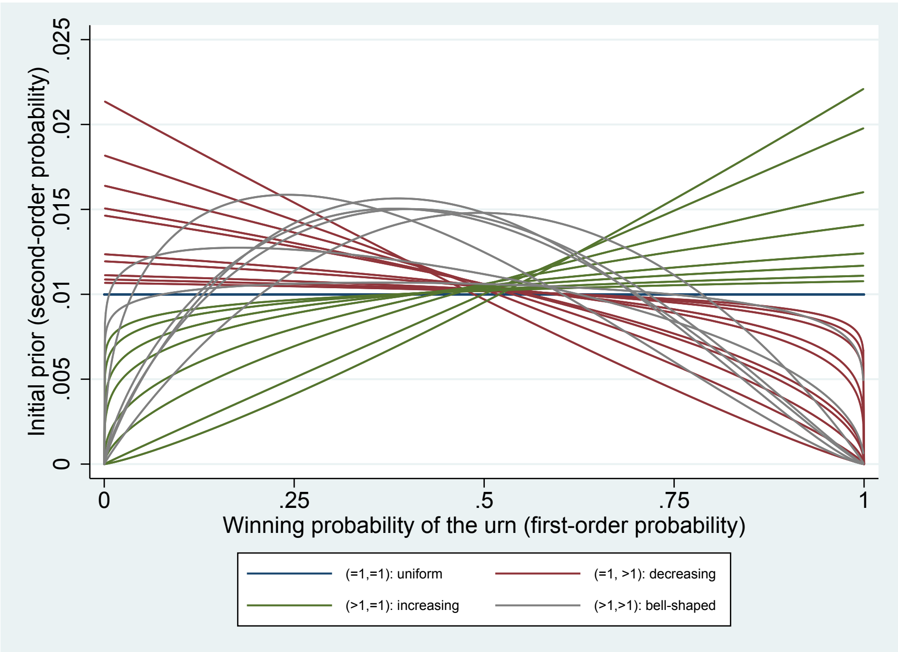

Figure 5: LS4: distribution of the initial prior

These diagrams illustrate the estimation results of the initial priors described in LS4. LS4 assumes that a subject starts with a *beta*-distributed initial prior, employs the Bayes rule to update it (obtaining the posteriors), and adjusts her belief using the posterior as a reference. The shape of the initial prior is characterized by parameter bundle (α_i^c, β_i^c). We estimate subject-specific (α_i^c, β_i^c) bundles based on equation (32), together with the responsiveness parameter γ^c. $\hat{\alpha^c}_i$ and $\hat{\beta}_i^c$ are restricted $\hat{\alpha}_i^c, \hat{\beta}_i^c \geq 1$, . Diagram (a) illustrates the distribution of the 40 pairs of $(\hat{\alpha}_i^c, \hat{\beta}_i^c)$. The size of the bubble represents the frequency of $(\hat{\alpha}_i^c, \hat{\beta}_i^c)$ estimated at this value. Diagram (b) illustrates the PMFs of the initial prior recovered from $(\hat{\alpha}_i^c, \hat{\beta}_i^c)$ using equation (16) and (30). Sample size $N = 40$. One sample drops out since she never updates her belief and thus (α_i^c, β_i^c) cannot be estimated.

