HOUSE OF FINANCE

WORKING PAPER SERIES

Wenhui Li – Christian Wilde

Belief Formation and Belief Updating under Ambiguity: Evidence from Experiments

SAFE Working Paper No. 251

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

Wenhui Li and Christian Wilde†

August 13, 2020

Abstract

Decisions under ambiguity depend on the beliefs regarding possible scenarios and the attitude towards ambiguity. This paper exclusively focuses on beliefs, and beliefs are measured independently from attitudes, in contrast to many previous studies. We use laboratory experiments to estimate the subjective belief formation and belief updating process in an ambiguous environment. As a main contribution, we recover the entire belief distribution of individual subjects and scrutinize how beliefs are updated in response to new information. For 70% of the subjects, we can reject the *objective equality hypothesis* that one's initial prior follows a uniform distribution. A further investigation of biases in initial beliefs reveals that 66% of the subjects display neither pessimism nor optimism in initial beliefs. Overall, the *unbiased belief hypothesis* cannot be rejected. The recovered belief updating rules reveal that the *Bayesian updating hypothesis* can be rejected for 84% of the subjects. Among them, most subjects under-react to new information compared to what Bayes' rule implies. Finally, we find that beliefs are heterogeneous and cannot be characterized by a single distribution that fits for all subjects.

Keywords: ambiguity, belief distribution, belief updates, learning strategy, Bayes' rule, laboratory experiments

JEL code: D81, D83

*We thank Peter Ockenfels for very 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 Leibniz Institute for Financial Research SAFE.

†Correspondence: Leibniz Institute for Financial Research SAFE, Goethe University Frankfurt, Theodor-W.-Adorno-Platz 3, 60323 Frankfurt am Main, Germany. Emails: w.li@safe.uni-frankfurt.de; wilde@finance.uni-frankfurt.de.

1 Introduction

One central question in the decision theory is how the presence of ambiguity affects individuals' decision-making. Ambiguity may affect decision-making through two aspects: how is the belief about the possible scenarios in the ambiguous environment formed and updated, and what is the attitude towards ambiguity. These two aspects: beliefs and attitudes, work very differently in shaping the final decisions. A sizable amount of theoretical literature establishes concepts that separate beliefs and attitudes (Ahn 2008; Brennan 1998; Cao et al. 2005; Chen and Epstein 2002; Epstein and Schneider 2007; Galaabaatar and Karni 2013; Gollier 2011; Karni 2018; Klibanoff et al. 2005). However, empirically it is very difficult to distinguish beliefs from attitudes and recover them separately. In fact, most empirical/experimental literature jointly estimates beliefs and attitudes in one package (Abdellaoui et al. 2011; Baillon et al. 2018). In addition, the literature which elicits beliefs usually pays more attention to belief updating, with initial beliefs being less discussed. To supplement these gaps, this empirical paper may contribute to the current literature in the following ways. First, this paper isolates beliefs from attitudes in an ambiguous environment by adopting a special laboratory experiment design. Second, we use experimental data to estimate each subject's entire initial belief distributions in situations involving ambiguity. Third, we investigate whether initial beliefs are biased towards good or bad outcomes, i.e. whether there exists pessimism or optimism in subjects' initial beliefs. Fourth, we investigate how subjects update their beliefs in response to new information.

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 (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. 2008; 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 urn, whose winning probability is unknown to the players. As designed, the ambiguous lottery urn contains 100 balls. A ball is either white or black. Neither the number of white balls nor the number of black balls is known to the players. In this way, a completely ambiguous environment is operationalized. The ambiguous lottery can be translated into a package of multiple single lotteries. A single lottery in the package has a known winning probability, corresponding with a possible scenario of the number of white balls. This winning probability is defined as the first-order probability, representing the first source of uncertainty, commonly known as risk. Ambiguity, the second source of uncertainty, arises when the realization likelihood of each single lottery (defined as the second-order probability) is unknown. In this paper, we investigate beliefs by tracking down the subjective evaluations of the second-order probability. Practically, in laboratory experiments we ask subjects to guess the number of white balls in the ambiguous lottery urn at some time points. At each time point, a certain number of draws from the ambiguous urn (with replacement) have been implemented, and are provided as information to the subjects. It means that, for each subject, we record her guesses about the composition of the urn along a sequence of draw implementations from the urn. With such guesses, we recover her belief updating strategy and estimate her conceived second-order probabilities after each draw.

The first objective of this paper is to recover the initial subjective evaluation of the second-order probability under complete ambiguity at the very beginning of the experiment, known as the initial prior. Our interest is to recover the initial prior distributions and to capture the shape of the distributions. We first characterize the initial prior distribution by a *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 the uniform distribution, bell-shaped distributions and distributions with monotonic probability density function (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 as well. Due to its attractive features, e.g. simplicity of construction and diversity of distributions, we first assume that one's initial prior follows a *beta* distribution and subsequently recover the shape of the distribution by estimating the shape parameters. To proceed in the analysis of the initial priors, we first hypothesize that in a completely ambiguous environment, all scenarios are considered equally likely. We name it as the *objective equality hypothesis*. Under this hypothesis, the initial prior of a subject follows a uniform distribution. The analysis of initial priors aims to show whether the objective equality hypothesis can be rejected.

Compared to the discussion of posteriors and belief updating, direct analyses about initial priors are rare. One example is Cubitt et al. (2018). The authors apply the smooth model (Klibanoff et al. 2005) to estimate subjects' beliefs, represented by the second-order probabilities. They find that beliefs are heterogeneous and elicited mode values are found in the neighborhood of the true parameter value. Our paper differs from Cubitt et al. (2018) as follows: First, beliefs in Cubitt et al. (2018) are represented by second-order probabilities with binary supports. Another example with binary-support second order probability is Dominitz and Hung (2009). In contrast to them, we design the ambiguous lottery to be completely ambiguous. From the subjects' point of view, it means that the winning probability of the lottery lies between $[0, 1]$, incremented by 0.01. Thus, our experiment generates a richer support space of the second-order probability, inducing high degree of ambiguity. The high degree of ambiguity in our experiment with minimum information provision facilitates the understanding of beliefs in a more general case. Other examples with rather rich support space of the second order probability include Chew et al. 2017 and Coffman et al. (2019), who construct interval-like supports of the second-order probability, yet narrower than ours. Filippis et al. (2017) and Eil and Rao (2011) construct a rich space of second-order probability as we do, but initial priors are not the focus of their analyses.

The second objective of this paper is to analyze the tendency that subjects display towards pessimism/optimism in initial beliefs. In the ambiguity literature, the concept of pessimism/optimism is usually mingled with other concepts such as ambiguity attitude. In many cases, authors use the word pessimism/optimism interchangeably with the word ambiguity aversion/seeking (Binmore et al. 2012; Giraud and Thomas 2017). We argue that pessimism can be conceptualized as the perceptive tendency to bias toward bad scenarios in assigning probabilities to scenarios, and optimism as bias toward good scenarios. Accordingly, we define unbiased beliefs as beliefs with no bias toward bad or good scenarios. We clarify that pessimism/optimism/unbiasedness is, in essence, a characteristic of belief, rather than an attitude. It bases on the shape of belief, prior to and independent from the preference-related decision-making based on this belief. Attitude only enters and takes effect during the phase of preference-related decision-making. Therefore, pessimism/optimism/unbiasedness 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/unbiasedness is partitioned out as a clean independent measurement. The experiment design 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 (optimism) can be measured by the bias towards the worst (best) possible situation, deviating from a neutral benchmark. To investigate whether the subjects display pessimism/optimism in beliefs, we set up the second hypothesis: subjects have unbiased initial beliefs. We name it as the *unbiased belief hypothesis*.

A limited amount of empirical literature in ambiguity treats pessimism/optimism as a belief characterization. One example is Ahn et al. (2014). The authors conceptualize pessimism as overweighting the probabilities of low payoffs and underweighing the probabilities of high payoffs, and optimism as the opposite. Using experimental data, the authors conclude that 71% of the studied subjects are classified as neither pessimistic nor optimistic. However, the authors jointly estimate the parameters governing pessimism and the parameters governing ambiguity aversion so that there exists a risk that the identification of pessimism/optimism may be confounded by attitudes. The mainstream literature, on the contrary, treats pessimism/optimism as an attitude characterization, rather than a belief characterization. With a clearer separation between belief and attitude, our paper provides some new insights about pessimism/optimism in beliefs.

The third 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 new information. The experiment permits learning in the way that a subject always faces the same ambiguous urn through the entire experiment, and that draws with replacement from the ambiguous urn are implemented, generating new information for learning (for the detailed design, see the experiment design section). Therefore, she can update her belief about the urn composition, 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 update their beliefs employing Bayes' rule, and non-Bayesian updates (Epstein 2006; Epstein et al. 2008; Marinacci 2002), where subjects update their beliefs without applying Bayes' rule. To better understand the latent updating mechanism of each subject, this paper models a learning strategy which accommodates that subjects may deviate from Bayesian updating. This learning strategy assumes that a subject forms an initial prior which can be characterized by a *beta*-distribution, and updates her belief parameters (α and β) in response to new information. Depending on how a subject updates the parameters, she may follow Bayes' rule perfectly, but she may also decide to update the parameters more or less than Bayes' rule implies. For each subject, we use experimental data to estimate the parameters which govern one's initial prior and belief updating rule, and recover one's belief dynamics using the estimated parameters. This proposed learning strategy incorporates a variety of initial priors as well as belief updating rules. Such wide coverage seems to be diverse enough, since the factual belief updates of subjects are mostly well fitted by the proposed learning strategy. The recovery of belief dynamics allows us to test the third hypothesis that a subject perfectly employs Bayes' rule when updating beliefs. We name it as the *Bayesian updating hypothesis*.

Sizable empirical literature contributes to belief updating using experimental data. Whether subjects employ Bayes' rule for learning is widely studied in laboratory experiments. Some studies observe incompatibility with Bayesian updating: Filippis et al. (2017) argue that subjects learn from private signals as implied by Bayes' rule only when signals confirm their priors, but overweigh signals when they contradict their priors. Buser et al. (2018) and Dominitz and Hung (2009) observe that subjects update beliefs more conservatively than if Bayes' rule is employed, independent of signal type. Other studies observe consistency with Bayes' rule: Mobius et al. (2014) observe that subjects process learning in the way that priors are sufficient statistics for past information, supporting Bayes' rule. Eil and Rao (2011) also discover that subjects may conform to Bayes' rule more closely in response to good news than bad news. Other authors try to explain the employment of Bayes' rule in line with gender (Coffman et al. 2019) and an event's self-relevance degree (Ertac 2011). It seems that whether one employs Bayes' rule at all in belief updating, as well as to what extent one follows Bayes' rule, is hardly conclusive in the existing literature. Our paper, which proposes an imperfect Bayesian updating strategy, may add new insights to this discussion.

This paper reaches the following findings: (1) The initial prior distributions are mostly characterized by bell-shaped *beta*-distributions, rather than a uniform distribution. This result does not support the *objective equality hypothesis* that in a completely ambiguous environment all scenarios are considered equally likely. (2) The investigation of pessimism/optimism in initial beliefs supports the *unbiased belief hypothesis*. Subjects mostly demonstrate no bias towards either bad or good situations. Among biased initial beliefs, optimism is more likely to be seen than pessimism. (3) The recovered belief updating dynamics show that the *Bayesian updating hypothesis* can be rejected for 84% of the subjects. It means that most subjects respond to new information different from perfect Bayesian updaters. Among them, most subjects under-react to information in comparison to what Bayes' rule implies.

The rest of the paper is organized as follows: Section 2 presents the background and underlying assumptions related to the discussion of beliefs under ambiguity. Section 3 introduces the experiment design. Section 4 presents the descriptive analysis of the belief update data. Section 5 presents the learning strategy model and discusses the belief estimation results. Section 6 tests the three hypotheses. Section 7 concludes.

2 Background

This section introduces the main concepts and underlying assumptions related to ambiguity and belief updates.

Ambiguity and second-order probability. Let Ω denote a state space with $\omega \in \Omega$ as various states. A set of subsets of Ω, written as $\sigma(\Omega)$, are called events. A countably additive probability π is a mapping which associates events with likelihood measurements: $\pi : \sigma(\Omega) \to [0, 1]$. The set of all possible π is denoted as Δ. In other words, Δ contains all possible ways which map $\sigma(\Omega)$ into likelihood measurements.

The first source of uncertainty arises from that no likelihood measurement in π is equal to one. It means that whether a given event in $\sigma(\Omega)$ will happen is uncertain. π, which describes the likelihood of events, is defined as the first-order probability. This source of uncertainty is interpreted as risk. Δ contains multiple elements, indicating multiple ways to value π. Define a mapping which associates subsets of Δ to likelihood measurements: $\mu : \sigma(\Delta) \to [0, 1]$. The second source of uncertainty arises from that no likelihood measurement in μ is equal to one. This source of uncertainty is interpreted as ambiguity. μ is defined as the second-order probability, and assumed to be countably additive.

Belief. The term *belief* in this paper refers to the second-order probability μ. The paper aims to estimate the second-order probability μ for each subject. Consider the ambiguous lottery example again. Define there are only two states of the lottery, winning or losing, written as $\Omega = \{w, l\}$, where w denotes winning and l denotes losing. The winning probability of the lottery can be defined as the likelihood of winning if the lottery is played: π_w. Likewise, we can define $\pi_l \equiv 1 - \pi_w$ for the losing probability of the lottery. $\pi = \{\pi_w, \pi_l\}$ constructs the first-order probability. The first source of uncertainty, risk, arises from the assumption that π does not include an element equal to one. Yet there are multiple possibilities to construct π. For instance, π can be $\{0.1, 0.9\}$, $\{0.2, 0.8\}$, or other values. We can define all possible values of π as $\Delta = \{\pi^j : j = 1, 2, \cdots\}$. The second-order probability is denoted as $\mu : \sigma(\Delta) \to [0, 1]$. The assumption that μ does not include an element equal to one leads to the existence of ambiguity, i.e. $\mu_j \neq 1$ for any element indexed by j in $\sigma(\Delta)$. Back to the ambiguous lottery example, μ is the likelihood measurement for each possible winning probability of the lottery.

In our paper, the second-order probability μ is assumed to be subjective and non-degenerate for most of the subjects. Each individual can associate the subsets of Δ to some likelihood measurements based on her own understanding. Hence, μ for subject i is likely to be different from μ for subject i', $i \neq i'$. Non-degenerate indicates that $\mu_j \neq 1$ for any j, at least for most of the subjects. Therefore, μ is non-trivial and experiments are designed to estimate the subjective μ for each subject.

Belief updating. μ is not static over time. Subjects can update μ using the available relevant information. This process of updating μ is also called *learning*. Reconsider the ambiguous lottery example, μ may be updated if the lottery is repeatedly played and the results, winning or losing, are observable to the subjects. In order to recover μ and its updating process, the experiment is designed to elicit both the mode value of initial μ when no information about the winning probability of the lottery is provided, and the mode value of updated μ when some information is generated by repeatedly playing the lottery in front of the subjects. The initial μ is called the initial prior, while the updated μ are called posteriors. Each lottery play generates one piece of new information, and thus corresponds with one specific posterior.

Mode. μ is, in essence, a probability distribution, whose shape is of our main interest. However, in practice it is hard to design a truth-inducing incentivization method to elicit the entire distribution. Our experiment, instead, elicits the mode value of μ. The mode of μ is defined as a subset of Δ which is associated with the highest realization likelihood measurement: $M(\mu) := \{\pi^j \in \Delta : \underset{j}{\text{argmax}}\ \mu_j\}$. In theory, the $M(\mu)$ is not necessarily a singleton. For instance, if μ follows a uniform distribution, i.e. $\mu_j = \mu_k$ for all $j \neq k$, then $M(\mu)$ is any π^j. This paper allows a uniformly distributed μ.

Perceived degree of ambiguity and confidence. The perceived degree of ambiguity can be identified by the dispersion measurement of μ: how evenly are the realization likelihoods distributed across all elements of Δ. The more evenly they are distributed, the higher is the perceived degree of ambiguity. In contrast, the more squeezed the likelihood distribution is, the lower is the perceived degree of ambiguity. One common measurement of the perceived degree of ambiguity is the variance of μ. In line with the above logic, the higher the variance of μ is, the higher is the perceived degree of ambiguity.

Confidence is a concept describing the perceived degree of ambiguity based on subjective μ. If a subject has a rather dispersed μ distribution, it can be said that from her point of view, the perceived degree of ambiguity is high, and her confidence about which scenario is to realize is low. In contrast, if a subject has a rather squeezed μ distribution, it can be said that from her point of view, the perceived degree of ambiguity is low, and she is confident about which scenario is to realize. It is worth mentioning that confidence is a subjective concept, rather than a objective one. It reflects the characteristics of one's own belief.

Pessimism, optimism, and unbiased beliefs. The ambiguous lottery is designed such that winning pays out a positive financial reward, while losing pays out zero. Subjects should always prefer winning to losing. This preference holds independent from the utility function form, as long as it is monotonically increasing. Therefore, we define that favorable situations are situations with higher winning probability $\pi_w^j \in \Delta$, and unfavorable situations are situations with lower $\pi_w^j \in \Delta$. Roughly speaking, more likely to win is better than less likely to win. The best situation corresponds with $\pi_w = 1$ and the worst situation corresponds with $\pi_w = 0$.

As a characteristic of belief, pessimism is defined as a bias in belief towards unfavorable situations. It is characterized by a second-order probability μ which associates more likelihood to unfavorable situations than to favorable situations. In contrast, optimism is defined as a bias in belief towards favorable situations. It is characterized by a second-order probability μ which associates more likelihood to favorable situations than to unfavorable situations. Unbiased beliefs are beliefs which display no biases towards unfavorable or favorable situations. The midpoint between the most favorable situation and the most unfavorable situation might stand for a neutral benchmark, discriminating between favorable and unfavorable situations.

3 Experiment design

3.1 Guess game

We operationalize the ambiguous environment with 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 white balls nor the true proportion of black balls is known to any subject. It is also explained to the subjects that the number of white balls 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 a number between zero and 100 (zero and 100 are inclusive), corresponding to her own belief about the proportion of white balls in the urn. 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 subjects. Then the ball is immediately put back into the urn (draw with replacement). Guess games and draws are played/implemented for multiple times and the sequence is designed as follows: A subject first plays the initial guess game before any draw is implemented, denoted as G_0, where the subscript indexes the number of draws already implemented. Then the first draw is implemented. This sequence, one guess game followed by one draw implementation, repeats 15 times. In addition, in all Sessions except Session I, an additional guess game, G_{15}, is played after the 15^{th} draw[1]. Table 1 displays the complete experiment procedure. Following this design, when a subject plays guess game G_n, she has already observed n times of draws ($n = 0, 1, \cdots, 15$). The belief reported in guess game G_n can be seen as the updated belief based on the learning of the n-time draw history. The past draw history, if any, is displayed on screen for subjects' reference. Figure 1a displays the screen-shot of the initial guess game G_0 when no draws are implemented. Figure 1b, as an example, displays the screen-shot of guess game G_5, in which the history of five draws is displayed on the screen for subjects' reference. Table 2 reports the guess games played by each subject.

The true proportion of white balls in the ambiguous urn is fixed at 40 for all subjects. Basically, it is beneficial to choose a value close to 50. It would generate more balanced proportions of white draws and black draws along one's draw history. In contrast, choosing

[1]Subjects in Session I go through the sequence: Guess game G_0, first draw, G_1, second draw, $\cdots$, G_{14}, 15^{th} draw. Subjects in Session II-VII go through the sequence: G_0, first draw, G_1, second draw, $\cdots$, G_{14}, 15^{th} draw, *plus* G_{15}.

a value far away from 50 increases the probability that draws of one color are much more frequently observed than draws of the other color. Hence, if the true proportion of white balls is set close to 50, subjects' responses to both white draws and black draws are more evenly observed. Additionally, choosing a value close to 50 also increases the variation of draw history, conducive to the belief estimation. However, the value 50 itself is too prominent. Subjects may instinctively stick to this prominent point in the guess games. If true, we cannot tell whether they learn correctly, or they simply keep entering the prominent point ignoring learning. Therefore, we choose 40 as the true proportion of white balls. It is close to 50 but not too prominent, and not too hard to guess. The fact that all subjects are faced with the same ambiguous urn is to facilitate the comparability across subjects. Subjects are incentivized in the way that every time one enters the correct number of white balls (=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 their personal prior/posterior distribution (if any) in each guess game. The earning from the guess games of a subject is only announced to her at the very end of the experiment, hence the ambiguous feature of the urn sustains through the entire experiment.

It is worth mentioning 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 subjects actually play the lottery and may choose between alternatives. 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 ambiguous environment, the reported data in the guess games is only related to beliefs.

3.2 Incentive compatibility

In this part, we provide a brief discussion to support the incentive compatibility of our experiment design. The guess games are intended to investigate the mode values of subjective prior/posterior distributions. It is argued that eliciting the mode value of a discrete belief distribution is truth-telling, independent of subjective attitudes (Hurley and Schogren 2005). Below, we briefly show that the incentivization method in our experiment induces truth telling.

Suppose a subject has a prior/posterior (second order probability) when playing a guess game, denoted as μ. μ_j is defined as the PMF value of μ evaluated at the support j. $j = 0, 1, \cdots, 100$ denotes the 101 possible numbers of white balls in the urn. Under the applied incentivization method, the expected earning of a subject who enters j' in the guess game is $2\mu_{j'}$ Euro. To maximize the expected earning, one would choose j' such that $j' = \underset{j}{\operatorname{argmax}}\ \mu_j$, namely the mode value of μ. The same logic applies to each guess game $G_n : n = 0, 1, \cdots, 15$. Therefore, our experiment induces one to report her mode value of each prior/posterior in each guess game.