(a) Distribution of $(\hat{\alpha}_i^c, \hat{\beta}_i^c)$

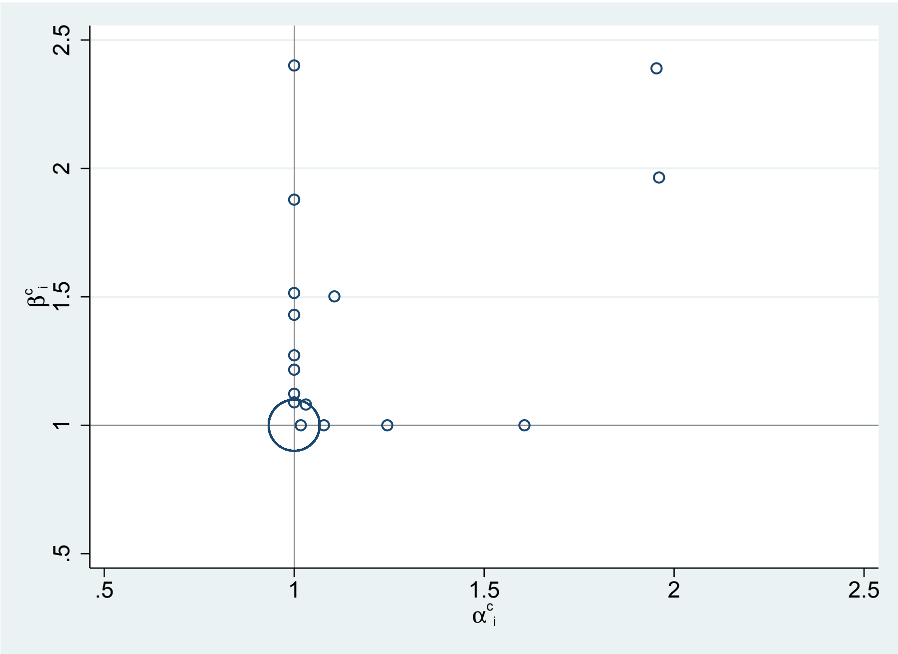

(b) PMF of the initial prior: recovered from $(\hat{\alpha}_i^c, \hat{\beta}_i^c)$

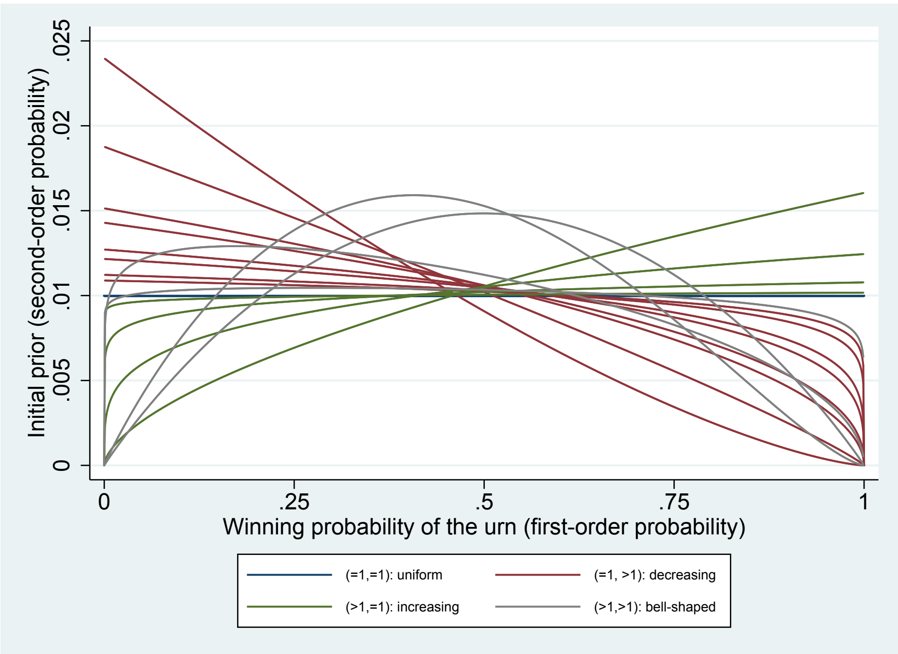

Figure 6: Goodness of fit: LS1-LS5

This diagram illustrates the absolute differences between the factual belief updates (the guess game entry $white_{it}$) and the belief updates simulated by one of the five learning strategies (LS1-LS5). It checks the goodness of fit of LS1-LS5. Five learning strategies are developed to capture subjects' belief update dynamics. LS2 assumes that a subject updates her belief by weighting between the Bayesian reference (Bayesian updates with uniformly distributed initial prior) and her current belief. LS3 assumes that a subject starts with a *beta*-distributed prior and employs the Bayes rule to update her belief. LS4 assumes that a subject update her belief by weighting between the Bayesian reference (Bayesian updates with *beta*-distributed initial prior) and her current belief. LS5 assumes that a subject upwards (downwards) adjusts her belief if a white (black) draw is observed, proportional to the maximum possible adjustment along the upward (downward) direction. The absolute difference diff(guess, LS1)$_t$-diff(guess, LS5)$_t$ are computed by equation (34)-(38)). These results corresponds with the *diff* columns in Table 9.