One concern may be the possibility that subjects hedge across guess games. It means that a subject may ignore her subjective prior and posteriors, distributing her guesses to cover the support domain as widely as possible. For example, a subject may disregard learning, and randomly enter 16 different numbers in the 16 guess games (i.e. entering guesses without repetition). In practice, such hedging strategy is possible, and it would undermine data quality because subjects enter numbers that do not correspond to their true beliefs. However, a dominant strategy is to choose 16 numbers, each of which is close to the conceived mode value of one's prior/posterior at that time, and simultaneously to ensure no repetition among the 16 numbers. The closer the numbers are to the mode value, the more possible it is to win something from the guess games. Therefore, a subject who plans to hedge, should have more incentive to hedge in such "clever" way rather than randomly choosing 16 numbers. It also means that, even if one hedges across the guess games, her guess game responses should closely reflect her conceived mode values. This mitigates our concern related to hedging. In fact, as the data show, only six out of 102 subjects enter unrepeated numbers in the guess games. For the rest of the subjects, repeated guesses show up. Hence, hedging across guess games tends to be negligible in our sample, both from a theoretical and an empirical perspective.

Another concern may be the possibility that subjects hedge across different parts of the experiment. If true, one's elicited beliefs may be different from one's true beliefs, and thus the data quality may be undermined. Such possibility needs to be checked, since apart from the guess games, our experiment also contains other parts. This may lead to hedging. In some choice games, an asset is constructed, whose payoff is determined by a random draw from the identical ambiguous urn presented in the guess games: a white draw pays out a positive financial reward, while a black draw always pays out zero. Subjects are asked to report their conceived reservation values of the asset. It is possible that subjects hedge between the guess games and such choice games: For instance, a subject who reports high reservation values in the choice games to benefit from a higher-than-believed proportion of white balls, may enter small numbers in the guess games to benefit from a lower-than-believed proportion of white balls. This strategy, betting on many white balls in one game and betting on many black balls in the others, constitutes hedging. Thus, the elicited data on beliefs may be biased. For this matter, we check whether hedging exists in our data. If hedging exists, following the logic above, a negative correlation between the responses in the guess games and the reservation values reported in the relevant choice games should be observed. Moreover, since subjects naturally benefit from a high number of white balls in the choice games, responses close to zero should prevail in the guess games for maximum benefit from hedging, if hedging exists in the first place. A check of the reported responses reveals that hedging does not seem to occur in our sample. First, such correlation is significantly positive, rather than negative. Second, responses close to zero are rarely seen in actual data: only nine out of 1619 responses (0.56%) are equal to zero. 74 responses (4.57%) are in the interval $[0, 19]$. In fact, the responses equal to zero are as many as the responses equal to 100, implying that no bias towards extremely low responses exists. In all, it seems that hedging across different parts of the experiment is negligible in our sample.

Our experiment is intended to only elicit the mode values of one's belief distributions. It seems that, for our study, eliciting the mode is more appropriate than other alternatives. One alternative would be to ask subjects to guess an interval which nests the true parameter value. Coffman et al. (2019) adopts a similar approach to elicit one's relative ranking of cognitive ability in experiment population. Although the chance to earn increases, one problem is that the elicited *interval* modes are not as informative as elicited *singleton* modes. Another alternative would be to elicit distributions directly, with incentivization following the quadratic scoring rule (QSR) (Eil and Rao 2011; Ertac 2011). However, the incentive compatibility of QSR relies on the assumption of risk neutrality. Countless experimental evidence already shows that subjects are mostly risk averse in laboratory experiments (Holt and Laury 2002). At last, it is practically easier for subjects to report mode values than to report distributions. Considering all these factors, we finally choose to elicit the mode of one's prior/posterior distribution, rather than other alternatives.

3.3 Other information

The laboratory experiment is computerized by Z-tree (Fischbacher 2007). Seven sessions of this experiment with 102 subjects in total have been conducted. The 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 Frankfurt. Table 2 summarizes the distribution of subjects in each session. In Session I-VI, subjects are assigned into markets. Each market implements its draws independently. A subject only observes the draw information of her own market. The design of market-specific draw implementation is to meet the requirement of market-wide asset trading. Session VII is without asset trading. 19 subjects participate in this session and each of them observes her own draw history, yielding 19 additional paths of draw history. In total, during the seven sessions of the experiment, 30 different paths of draw history are generated and 102 belief update dynamics are recorded. The variation of the data tends to enhance the explanatory power of our findings.

Apart from the guess games, a complete experiment session also contains several other parts, including choice games, asset trading (except in Session VII), as well as some quizzes and practice rounds for training, and a short questionnaire. Since the choice games and the asset trading parts are not of interest in this paper, we suppress an extensive introduction of these parts, except for certain points worth to be mentioned: (1) As explained above, the ambiguous urn used in the guess games is also used in other parts to determine the asset payoff: each time a white ball is drawn, the asset pays out a positive dividend; each time a black ball is drawn, the dividend is zero. 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. (2) Some quizzes, a demonstration of a physical ambiguous lottery urn (in fact, a big box containing 100 chips with "white" or "black" written on), and trial draws from risky urns are implemented to help the subjects fully understand the concept of a lottery and the guess game. The content of the ambiguous urn is not observable to any subjects, and no draws from the ambiguous urn are implemented during the demonstration. Therefore, when the first guess game G_0 is played, the lottery urn is completely ambiguous, as intended.

A complete experiment session lasts about 2 hours and 15 minutes on average. The total earning from all parts is on average 31 Euro.

4 Descriptive analysis

The guess game response of subject i after n draws are observed is denoted as $white_{i,n}$. Figure 2 presents the frequency distribution of the responses from the first guess game ($white_{i,0}$). Each subject plays 15 (in Session I) or 16 (in all other sessions) guess games, thus generating 15 or 16 responses. Subjects are assigned into markets, as Table 2 informs. In Market 1-11, draw implementations are market-specific. It means that each market generates its own draw history, and subjects in the same market observe the same draw history. For Market 1-11, we compute the market-average guess game response at each n value, defined as:

$$\overline{white}_{m,n} = \frac{1}{N_m} \sum_{i \in m} white_{i,n} \tag{1}$$

$$m = 1, 2, \cdots, 11; \quad i \in \{1, 2, \cdots, N_m\}; \quad n = 0, 1, \cdots, T$$

where N_m denotes the number of subjects in market m, and T denotes the highest number of draws[2]. For the 19 subjects in Market 12, each subject observes her own draw history, thus generating 19 independent paths of draw history. We can directly report $white_{i,n}$ for each i in Market 12. All paths of draw history and the corresponding guess game responses, market-average or subject-specific, are illustrated in Figure 3.

In total, our experiment generates 30 different paths of draw history. Figure 3 illustrates each path in one sub-graph. Sub-graphs 1-11 represent the eleven paths in Markets 1-11, respectively. Sub-graphs 12-30 represent the other 19 paths, each for one subject in Market 12. Each gray bar represents a white draw, otherwise, a black draw. It can be seen that, across the 30 draw history paths, there are in total almost as many white draws as black draws. It generates a rather balanced data set, with subjects' responses to both white draws and black draws being evenly observed. On the other hand, some paths generate relatively more white draws, others generate relatively more black draws. This increases the variety of draw history paths.

The blue lines in each sub-graph represent subjects' responses in the guess games. Sub-graphs 1-11 illustrate the values of $\overline{white}_{m,n}$ for Markets 1-11, respectively. Sub-graphs 12-30 illustrate the values of $white_{i,n}$ for the 19 subjects in Market 12, respectively. Ups and downs along the blue lines are observed in almost all sub-graphs. It implies that most of the subjects actively respond to the draw information, and learning is prevalent and observable. Another observation is that most subjects update their beliefs rationally. In theory, a subject who updates her beliefs rationally should act as follows: each time a white draw is observed, she does not downwards adjust her current belief; Each time a black draw is observed, she does not upwards adjust her current belief. This definition can be understood as weakly rational. Following this definition, we compute the rate of rational belief updates for each subject, defined as the proportion of rational belief updates out of all 15 updates (14 for Session I) of a given

[2] In Market 1, 14 draws (with replacement) are implemented, and accordingly 15 guess games ($G_0, G_1, \cdots, G_{14}$) are played, hence $T = 14$ for Market 1. For Markets 2-11, in each market 15 draws (with replacement) are implemented, and accordingly 16 guess games ($G_0, G_1, \cdots, G_{14}, G_{15}$) are played, hence $T = 15$ for Markets 2-11. See Table 2 for details.

subject. The proportion of rational belief updates is reported in Figure 4 in form of a frequency distribution. It can be seen that most subjects update their beliefs rationally. More than 60% subjects have a rate of rational belief updates no less than 0.9. i.e. over 13 out of 15 updates meet the rationality definition. Only eight subjects have rates lower than 0.5. The overall mean value is 0.84. The fact that subjects do learn from the draw information and update beliefs rather rationally motivates us to study their learning strategies thoroughly. To proceed, we construct a learning strategy which may be used to model subjects' belief updates.

5 Learning strategy

In this chapter, we propose a learning strategy to model subjects' belief updating dynamics. This learning strategy develops the thought of the conjugate prior theory in the Bayesian analyses (Diaconis and Ylvisaker 1979; Gelman et al. 2004; Schlaifer and Raiffa 1961). The conjugate prior theory assumes that a subject forms an initial prior which can be characterized by a *beta*-distribution, and employs Bayes' rule to update the prior using new information. The learning strategy of this paper inherits the assumption on the initial prior, but eases the assumption on the updating rule so that a subject can freely choose either to follow Bayes' rule or to deviate from Bayes' rule. By construction, this learning strategy embraces rich flexibility in the way of learning: it covers a variety of initial prior distributions, supported by the variety of *beta*-distribution characterizations. In addition, this learning strategy covers both perfect Bayesian updating (as a special case) and imperfect Bayesian updating (i.e. deviation from Bayes' rule). To better organize the detailed introduction, we first introduce the Bayesian special case of this learning strategy, in which a subject perfectly follows Bayes' rule. Then, based on this Bayesian special case, we extend the learning strategy into a more general form, i.e. imperfect Bayesian updating, in which a subject may deviate from Bayes' rule.

5.1 Learning strategy: the Bayesian special case

In the laboratory experiments, the urn used to operationalize the ambiguous environment in the guess games is completely ambiguous: there are 101 possible scenarios regarding the true composition of the urn. These 101 scenarios can be indexed by the possible proportion of white balls (in percent), denoted by θ, where $\theta = 0, 0.01, 0.02, \cdots, 1$. The prior/posterior distribution of a subject at any point in time is given by the probabilities that she assigns to each of the 101 scenarios.