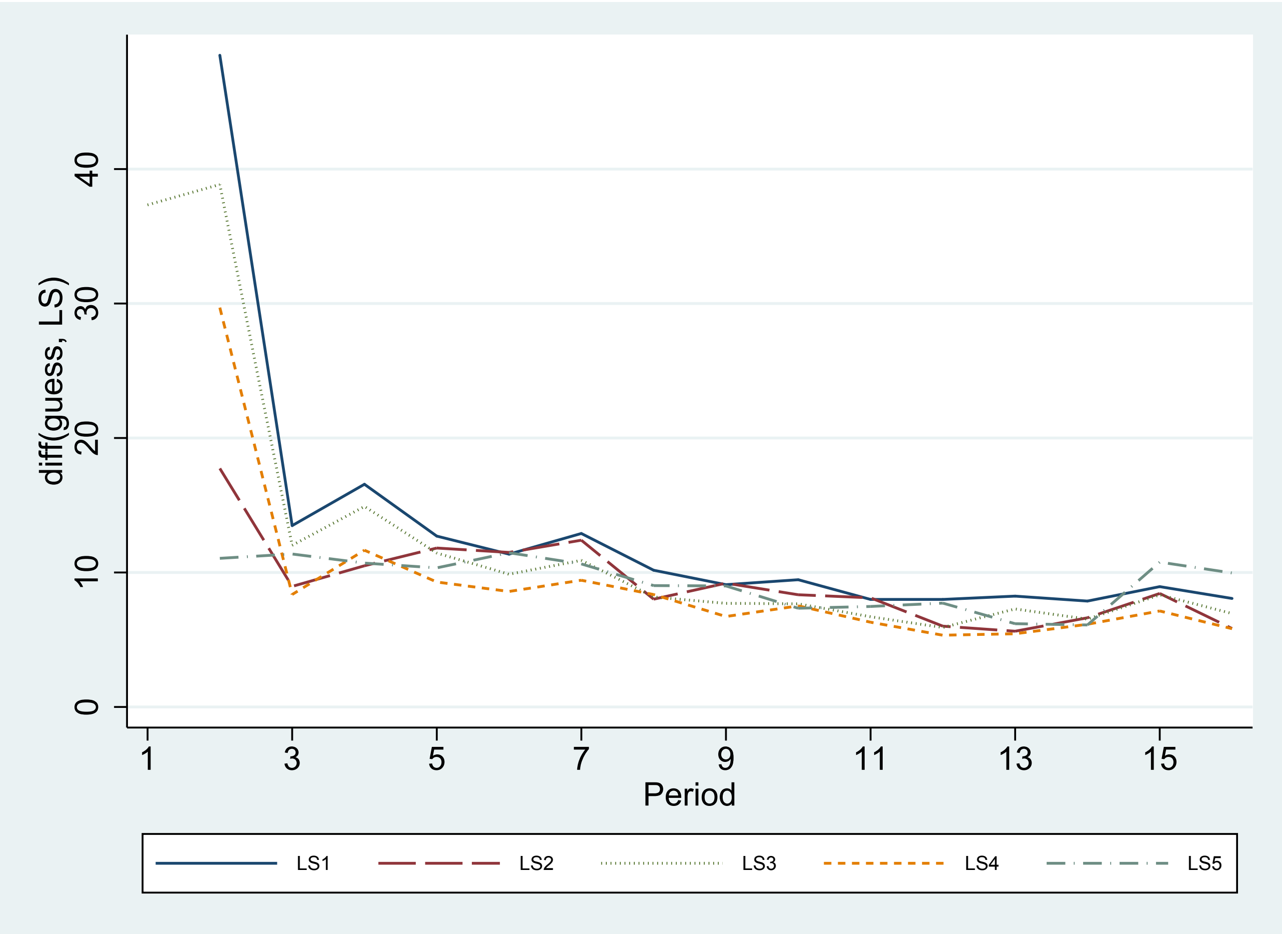

Figure 7: Speed of convergence of LS1-LS5

This figure illustrates how fast one belief update process converges to a constant number. The five belief updates generated by five different learning strategies are considered, respectively. For a given learning strategy (LS1-LS5), the belief adjustments between every two neighboring periods of subject i are first computed based on equation (39), and then averaged over all subjects. The results are illustrated by the line Diff(LS1)_t- Diff(LS5)_t, respectively. As a reference, the belief adjustment based on factual guess game data is computed by equation (41)-(42) and illustrated by the red solid line Diff(guess)_t