The Bayesian special case of the learning strategy assumes that a subject forms an initial prior characterized by a *beta*-distribution, and employs Bayes' rule to update her priors using new information. For simplicity, we denote this special case as LS$'$. According to Bernstein von Mises Theorem, LS$'$ generates belief updates that asymptotically converge to the true parameter value. The *beta*-distribution is chosen for initial priors since it characterizes a wide range of distributions defined in the interval $[0, 1]$, parameterized by only two shape parameters α and β. Figure 5 illustrates some examples. In case of $\alpha = \beta = 1$, the distribution becomes a uniform distribution. In case of $\alpha = 1,\ \beta > 1$, the probability density function (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 $\alpha,\ \beta > 1$, the PDF has a bell shape[3]. The higher the value of $(\alpha + \beta)$ is, the more squeezed is the distribution.

Although the *beta*-distribution mostly applies to continuous distributions, the 101 discrete scenarios of the white-ball proportion is to some extent dense enough. Later we discretize the *beta*-distributions, translating the PDF representations of the distributions into probability mass function (PMF) representations, to match the 101 discrete supports. Starting with a continuous prior distribution does little harm, since our main interest is the mode of the distribution, and the maximum discrepancy between the mode read from the distribution described by a continuous PDF and the mode read from the distribution described by a discrete PMF is only 0.005 (out of 1). We restrict $\alpha \geq 1$ and $\beta \geq 1$ to guarantee the uniqueness of mode value whenever a mode exists.

Let $(\alpha_{i,0}, \beta_{i,0})$ denotes the shape parameter bundle which governs the initial prior of subject i (the subscript "0" denotes that no draws are yet observed). The initial prior (in form of PDF) can be written as:

$$Prior(\theta|\alpha, \beta) = \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)}\theta^{\alpha-1}(1-\theta)^{\beta-1} \tag{2}$$
$$\theta \in [0, 1]; \quad i \in \{1, 2, \cdots, N\}; \quad \alpha, \beta \geq 1$$

where $\Gamma(\cdot)$ denotes the gamma function. For readability, we suppress the subscripts of $\alpha_{i,0}$ and $\beta_{i,0}$ for the time being. Suppose after n draws, subject i observes $k_{i,n}$ units of white draws. The posterior of subject i following Bayes' rule reads:

$$Posterior(\theta|n, k; \alpha, \beta) = \frac{Prior(\theta|\alpha, \beta) \times Prob(k|n, \theta)}{\int_{\theta'} Prior(\theta'|\alpha, \beta) \times Prob(k|n, \theta')d\theta'} \tag{3}$$

where

$$Prob(k|n, \theta) = \binom{n}{k}\theta^k(1-\theta)^{n-k} \tag{4}$$

$$\binom{n}{k} = \frac{n!}{k!(n-k)!} \tag{5}$$

For readability, we suppress the subscript of $k_{i,n}$ as well. It can be shown that the posteriors are also characterized by a *beta*-distribution, with updated shape parameter bundle $(\alpha + k,\ \beta + n - k)$. This property is summarized in the proposition below:

Proposition 1. *Suppose that Subject i, with an initial prior characterized by a beta-distribution with shape parameter bundle (α, β), updates her initial prior into a posterior employing Bayes' rule. After n draws are observed with k of them being white draws, her posterior can then be characterized by a beta-distribution with shape parameter bundle $(\alpha + k,\ \beta + n - k)$.*

[3] For a *beta*-distribution with shape parameters $\alpha,\ \beta \geq 1$, its mean value is computed by $\frac{\alpha}{(\alpha+\beta)}$. Its mode value is computed by $\frac{\alpha-1}{(\alpha+\beta-2)}$ if $\alpha,\ \beta > 1$; the mode is equal to zero for $\alpha = 1,\ \beta > 1$; the mode is equal to one for $\alpha > 1, \beta = 1$; and the mode takes any value in [0, 1] for $\alpha = \beta = 1$. Its variance is computed by $\frac{\alpha\beta}{(\alpha+\beta)^2(\alpha+\beta+1)}$.

The proof can be found in Appendix A. Given this property, we can also visualize LS′ in the following way: A subject forms an initial prior characterized by a *beta*-distribution with (α, β). Then the draw implementation starts. Whenever the subject observes a white draw, she adds one to the α term. Whenever she observes a black draw, she adds one to the β term. Accumulatively, the subject in fact adds the number of white draws to the α term and adds the number of black draws to the β term: after n draws, with k of them being white draws and $n-k$ of them being black draws, the α term becomes $\alpha+k$ and the β term becomes $\beta+n-k$. Her posterior is thus updated to a new *beta*-distribution characterized by $(\alpha+k,\ \beta+n-k)$. Thus, this updating process can be summarized by the following equations:

$$\alpha_{i,n} = \alpha_{i,0} + k_{i,n} \tag{6}$$

$$\beta_{i,n} = \beta_{i,0} + n - k_{i,n} \tag{7}$$

$$k_{i,0} = 0;\ \text{all } \alpha, \beta \text{ terms} \geq 1;\ n = 1, 2, \cdots, T;\ i \in \{1, 2, \cdots, N\}$$

where $\alpha_{i,n}$ and $\beta_{i,n}$ denote the (updated) shape parameters governing subject i's prior/posterior after n draws are observed. $\alpha_{i,0}$ and $\beta_{i,0}$ denote the shape parameters governing i's initial prior (before any draw is observed, i.e. $n=0$). $k_{i,n}$ denotes the number of white draws that subject i observes out of n draws. Naturally, $k_{i,0}=0$: when no draws are implemented, no white draws are observed.

This pattern that adds the frequency of white draws to the initial $\alpha_{i,0}$ and adds the frequency of black draws to the initial $\beta_{i,0}$ also reflects how the shape of the posterior distribution is updated. If white draws are more frequently observed than black draws, $\alpha_{i,0}$ receives relatively more add-ups than $\beta_{i,0}$. According to the property of *beta*-distribution, this results in a more left-skewed distribution with a higher mode value. On the other hand, if black draws are more frequently observed than white draws, $\beta_{i,0}$ receives relatively more add-ups than $\alpha_{i,0}$, causing the posterior *beta*-distribution to be more right-skewed and the resulting mode value is lower. 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 draws accumulate, n increases for sure and $k_{i,n}$ is on the track of increase. Hence, both $\alpha_{i,n}$ and $\beta_{i,n}$ are on the track of increase. The PDF of a *beta*-distribution with a larger parameter bundle concentrates more densely on some narrower bandwidth. The bell-shape PDF curve looks thinner and taller. This is also intuitive, since when learning evolves, subjects tend to feel more confident about their beliefs, and the posterior distribution becomes more squeezed. Another good feature of the Bayesian updates is that the marginal adjustment diminishes when time evolves. It means that when n increases, subjects are less sensitive to an additional draw, and thus adjust their beliefs more mildly. Consider the belief adjustment at $n=1$ and at $n=15$ for instance. Adding one piece of new information to the existing one piece can largely affect the general outlook of the information set and thus may cause large adjustment in beliefs. On the contrary, adding one piece of new information to the existing fifteen pieces can hardly generate a similar effect. The importance of new information decreases when the total amount of information increases. Hence, the diminishing marginal adjustment should be taken into account when modeling the subjective belief updating process. This is achieved by this learning strategy.

It is also worth mentioning that the Bayesian special case of this learning strategy nests the maximum likelihood updating, one of the most prominent strategies in updating: In case that a subject has a uniformly distributed initial prior (i.e. $\alpha_{i,0} = \beta_{i,0} = 1$) and updates her priors following Bayes' rule, the mode value of her posterior distribution is always equal to the maximum likelihood update after any particular number of draws ($n = 1, 2, \cdots$) of the same draw history. The proposition below formally presents this relation:

Proposition 2.

$$\underset{\theta}{\operatorname{argmax}}\, Posterior(\theta|n, k_{i,n}, \alpha_{i,0}, \beta_{i,0}) = \frac{k_{i,n}}{n} \tag{8}$$

$$\alpha_{i,0} = \beta_{i,0} = 1; \quad n = 1, 2, \cdots, T; \quad i \in \{1, 2, \cdots, N\}$$

The formal proof of this proposition can be found in Appendix B. The LHS of Equation (8) denotes the mode of subject i's posterior distribution after observing $k_{i,n}$ units of white draws out of n draws. The RHS denotes the maximum likelihood update based on the same draw history. Since the guess games elicit the mode values of subjects' prior/posterior distributions, the LHS of Equation 8 is observable. Proposition 2 implies that if a subject adopts maximum likelihood method in belief updating, her belief dynamics can be equivalently recovered by LS′ with $\hat{\alpha}_{i,0} = \hat{\beta}_{i,0} = 1$. In this case, we observe the maximum likelihood updates as entries in the guess games, or equivalently the mode values of the posterior distributions recovered from LS′. In other words, Proposition 2 states that maximum likelihood updating is nested within LS′.

5.2 Learning strategy: the general form

The general form of the learning strategy extends the idea of LS′. The learning strategy assumes that a subject forms an initial prior characterized by a *beta*-distribution and updates her priors by imperfectly employing Bayes' rule. This means that a subject may under-react or over-react to new information compared to what Bayes' rule implies in belief updating. Of course, she may also perfectly employ Bayes' rule, then the learning strategy degenerates to its Bayesian special case. The updating rule can be modeled as follows:

$$\alpha_{i,n} = \alpha_{i,0} + \gamma_i^w \cdot k_{i,n} \tag{9}$$

$$\beta_{i,n} = \beta_{i,0} + \gamma_i^b \cdot (n - k_{i,n}) \tag{10}$$

$$\text{all } \alpha,\ \beta \geq 1; \quad \gamma_i^w, \gamma_i^b \geq 0; \quad n = 1, 2 \cdots, T; \quad i \in \{1, 2, \cdots, N\}$$

where $k_{i,0} = 0$. Equation (9) describes that subject i updates the belief parameter $\alpha_{i,n}$ by adding the (scaled) number of white draws to her initial $\alpha_{i,0}$. The number of white draws ($k_{i,n}$), before added to the initial prior, may be scaled by a multiplier: γ_i^w. In case of $\gamma_i^w = 1$, subject i responds to the white draws in the way that she adds exactly $k_{i,n}$ (without scaling) to her initial parameter $\alpha_{i,0}$ after observing n draws. This implies perfect employment of Bayes' rule (as in the Bayesian special case). In case of $\gamma_i^w < 1$, subject i responds to the white draws in the way that she adds a scaled-down $k_{i,n}$ to her initial $\alpha_{i,0}$. This means that she under-reacts to what Bayes' rule implies. In case of $\gamma_i^w > 1$, subject i responds to the white draws in the way that she adds a scaled-up $k_{i,n}$ to $\alpha_{i,0}$, implying an over-reaction to what Bayes' rule implies.