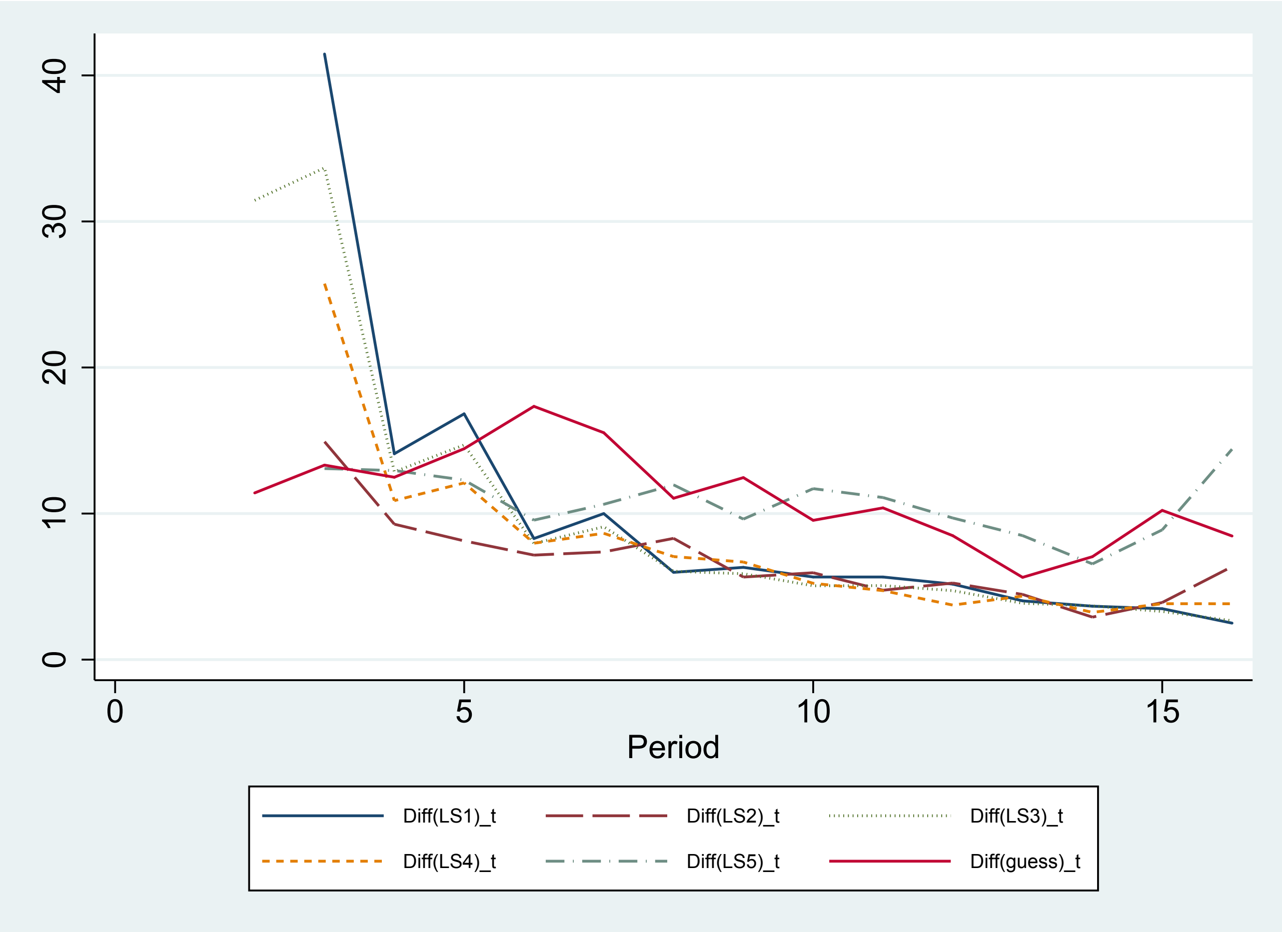

Figure 8: Ranking of LS by subject

This diagram illustrates the by-subject ranking of the five learning strategies, measured by the root mean squared error (RMSE). RMSE is used to quantify, for a given subject, the explanatory power of a given LS to her factual belief updates. The lower the RMSE is, the better the corresponding LS explains the factual belief updates. All RMSEs are computed based on equation (43). Subject 22 has no entry since he/she never updates the belief.

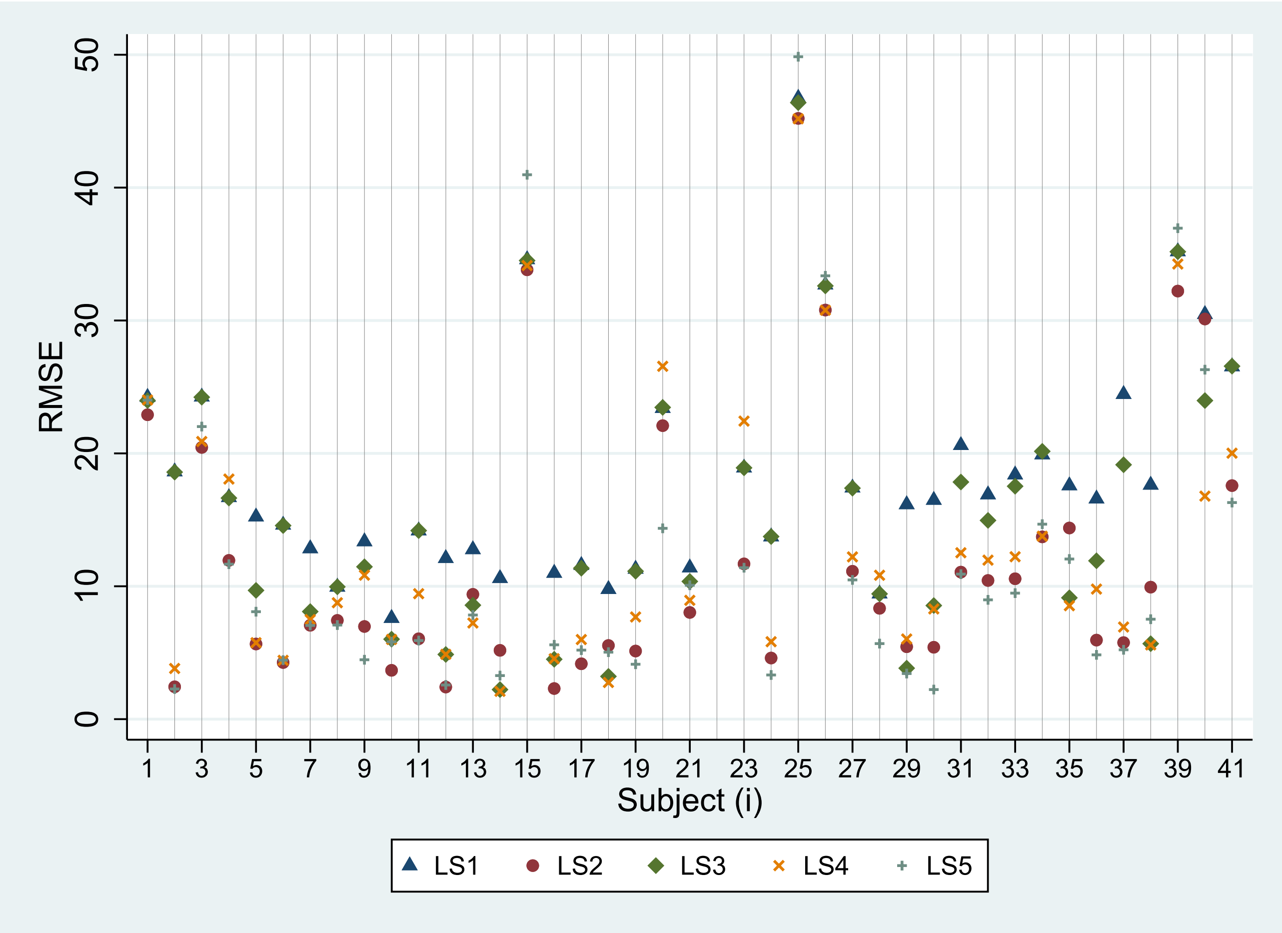

Figure 9: Initial prior distributions of six benchmark strategies B1-B6

This diagram illustrates the PMF of six benchmark strategies B1-B6. B1 has a uniformly distributed prior (the blue solid line). B2 illustrates a symmetric triangle-shape prior; B3-B6 illustrate four priors which are characterized by *beta*-distribution with various shape parameter values. B3 sets $\alpha = \beta = 2$; B4 sets $\alpha = \beta = 3$; B5 sets $\alpha = \beta = 10$; B6 sets $\alpha = \beta = 50$