Analogously, Equation (10) describes that subject i updates the belief parameter $\beta_{i,n}$ by adding the (scaled) number of black draws to her initial $\beta_{i,0}$. $\gamma_i^b = 1$ implies subject i responds to black draws as Bayes' rule implies (as in the Bayesian special case), $\gamma_i^b < 1$ implies under-reaction compared to Bayes' rule, and $\gamma_i^b > 1$ implies over-reaction compared to Bayes' rule. Take note that both γ_i^w and γ_i^b are subject-specific and time-invariant. This implies that each subject has her own updating rule governed by the two parameters γ_i^w and γ_i^b, and her responses to a white draw/black draw is constant over time. Moreover, γ_i^w is not necessarily equal to γ_i^b, implying that i's responses to white draws may differ from her responses to black draws (i.e. with different scalings on $k_{i,n}$ and $n - k_{i,n}$). In case of $\gamma_i^w = 0$ ($\gamma_i^b = 0$), subject i never updates her beliefs when observing white (black) draws. In case of $\gamma_i^w = \gamma_i^b = 0$, i never learns at all: her initial belief distribution, characterized by the *beta*-distribution with shape parameters $(\alpha_{i,0},\ \beta_{i,0})$, perpetuates.

The mechanism of the learning strategy can also be interpreted in the following way: subject i starts with an initial prior which can be characterized by a *beta*-distribution with shape parameter bundle $(\alpha_{i,0},\ \beta_{i,0})$. After observing $k_{i,n}$ white draws out of n draws ($n \geq 1$), she acts *as if* she observes $\gamma_i^w \cdot k_{i,n}$ white draws, and updates her belief perfectly employing Bayes' rule. Analogously, after observing $(n - k_{i,n})$ black draws, she acts *as if* she observes $\gamma_i^b \cdot (n - k_{i,n})$ black draws, and updates her belief perfectly employing Bayes' rule. Thus, according to Proposition 2, her updated belief after n draws can also be characterized by a *beta*-distribution, with updated parameter bundle $(\alpha_{i,0} + \gamma_i^w k_{i,n},\ \beta_{i,0} + \gamma_i^b (n - k_{i,n}))$, or simply $(\alpha_{i,n},\ \beta_{i,n})$ as defined in Equation (9) and (10), respectively. In other words, the set $\{(\alpha_{i,n},\ \beta_{i,n}) : n = 0, 1, \cdots, T\}$ captures subject i's complete belief dynamics, and her entire belief distributions after each draw can be recovered from $(\alpha_{i,n},\ \beta_{i,n})$.

5.3 Parameter estimation

To recover a subject's belief dynamics under this learning strategy, it is sufficient to estimate the parameter set $\{\alpha_{i,0},\ \beta_{i,0},\ \gamma_i^w,\ \gamma_i^b\}$. Here, we directly present the regression equation for the parameter estimation. The deduction of this equation can be found in Appendix C. The regression equation reads:

$$\frac{white_n}{100} = M_n \left[\frac{1 - M_n}{1 - white_n/100} \right]^{\frac{\beta_0 + \gamma^b (n - k_n) - 1}{\alpha_0 + \gamma^w k_n - 1}} + \epsilon_n \tag{11}$$

$$\text{where} \quad M_n \equiv \frac{\alpha_0 + \gamma^w k_n - 1}{\alpha_0 + \gamma^w k_n + \beta_0 + \gamma^b (n - k_n) - 2} \tag{12}$$

where ϵ_n denotes the error term. For readability, all i subscripts (indexing subjects) are suppressed for the time being. In Equation (11), $white_n$ (guess game responses), n (number of draws), and k_n (number of white draws) are observable. All other parameters, $\alpha_0,\ \beta_0,\ \gamma^w$ and

γ^b, are to be estimated. To restrict $\alpha_0, \beta_0 \geq 1$ and $\gamma^w, \gamma^b \geq 0$, exponential functions are applied:

$$\alpha_0 = 1 + \exp(a) \tag{13}$$

$$\beta_0 = 1 + \exp(b) \tag{14}$$

$$\gamma^w = \exp(r^w) \tag{15}$$

$$\gamma^b = \exp(r^b) \tag{16}$$

We plug Equation (13)-(16) into Equation (11) and estimate parameters $\{a,\ b,\ r^w,\ r^b\}$ for each subject, using her responses in the guess games as well as the observed draw history. This means that, for each subject, at least 15 data entries are used to estimate four unknown parameters. The estimation applies the nonlinear regression method. Equation (11) does not allow $white_n = 100$. In practice, for subjects who respond 100 in the guess game, we reset $white_n = 99$. This only affects nine data entries (out of 1619) and seven subjects (out of 102).

With the estimated $\{\hat{a},\ \hat{b},\ \hat{r}^w,\ \hat{r}^b\}$, for each subject we recover the parameters $\{\hat{\alpha}_0,\ \hat{\beta}_0,$ $\hat{\gamma}^w,\ \hat{\gamma}^b\}$ based on Equation (13)-(16), respectively. For 99 out of 102 subjects, we obtain the estimation results from the regression. For three subjects, the regression is not possible, since there is no data variation within subjects. These three subjects never update their beliefs: two subjects constantly respond 50 in all guess games, and one subject constantly responds 70. For these three subjects, we manually parameterize their belief dynamics as follows: for the two subjects constantly responding 50, $\hat{\alpha}_{i,0} = \hat{\beta}_{i,0} \to +\infty$ and $\hat{\gamma}_i^w = \hat{\gamma}_i^b = 0$. For the subject constantly responding 70, $\hat{\alpha}_{i,0},\ \hat{\beta}_{i,0} \to +\infty$ and $(\hat{\alpha}_{i,0} - 1)/(\hat{\alpha}_{i,0} + \hat{\beta}_{i,0} - 2) = .7$, as well as $\hat{\gamma}_i^w = \hat{\gamma}_i^b = 0$. These parameterizations mean that, without observing any draw, they have already had very clear beliefs about the urn composition, and that they do not update their beliefs along the draws. In addition, such parameterizations also perfectly recover each elicited mode.

For each subject, using the results of $\{\hat{\alpha}_{i,0},\ \hat{\beta}_{i,0},\ \hat{\gamma}_i^w,\ \hat{\gamma}_i^b\}$, we can further recover one's entire belief distribution after each draw. We proceed in the following steps: first, we derive $\hat{\alpha}_{i,n}$ and $\hat{\beta}_{i,n}$ based on Equation (9)(10), respectively. Second, we discretize the support of the belief distributions, $\theta \in [0,1]$, to the set $\{\theta | \theta = 0, 0.01, 0.02, \cdots, 1\}$ and compute the probability mass (PMF) evaluated at each element of the set:

$$\text{pmf}(\hat{\alpha}_{i,n}, \hat{\beta}_{i,n}, \theta) = \begin{cases} \text{cdf}(\hat{\alpha}_{i,n}, \hat{\beta}_{i,n}, 0.005); & \text{if } \theta = 0 \quad (17) \\ \text{cdf}(\hat{\alpha}_{i,n}, \hat{\beta}_{i,n}, \theta + 0.005) \\ \qquad - \text{cdf}(\hat{\alpha}_{i,n}, \hat{\beta}_{i,n}, \theta - 0.005); & \text{if } \theta \in \{0.01, 0.02, \cdots, 0.99\} \quad (18) \\ 1 - \text{cdf}(\hat{\alpha}_{i,n}, \hat{\beta}_{i,n}, 0.995); & \text{if } \theta = 1 \quad (19) \\ 1/101; & \text{for all } \theta \text{ if } \hat{\alpha}_{i,n} = \hat{\beta}_{i,n} = 1 \quad (20) \end{cases}$$

where $\text{pmf}(\hat{\alpha}_{i,n}, \hat{\beta}_{i,n}, \theta)$ denotes the probability mass of a *beta*-distribution with $(\hat{\alpha}_{i,n}, \hat{\beta}_{i,n})$ evaluated at the support θ. Analogously, $\text{cdf}(\hat{\alpha}_{i,n}, \hat{\beta}_{i,n}, \theta)$ denotes the value of the cumulative distribution function. This discretization guarantees that the summation of the probability mass within one distribution is always equal to one. In fact, $\text{pmf}(\hat{\alpha}_{i,n}, \hat{\beta}_{i,n}, \theta)$ recovers subject i's entire belief distributions (in form of PMF) after each draw $(n = 0, 1, \cdots, T)$.

5.4 Estimation results

In this chapter, we discuss the estimation results of the parameters as computed based on Equation (17)-(20). In general, the nonlinear regressions fit the observed data very well. Figure 6 illustrates the frequency distribution of the R^2 values of the regressions across subjects. It shows that most regressions achieve an R^2 very close to one, implying excellent goodness of fit. More importantly, for most subjects, the belief updating dynamics recovered from the estimated parameters fit the observed guess game data (the elicited modes) very well. As an example, Figure 7 illustrates the belief dynamics of Subject 46. Each sub-graph represents her recovered belief distribution after a certain number of draws. A red vertical line, which represents her guess game response (i.e. the elicited mode value of the belief distribution, read from the X-axis), is added in each sub-graph. As can be seen, the mode of the estimated belief distribution is very close to the elicited mode in each sub-graph. This indicates that the belief estimation achieves excellent goodness of fit in terms of fitting the mode value. This consolidates the quality of the belief estimation.

A detailed picture of the parameter estimation results is given in Table 3 and Figure 8-9. Table 3 Panel A summarizes the results of $(\hat{\alpha}_{i,0},\ \hat{\beta}_{i,0})$, i.e. subjects' initial priors. Figure 8a and 8b plot $\hat{\beta}_{i,0}$ against $\hat{\alpha}_{i,0}$. In both diagrams, the bubble size represents the number of subjects. Ten subjects are estimated by $\hat{\alpha}_{i,0} = \hat{\beta}_{i,0} = 1$, represented by the largest bubble in both diagrams. Figure 8a represents the sub-samples with $\hat{\alpha}_{i,0},\ \hat{\beta}_{i,0} \leq 70$ (N=94). For better visibility, Figure 8b restricts to the sub-samples with $\hat{\alpha}_{i,0},\ \hat{\beta}_{i,0} \leq 2$ (N=59). The results show that 90 out of 102 subjects have $\hat{\alpha}_{i,0}$ and $\hat{\beta}_{i,0}$ which are both greater than one. This implies that nearly 90% of the subjects have bell-shaped distributed initial priors, a dominant majority among all shapes. Ten subjects have uniformly distributed initial priors, and two subjects have initial priors with increasing PDF representations. In terms of mode value of the initial prior, 52 subjects have a mode greater than 0.5, 38 subjects have a mode less than 0.5, and two subjects have a mode equal to 0.5. For the ten subjects with uniformly distributed priors, the mode is not specified, but the distribution is symmetric around 0.5. This implies that on average the initial prior distributions are slightly left-skewed, and more probability mass is assigned to scenarios with a proportion of white balls greater than 0.5.