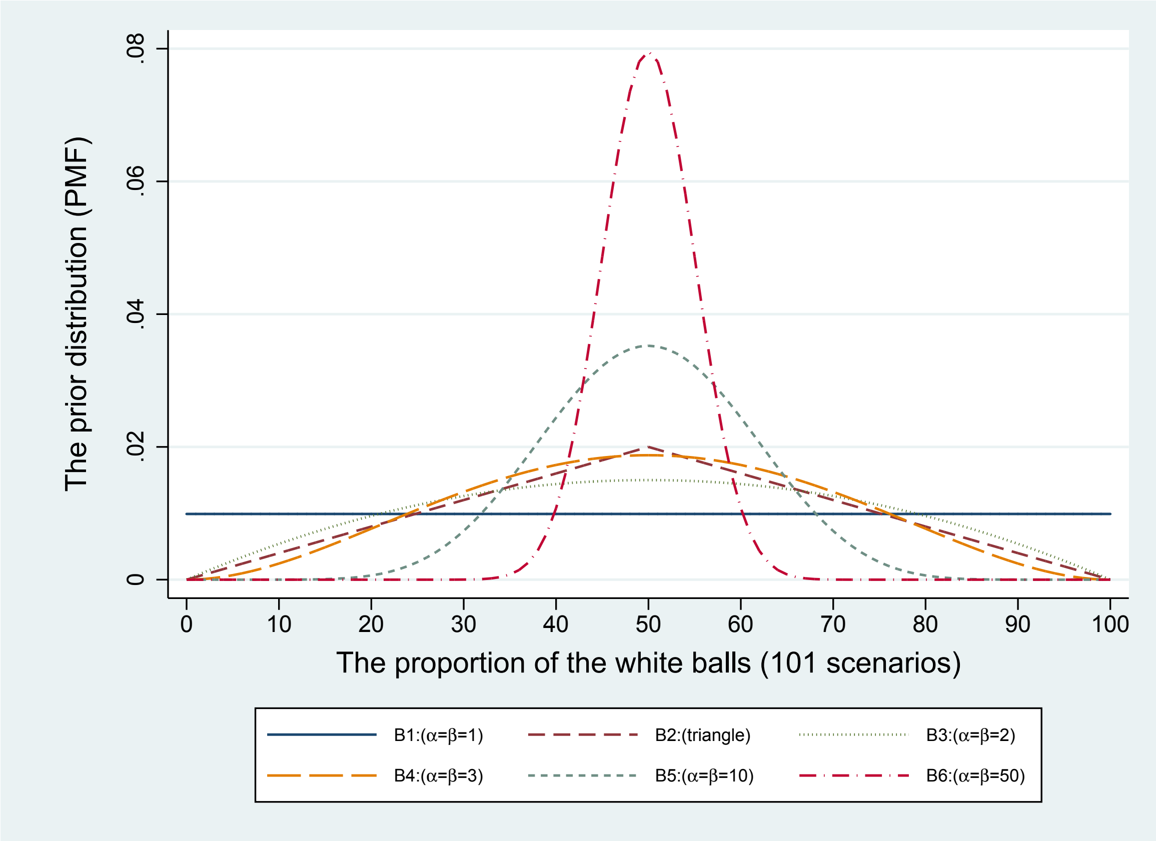

Figure 10: Pessimism/optimism of subjects: in comparison with benchmark B1-B6

Diagram (a) illustrates the absolute difference between the factual guess game entries and six benchmark strategies (B1-B6), respectively. All benchmark strategies employ Bayes rule to update beliefs. They differ from each other only in terms of the initial prior. B1 is constructed with a uniform distribution. B2 assumes a symmetric triangle-shaped initial prior distribution (see Figure 9); B3-B6 assumes the initial priors characterized by the *beta*-distribution with $\alpha = \beta = 2, 3, 10, 50$, respectively. The belief updates of B1-B6 are computed by equation (44)-(45). The per period average absolute difference diff(guess, B1)$_t$-diff(guess, B6)$_t$ are computed by equation (46). sd.columns report the corresponding standard deviation. The absolute difference can be interpreted as the goodness of fit of B1($\cdots$ B6) when regressing the guess game responses. The illustration corresponds with the results in Table 13 Panel A.
Diagram (b) illustrates the exact differences between the factual guess game entries and the six benchmark strategies (B1-B6), respectively, computed by equation (47). sd.columns report the corresponding standard deviation. The exact difference can be interpreted as the pessimism/optimism of subjects, comparing with the chosen benchmark (B1 $\cdots$ B6). Negative value indicates, on average, lower value of the updates (towards zero) than the benchmark value, hence implying pessimism. Positive values, optimism. The illustration corresponds with the results in Table 13 Panel B.

(a) Goodness of fit: **absolute** difference between guess games responses and benchmark B1-B6 updates

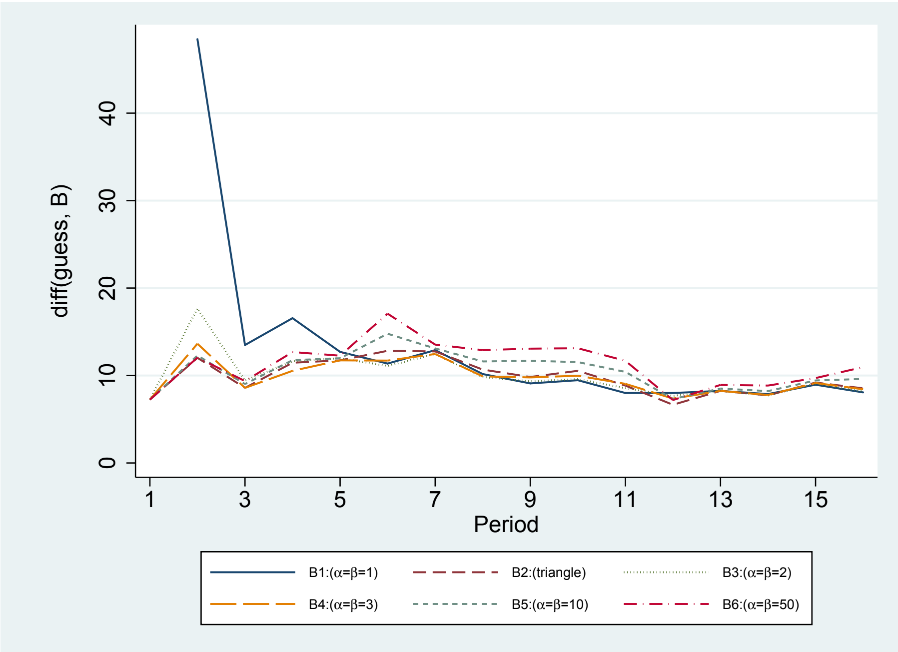

(b) pessimism/optimism: **exact** difference between guess games responses and benchmark B1-B6 updates