The initial prior distributions recovered from $(\hat{\alpha}_{i,0},\ \hat{\beta}_{i,0})$ are depicted in Figure 8c (in form of PMF) and in Figure 8d (in form of CDF). Each curve represents one subject. As can be seen, the curves cover a variety of distributions: uniform distribution (the horizontal line), bell-shaped distributions, and extremely squeezed distributions (the vertical lines). In terms of mode, values less than 0.5, greater than 0.5, and around 0.5 are all observed. This implies that there exists large heterogeneity in the recovered initial prior distributions.

The parameters $\hat{\gamma}_i^w$ and $\hat{\gamma}_i^b$ govern subject i's updating rule, i.e. the response to white draws and black draws, respectively. The results are reported in Table 3 Panel B. Figure 9a and 9b plot $\hat{\gamma}_i^b$ against $\hat{\gamma}_i^w$. In either figure, the bubble size represents the number of subjects. Ten subjects are estimated by $\hat{\gamma}_i^w = \hat{\gamma}_i^b = 0$, represented by the largest bubble in both diagrams. For better visibility, Figure 9a and 9b restrict to the sub-samples with $\hat{\gamma}_i^w,\ \hat{\gamma}_i^b \leq 5$ (N=93) and $\hat{\gamma}_i^w,\ \hat{\gamma}_i^b \leq 1$ (N=79), respectively. A $\hat{\gamma}_i^w$ $(\hat{\gamma}_i^b)$ not equal to one indicates that subject i deviates from Bayes' rule in the way she responds to the white (black) draws. The results show that no subject perfectly follows Bayes' rule in belief updating. A large proportion of subjects under-react to the draw information in comparison to what Bayes' rule implies. In fact, 75% of the subjects have $\hat{\gamma}_i^w < 1$ and $\hat{\gamma}_i^b < 1$. In addition, subjects tend to respond to white draws and black draws symmetrically. This is supported by both diagrams, in which most bubbles are located close to the 45-degree line. This implies that, on average, subjects tend to react to good news and bad news nearly symmetrically in belief updating.

Table 3 Panel C reports the cross tabulation of subjects' initial priors and their responses to white draws. Analogously, Panel D reports the cross tabulation of subjects' initial priors and their responses to black draws. The numbers indicate that subjects who have bell-shaped initial priors most likely under-react to new information compared to what Bayes' rule implies. On the other hand, most subjects who have a uniformly distributed initial prior do not learn at all along the draws.

Using the estimated parameters, we recover the belief updating dynamics of each subject (as Figure 7 for each subject). To understand one's belief updating dynamic more deeply, we proceed to test three hypotheses closely related to subjects' initial belief characterizations and their updating rules.

6 Hypothesis testing

In the previous sections, for each subject, we use her elicited mode values reported in the guess games to estimate parameters which govern her belief formation and belief updates. Consequently, we characterize, for each subject, her initial prior distribution and her updating rule. Such characterization enables us to test the following three hypotheses:

(a) **Objective equality hypothesis**: in a completely ambiguous environment, all possible scenarios are considered equally likely;

(b) **Unbiased belief hypothesis**: in a completely ambiguous environment, one's belief demonstrates no bias towards either bad or good outcomes; and

(c) **Bayesian updating hypothesis**: subjects employ Bayes' rule when updating their beliefs.

Hypothesis (a) and Hypothesis (b) both, in different ways, relate to the shape of the initial prior distribution. While Hypothesis (a) relates to the dispersion of the initial prior distribution, Hypothesis (b) relates to the symmetry of the initial prior distribution. Hence, both hypotheses can be tested by scrutinizing the characterizations of subjects' initial prior distributions, which are governed by $\hat{\alpha}_{i,0}$ and $\hat{\beta}_{i,0}$. Hypothesis (c) is related to subjects' belief updating rules. The estimated parameters which govern the updating rule (i.e. $\hat{\gamma}_i^w$ and $\hat{\gamma}_i^b$) can be used to test this hypothesis. To proceed, we test hypothesis (a)-(c) sequentially.

6.1 Objective equality hypothesis

Based on the parameter estimation results of each subject, we examine the *objective equality hypothesis*, which claims that in a completely ambiguous environment, the initial prior is uniformly distributed, i.e. all possible scenarios are considered equally likely to realize. Table 4 Panel A summarizes the initial prior distributions across subjects. Subjects are categorized into groups according to the shapes of their initial prior distributions. The "without significance test" columns report the grouping based on the results of $\hat{\alpha}_{i,0}$ and $\hat{\beta}_{i,0}$ as reported in Table 3 Panel A: ten subjects have a uniformly distributed initial prior (Row 1), two subjects have initial priors with increasing PDF representations (Row 3), and 87 subjects have bell-shaped initial prior distributions (Row 4). The three subjects who never update their beliefs are estimated by some extremely squeezed initial prior distributions (Row 5). No subjects have an initial prior with a decreasing PDF representation (Row 2).

The results show that 87 out of 102 subjects are found to have bell-shaped distributed initial priors, a dominating majority. Three subjects have extremely squeezed distributions, which can also be seen as special cases of bell-shaped distributions. Therefore, for nearly 90% of the subjects, their initial priors are characterized by bell-shaped distributions. In contrast, only ten subjects form a uniformly distributed initial prior, accounting for less than 10% of the sample. These results indicate that most subjects have non-uniform initial priors.

To examine the *objective equality hypothesis* more rigorously, for each subject we test whether $H_0 : \alpha_{i,0} = 1$ & $\beta_{i,0} = 1$ can be rejected at 5% in a joint test. In other words, we test whether the *objective equality hypothesis* which claims uniformly distributed initial priors can be rejected. The last two columns of Table 4 Panel A report the grouping based on the test results. The results show that for 71 subjects (70% of the sample), the *objective equality hypothesis* can be rejected. This can also be seen in Figure 8a and 8b: the blue solid bubbles, which represent subjects with non-uniformly distributed initial priors, account for a dominating proportion in both diagrams. In conclusion, most subjects tend not to believe that all scenarios are equally likely, even in a completely ambiguous environment.

6.2 Unbiased belief hypothesis

In the previous section, we find that beliefs are not such that all possible scenarios are considered equally likely, although they are objectively equally likely according to the information available. Rather, subjects tend to believe that certain scenarios are more likely than others. In the next step, we investigate which particular scenarios are considered more likely and whether beliefs are biased towards good or bad scenarios.

The *unbiased belief hypothesis* claims that one's belief demonstrates no bias towards either bad or good scenarios. Rejection of the hypothesis implies pessimism or optimism. In this context, pessimism describes a tendency that a subject has beliefs that are biased towards bad scenarios in an ambiguous environment, while optimism describes a bias towards good scenarios. In the analysis, we focus on pessimism/optimism/unbiasedness in initial beliefs, i.e. beliefs in a completely ambiguous environment, before any information has been revealed that could rule out certain scenarios.

One way to diagnose pessimism or optimism in one's initial belief is to compare one's initial belief, represented by the mode value of the initial prior, with an unbiased benchmark. Any negative deviation from the unbiased benchmark is diagnosed as pessimism, and any positive deviation as optimism. One's initial mode value is simply $white_{i,0}$, the response of subject i in guess game G_0. Figure 2 illustrates the frequency distribution of $white_{i,0}$ across 102 subjects. A natural benchmark for an unbiased belief is given by the mid-point, i.e. 50. Values above 50 indicate optimism, while values below 50 indicate pessimism. The mean value of the initial modes across 102 subjects is equal to 49.64 (N=102, sd.=11.08, min=23, max=70). Among the 102 subjects, 34 subjects choose $white_{i,0} < 50$, 42 subjects choose $white_{i,0} = 50$, and 26 subjects choose $white_{i,0} > 50$. It appears that more subjects display pessimism. However, the t-test shows that $white_{i,0}$ is not significantly different from 50 ($H_0 : white_{i,0} = 50$, p-value=0.742). For the 34 subjects with $white_{i,0} < 50$, the mean value of $white_{i,0}$ is equal to 38.13, with median equal to 40. The lowest observed value is 23. It seems that pessimism in the initial beliefs is discernible, but there is no extreme pessimism. Analogously, for the 26 subjects with $white_{i,0} > 50$, the mean value of $white_{i,0}$ is equal to 64.12, with median equal to 63. The highest observed value is 77. Optimism seems to be discernible, but again there is no extreme optimism. Yet, pessimism and optimism tend to display similar magnitudes in initial beliefs. In conclusion, subjects on average seem to have unbiased initial beliefs, although heterogeneity exists among individual subjects, with some pessimism and some optimism.

An alternative way to diagnose the pessimism or optimism in one's initial belief is to investigate the shapes of the recovered initial prior distributions. The benchmark is rescaled from 50 in $[0, 100]$ to 0.5 in $[0, 1]$. In essence, the benchmark is identical as above. For each subject, the shape of her initial prior distribution is governed by $\hat{\alpha}_{i,0}$ and $\hat{\beta}_{i,0}$. In case of $\hat{\alpha}_{i,0} = \hat{\beta}_{i,0}$, the distribution is symmetric, meaning that one equally weighs bad outcomes and good outcomes. This is categorized as an unbiased belief. In case of $\hat{\alpha}_{i,0} < \hat{\beta}_{i,0}$, the distribution is right-skewed, meaning that one assigns higher likelihood to bad outcomes than to good outcomes. This is categorized as pessimism. In case of $\hat{\alpha}_{i,0} > \hat{\beta}_{i,0}$, the distribution is left-skewed, meaning that one assigns higher likelihood to good outcomes than to bad outcomes. This is categorized as optimism. Among the three subjects who never update their beliefs, the two who choose $white_{i,0} = 50$ are categorized as having unbiased beliefs, and the one who chooses $white_{i,0} = 70$ is categorized as having optimistic beliefs. Table 4 Panel B summarizes the categorization. The results reported in the "without significance test" columns are based on the estimated parameters ($\hat{\alpha}_{i,0}$ and $\hat{\beta}_{i,0}$). Twelve out of 102 subjects have unbiased initial beliefs. We observe more cases of optimism (52 subjects) than cases of pessimism (38 subjects). Yet the difference is relatively small, accounting for 13% of the sample.

To examine the *unbiased belief hypothesis* more rigorously, for each subject we test whether $H_0 : \alpha_i = \beta_i$ can be rejected at 5%. In case that the null hypothesis can be rejected, the initial prior is considered significantly biased: significant pessimism is diagnosed in case of $\hat{\alpha}_{i,0} < \hat{\beta}_{i,0}$ or significant optimism is diagnosed in case of $\hat{\alpha}_{i,0} > \hat{\beta}_{i,0}$. The last two columns of Table 4 Panel B report the grouping based on the test results. The results show that for 66 subjects (nearly 66% of the sample), the *unbiased belief hypothesis* cannot be rejected. This implies that 66% of the sample tend to have unbiased initial beliefs. In contrast, 13% of the sample display pessimism in initial beliefs, while 23% display optimism. These results imply that, on an individual level, the *unbiased belief hypothesis* cannot be rejected for most subjects. Most subjects tend to have unbiased initial beliefs, with slightly more cases of optimism than cases of pessimism.

An additional diagnosis is to compare the average probability mass assigned to $\theta < 0.5$ (all scenarios with winning probability less than 0.5: bad scenarios) with the average probability mass assigned to $\theta > 0.5$ (all scenarios with winning probability greater than 0.5: good scenarios) of all subjects. The results (not shown in the table) show that the probability mass for $\theta < 0.5$ is equal to 0.455 (subject-average), while the probability mass for $\theta > 0.5$ is equal to 0.513 (subject-average). The difference, however, is insignificant: using the overall sample data, the null hypothesis that the probability mass for $\theta < 0.5$ is equal to the probability mass for $\theta > 0.5$ cannot be rejected (two-sided t-test, P-value=0.103, N=102). This implies that on average subjects' initial beliefs display neither pessimism nor optimism, and the *unbiased belief hypothesis* cannot be rejected on a sample-population level.

The three diagnoses above lead to the same conclusion: on the level of the entire sample-population, subjects on average tend to have unbiased initial beliefs, and the *unbiased belief hypothesis* cannot be rejected. On an individual level, unbiased initial beliefs still dominate, but heterogeneity exists: both pessimistic and optimistic initial beliefs are observed among subjects (at discernible but not extreme degrees). Our findings are similar to those in Ahn et al. (2014), who also study pessimism/optimism under ambiguity. The authors argue that more than 70% of the studied subjects are classified as neutral (equivalent to the definition of "unbiased" in our paper), 10% as pessimistic, and 14% as optimistic. It is obvious that subjects with unbiased beliefs account for a dominating proportion in their paper. This is similar to our finding that subjects with unbiased initial beliefs account for the largest proportion.

6.3 Bayesian updating hypothesis

So far, we have investigated initial beliefs. We now turn to the question how these beliefs are updated as new information becomes available. We examine the *Bayesian updating hypothesis*, which claims that a subject employs Bayes' rule when updating beliefs. Table 4 Panel C summarizes the estimated updating rules, governed by $\hat{\gamma}_i^w$ and $\hat{\gamma}_i^b$, across subjects. $\hat{\gamma}_i^w = \hat{\gamma}_i^b = 1$ indicates perfect Bayesian updating. Other cases indicate imperfect Bayesian updating. The "without significance test" columns report the grouping based on the estimated parameter values. In total, 99 subjects, a dominant proportion of the sample, are found to update beliefs by imperfectly employing Bayes' rule (Row 2). Three subjects never update their beliefs in the experiment (Row 3). No subject is found to be perfect Bayesian updaters (Row 1). These numbers speak against the *Bayesian updating hypothesis*. The fact that $\hat{\gamma}_i^w$ and $\hat{\gamma}_i^b$ are mostly lower than one implies that subjects mostly under-react in response to both white draws and black draws compared to what Bayes' rule suggests.

To examine the *Bayesian updating hypothesis* more rigorously, we test whether $H_0 : \gamma_i^w = 1$ & $\gamma_i^b = 1$ can be rejected at 5% in a joint test. In case that the null hypothesis can be rejected, the updating rule is considered to deviate significantly from Bayes' rule. The last two columns of Table 4 Panel C report the grouping based on the test results. The results show that 83 subjects of those who employ imperfect Bayesian updating, significantly deviate from Bayes' rule. Included the three subjects who never update beliefs, for 84% of the sample, the *Bayesian updating hypothesis* can be rejected. Only 16% of the sample do not significantly deviate from Bayes' rule, an obvious minority. These results can be more directly seen in Figure 9a and 9b: the blue solid bubbles, which represent the imperfect Bayesian updaters, are the majorities in both diagrams. Among these imperfect Bayesian updaters, most seen are subjects who under-react to new information compared to what Bayes' rule implies, since most blue solid bubbles are below $\hat{\gamma}_i^w = 1$ and $\hat{\gamma}_i^b = 1$.

Overall, we find that initial beliefs do not follow a uniform distribution, i.e. some scenarios are considered more likely than others. Moreover, there does not seem to be a dominant bias towards good or bad scenarios. Rather, higher probability is assigned to intermediate scenarios. Finally, we find that subjects mostly do not apply Bayes' rule for belief updating. They rather perform imperfect Bayesian updating, mostly under-reacting to new information compared to what Bayes' rule implies.

To conclude this section, we discuss an important observation in the light of the three hypothesis tests: the heterogeneity of initial beliefs and belief updating rules across subjects. The evidence used to test the first two hypotheses clearly shows that there exists large heterogeneity in the characterization of subjects' initial beliefs. Such heterogeneity is reflected in the variety of shapes of the initial belief distributions, for instance with respect to dispersion. The heterogeneity of initial beliefs is also discernible in the symmetry of the initial belief distributions, indicating whether a subject is pessimistic, optimistic, or has an unbiased initial belief. All three cases are clearly observed in our sample. At last, the evidence used to test the third hypothesis also shows that belief updating rules are also heterogeneous across subjects: the degree to which subjects stick to current beliefs or follow Bayes' rule varies substantially across subjects. All in all, it can be concluded that heterogeneity prevails along both belief formation and belief updating.

7 Conclusion

This paper manages to distinguish beliefs from attitudes in situations involving ambiguity. The initial belief formation and the belief updating process are directly tracked down by a simple and clear-cut experiment design. Using the elicited mode of each belief distribution, we recover all belief distributions with full characterizations for all subjects along the learning process. The results show that, for 70% of the subjects, their initial belief distributions are non-uniform. This finding supports the rejection of the *objective equality hypothesis* that one's initial belief follows a uniform distribution, i.e. that one conceives each situation in an ambiguous environment equally likely. The modes of the estimated initial belief distributions locate, for most subjects, not far away from the mid-point of the distribution: 0.5. In addition, the initial belief distributions are mainly symmetric around the mid-point 0.5. Both findings imply that subjects assign almost equal probability masses to good scenarios as to bad scenarios. This supports the *unbiased belief hypothesis*. As for subjects' belief updating rules, for 84% of the subjects, the *Bayesian updating hypothesis* can be rejected, implying that they deviate from Bayes' rule when updating beliefs.

The findings in the paper reveal that beliefs play an important role in analyses related to ambiguity. The rejection of the *objective equality hypothesis* implies that beliefs matter. Disregarding beliefs in studies on ambiguity aversion may lead to wrong estimations for attitudes towards ambiguity. Moreover, it is not enough to simply assume a prominent belief distribution (e.g. a uniform or any other distribution). The heterogeneity of beliefs found in this paper indicates that there exists no single belief distribution that fits for all subjects. Such heterogeneity exists not only in initial beliefs, but in belief updating rules as well. The rejection of the *Bayesian updating hypothesis* hints that most subjects do not employ Bayes' rule when updating beliefs. Belief estimation which directly assumes the employment of Bayes' rule may build up an erroneous basis. As a result, all further analyses built on this erroneous belief basis (e.g. attitude estimation) may be biased. The heterogeneity in terms of initial beliefs and belief updating consolidates the necessity of belief analysis as an independent part from attitude analysis in settings involving ambiguity. Ignorance of the possible effects of beliefs on one's decisions may lead to biased results when analyzing decision-making under ambiguity.

References

Abdellaoui, M., Baillon, A., Placido, L., and Wakker, P. P. (2011). The rich domain of uncertainty: Source functions and their experimental implementation. *American Economic Review*, 101(2):695–723.

Ahn, D., Choi, S., Gale, D., and Kariv, S. (2014). Estimating ambiguity aversion in a portfolio choice experiment. *Quantitative Economics*, 5(2):195–223.

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

Baillon, A., Huang, Z., Selim, A., and P. Wakker, P. (2018). Measuring ambiguity attitudes for all (natural) events. *Econometrica*, 86:1839–1858.

Binmore, K., Stewart, L., and Voorhoeve, A. (2012). How much ambiguity aversion? *Journal of Risk and Uncertainty*, 45(3):215–238.

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.

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.

Coffman, K. B., Collis, M., and Kulkarni, L. (2019). Stereotypes and belief updating. Working paper no. 19-068, Harvard Business School.

Cubitt, R., van de Kuilen, G., and Mukerji, S. (2018). The strength of sensitivity to ambiguity. Technical Report 3.

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.

Eil, D. and Rao, J. M. (2011). The good news-bad news effect: Asymmetric processing of objective information about yourself. *American Economic Journal: Microeconomics*, 3(2):114–138.

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

Ertac, S. (2011). Does self-relevance affect information processing? experimental evidence on the response to performance and non-performance feedback. *Journal of Economic Behavior & Organization*, 80(3):532–545.

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. (2008). Objective and Subjective Rationality in a Multiple Prior Model. Carlo Alberto Notebooks 73, Collegio Carlo Alberto.

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.

Gollier, C. (2011). Portfolio choices and asset prices: The comparative statics of ambiguity aversion. *The Review of Economic Studies*, 78(4):1329–1344.

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

Holt, C. A. and Laury, S. K. (2002). Risk aversion and incentive effects. *The American Economic Review*, 92(5):1644–1655.

Hurley, T. M. and Schogren, J. F. (2005). An experimental comparison of induced and elicited beliefs. *Journal of Risk and Uncertainty*, 30(2):169–188.

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.

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

Mobius, M., Niederle, M., Niehaus, P., and Rosenblat, T. (2014). Managing self-confidence.

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: Experiment procedure

This table reports the main procedures of the experiment design. All procedures that show up in at least one of the experiment sessions are included. Some procedures may show up only in some sessions, instead of all sessions. The procedures are implemented from top to bottom, and from left to right as in the table. The main structure and time-line of the design sustains for all sessions. n denotes the number of draws already implemented at the time the decisions are made.

	n=0	**n=1**	$\cdots$	**n=14**	**n=15**	
Choice games	G_0	G_1	$\cdots$	G_{14}	G_{15}	Other choice games
	↓	↓		↓		↓
	Asset Trading	Asset Trading	$\cdots$	Asset Trading		Choice game resolution
	↓	↓				Questionnaire
	1^{st} draw*	2^{nd} draw	$\cdots$	15^{th} draw		Payment

*Draw: in each draw, one ball is drawn out from the ambiguous urn, its color is displayed, and then the ball is put back into the urn (draw with replacement). In Session I-VI, each market implements its own draw. In Session VII, each subject implements its own draw. The ambiguous urn design is identical universally in all sessions.

n denotes the number of draws already implemented. G_n denotes the guess game, in which subjects guess the number of white balls in the ambiguous urn after n draws with replacement are already observed.