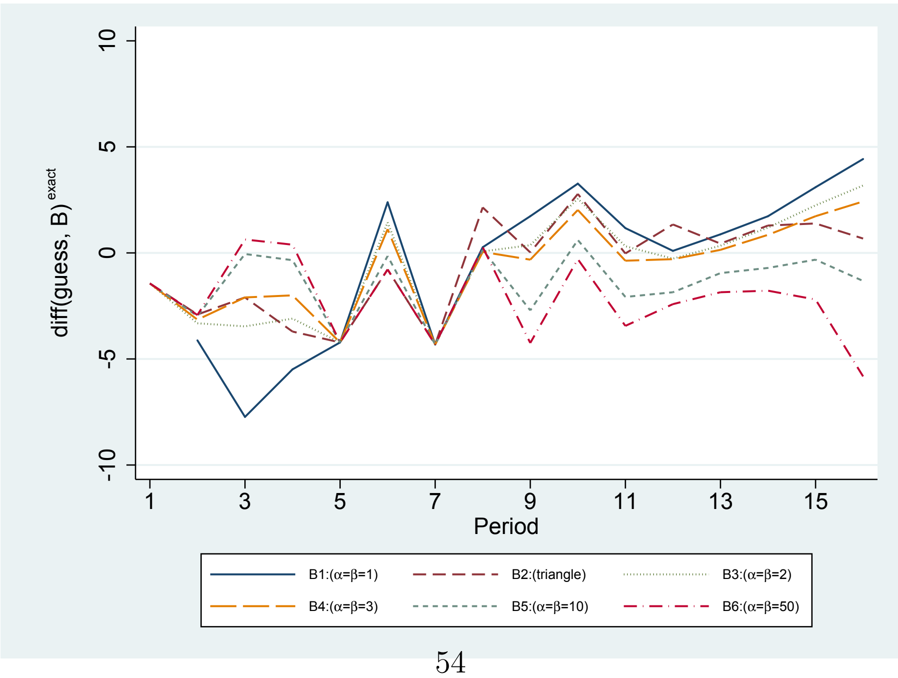

Recent Issues

No. 250	Nathanael Vellekoop, Mirko Wiederholt	Inflation Expectations and Choices of Households
No. 249	Yuri Pettinicchi, Nathanael Vellekoop	Job Loss Expectations, Durable Consumption and Household Finances: Evidence from Linked Survey Data
No. 248	Jasmin Gider, Simon N. M. Schmickler, Christian Westheide	High-Frequency Trading and Price Informativeness
No. 247	Mario Bellia, Loriana Pelizzon, Marti G. Subrahmanyam, Jun Uno, Draya Yuferova	Paying for Market Liquidity: Competition and Incentives
No. 246	Reint Gropp, Felix Noth, Ulrich Schüwer	What Drives Banks' Geographic Expansion? The Role of Locally Non-Diversifiable Risk
No. 245	Charline Uhr, Steffen Meyer, Andreas Hackethal	Smoking Hot Portfolios? Self-Control and Investor Decisions
No. 244	Mauro Bernardi, Michele Costola	High-Dimensional Sparse Financial Networks through a Regularised Regression Model
No. 243	Nicoletta Berardi, Marie Lalanne, Paul Seabright	Professional Networks and their Coevolution with Executive Careers: Evidence from North America and Europe
No. 242	Ester Faia, Vincenzo Pezone	Monetary Policy and the Cost of Wage Rigidity: Evidence from the Stock Market
No. 241	Martin Götz	Financial Constraints and Corporate EnvironmentalResponsibility
No. 240	Irina Gemmo, Martin Götz	Life Insurance and Demographic Change: An Empirical Analysis of Surrender Decisions Based on Panel Data
No. 239	Paul Gortner, Baptiste Massenot	Macroprudential Policy in the Lab