Table 2: Number of subjects per Session/Market

This table reports the number of subjects, the played guess games, the draw history information, and the other included experimental parts, in each market of each session.

Session ID	**Market ID**	**Subject ID**	**No. of subjects**	**Guess games played**	**Draw history**	**Other parts included**
I	1	1-13	13	$G_0, G_1, \cdots, G_{14}$	Path 1	choice games, asset trading
II	2	14-20	7	$G_0, G_1, \cdots, G_{15}$	Path 2	choice games, asset trading
	3	21-27	7		Path 3	
III	4	28-34	7	$G_0, G_1, \cdots, G_{15}$	Path 4	choice games, asset trading
	5	35-41	7		Path 5	
IV	6	42-48	7	$G_0, G_1, \cdots, G_{15}$	Path 6	choice games, asset trading
	7	49-55	7		Path 7	
V	8	56-62	7	$G_0, G_1, \cdots, G_{15}$	Path 8	choice games, asset trading
	9	63-69	7		Path 9	
VI	10	70-76	7	$G_0, G_1, \cdots, G_{15}$	Path 10	choice games, asset trading
	11	77-83	7		Path 11	
VII	12	84-102	19	$G_0, G_1, \cdots, G_{15}$	Path 12-30	choice games
			Total=102		*Total*=30	

Table 3: Estimation results of the learning strategy

This table summarizes the parameter estimation results of the learning strategy. The learning strategy assumes that a subject starts with a *beta*-distributed initial prior and updates her belief by fully or partially reacting to what Bayes' rule implies. The *beta*-distributed initial prior is characterized by $\hat{\alpha}_{i,0}$ and $\hat{\beta}_{i,0}$, and the updating rule is governed by $\hat{\gamma}_i^w$ and $\hat{\gamma}_i^b$. They are estimated based on Equation (11). Panel A reports the estimated initial priors: $\hat{\alpha}_{i,0}$ and $\hat{\beta}_{i,0}$, summarized across all subjects (N=102). The two subjects who stick to 50 in all guess games are estimated by $\hat{\alpha}_{i,0} = \hat{\beta}_{i,0} \to +\infty$ (entering Row 4). The subject who sticks to 70 in all guess games is estimated by $\hat{\alpha}_{i,0}, \hat{\beta}_{i,0} \to +\infty$ and $\hat{\alpha}_{i,0} > \hat{\beta}_{i,0}$ (entering Row 5). Panel B reports the estimated updating rule: $\hat{\gamma}_i^w$ and $\hat{\gamma}_i^b$, summarized across all subjects (N=102). The three subjects who never update their beliefs are estimated by $\hat{\gamma}_i^w = \hat{\gamma}_i^b = 0$ (entering Row 7). Panel C (Panel D) reports the cross tabulation of subjects' initial priors and their responses to white (black) draws.

Panel A: Estimated initial prior: $\hat{\alpha}_{i,0}$ and $\hat{\beta}_{i,0}$

	count	%
(1) $\hat{\alpha}_{i,0} = \hat{\beta}_{i,0} = 1$: uniform distri.	10	9.80
(2) $\hat{\alpha}_{i,0} = 1$, $\hat{\beta}_{i,0} > 1$: decreasing PDF (mode=0)		
(3) $\hat{\alpha}_{i,0} > 1$, $\hat{\beta}_{i,0} = 1$: increasing PDF (mode=1)	2	1.96
(4) $\hat{\alpha}_{i,0} = \hat{\beta}_{i,0} > 1$: bell-shaped distri. (mode=0.5)	2	1.96
(5) $\hat{\alpha}_{i,0} > \hat{\beta}_{i,0} > 1$: bell-shaped distri. (mode$>$ 0.5)	50	49.02
(6) $\hat{\beta}_{i,0} > \hat{\alpha}_{i,0} > 1$: bell-shaped distri. (mode$<$ 0.5)	38	37.25
Total	102	100

Panel B: Estimated updating rule: $\hat{\gamma}_i^w$ and $\hat{\gamma}_i^b$

	$\hat{\gamma}_i^w$		$\hat{\gamma}_i^b$	
	count	%	count	%
(7) $\hat{\gamma} = 0$: full-stickiness to initial belief	12	11.76	13	12.75
(8) $0 < \hat{\gamma} < 1$: under-reaction to Bayes' rule	69	67.65	68	66.67
(9) $\hat{\gamma} = 1$: employment of Bayes' rule				
(10) $\hat{\gamma} > 1$: over-reaction to Bayes' rule	21	20.59	21	20.59
Total	102	100	102	100

Table 3: Estimation results of the learning strategy (Continued)

Panel C: Initial prior ($\hat{\alpha}_{i,0}$, $\hat{\beta}_{i,0}$) and response to white draws ($\hat{\gamma}_i^w$)

	$\hat{\gamma}_i^w = 0$	$0 < \hat{\gamma}_i^w < 1$	$\hat{\gamma}_i^w = 1$	$\hat{\gamma}_i^w > 1$	Total
(11) $\hat{\alpha}_{i,0} = \hat{\beta}_{i,0} = 1$	7	3			10
(12) $\hat{\alpha}_{i,0} = 1$, $\hat{\beta}_{i,0} > 1$					0
(13) $\hat{\alpha}_{i,0} > 1$, $\hat{\beta}_{i,0} = 1$				2	2
(14) $\hat{\alpha}_{i,0} = \hat{\beta}_{i,0} > 1$	2				2
(15) $\hat{\alpha}_{i,0} > \hat{\beta}_{i,0} > 1$	3	36		11	50
(16) $\hat{\beta}_{i,0} > \hat{\alpha}_{i,0} > 1$		30		8	38
Total	12	69	0	21	102

Panel D: Initial prior ($\hat{\alpha}_{i,0}$, $\hat{\beta}_{i,0}$) and response to black draws ($\hat{\gamma}_i^b$)

	$\hat{\gamma}_i^b = 0$	$0 < \hat{\gamma}_i^b < 1$	$\hat{\gamma}_i^b = 1$	$\hat{\gamma}_i^b > 1$	Total
(17) $\hat{\alpha}_{i,0} = \hat{\beta}_{i,0} = 1$	8	2			10
(18) $\hat{\alpha}_{i,0} = 1$, $\hat{\beta}_{i,0} > 1$					0
(19) $\hat{\alpha}_{i,0} > 1$, $\hat{\beta}_{i,0} = 1$				2	2
(20) $\hat{\alpha}_{i,0} = \hat{\beta}_{i,0} > 1$	2				2
(21) $\hat{\alpha}_{i,0} > \hat{\beta}_{i,0} > 1$	3	35		12	50
(22) $\hat{\beta}_{i,0} > \hat{\alpha}_{i,0} > 1$		31		7	38
Total	13	68	0	21	102

Table 4: Initial prior distribution, pessimism/optimism in initial belief, and updating rule

This table summarizes the initial prior and the updating rule of each subject's learning strategy. "Without significance test" columns report results based on the estimated parameters. Panel A reports subjects' shapes of the initial prior distributions: $\hat{\alpha}_{i,0} = \hat{\beta}_{i,0} = 1$ implies uniform distribution. In significance test, subjects for whom $H_0 : \alpha_{i,0} = 1$ & $H_0 : \beta_{i,0} = 1$ can be rejected at 5% are categorized as having non-uniform distributions. Panel B reports unbiasedness/pessimism/optimism in subjects' initial beliefs: $\hat{\alpha}_{i,0} = \hat{\beta}_{i,0}$ ($\hat{\alpha}_{i,0} < \hat{\beta}_{i,0}$, $\hat{\alpha}_{i,0} > \hat{\beta}_{i,0}$) implies unbiased initial belief (pessimism, optimism). In significance test, subjects for whom $H_0 : \alpha_{i,0} = \beta_{i,0}$ can be rejected at 5% are categorized as displaying pessimism (if $\hat{\alpha}_{i,0} < \hat{\beta}_{i,0}$) or optimism (if $\hat{\alpha}_{i,0} > \hat{\beta}_{i,0}$). Panel C reports subjects' updating rules: $\hat{\gamma}_i^w = \hat{\gamma}_i^b = 1$ implies perfect Bayesian. In significance test, subjects for whom $H_0 : \gamma_i^w = 1$ & $\gamma_i^b = 1$ can be rejected at 5% are categorized as imperfect Bayesian. In all three panels, if a subject has no significant test results, she belongs to the same category in all columns (with or without the significance test).

Panel A: initial prior distribution

	without significance test		$H_0 : \alpha_{i,0} = 1$ & $\beta_{i,0} = 1$ can be rejected at 5%	
	count	%	count	%
(1) uniform distribution	10	9.80	(31)*	30.39
(2) decreasing PDF				
(3) increasing PDF	2	1.96	2	1.96
(4) bell-shaped distribution	87	85.29	66	64.71
(5) extremely squeezed PDF	3	2.94	3	2.94
Total	102	100	102	100

*For 31 subjects, $H_0 : \alpha_{i,0} = 1$ & $\beta_{i,0} = 1$ cannot be rejected at 5%. 16 subjects have no significance test results. These include ten subjects with $\hat{\alpha}_{i,0} = \hat{\beta}_{i,0} = 1$ (uniform distribution), two subjects with $\hat{\alpha}_{i,0} > 1, \hat{\beta}_{i,0} = 1$ (increasing PDF), three subjects who never update beliefs, and one subject with extremely large $\hat{\alpha}_{i,0}, \hat{\beta}_{i,0}$. For these 16 subjects, the standard errors of $\hat{\alpha}_{i,0}$ and $\hat{\beta}_{i,0}$ are close to zero, and thus no test results are available. These 16 subjects belong to the same category in all columns.

Panel B: initial belief: unbiased belief, pessimism, or optimism

	without significance test		$H_0 : \alpha_{i,0} = \beta_{i,0}$ can be rejected at 5%	
	count	%	count	%
(1) unbiased initial belief	12	11.76	(66)$	65.71
(2) pessimism in initial belief	38	37.25	13	12.75
(3) optimism in initial belief	52	50.98	23	22.55
Total	102	100	102	100

$For 66 subjects, $H_0 : \alpha_{i,0} = \beta_{i,0}$ cannot be rejected at 5%. 16 subjects have no significance test results. These include ten subjects with $\hat{\alpha}_{i,0} = \hat{\beta}_{i,0} = 1$ (uniform distribution), two subjects with $\hat{\alpha}_{i,0} > 1, \hat{\beta}_{i,0} = 1$ (increasing PDF), three subjects who never update beliefs, and one subject with extremely large $\hat{\alpha}_{i,0}, \hat{\beta}_{i,0}$. For these 16 subjects, the standard errors of $\hat{\alpha}_{i,0}$ and $\hat{\beta}_{i,0}$ are close to zero, and thus no test results are available. These 16 subjects belong to the same category in all columns.

Table 4: Initial prior distribution, pessimism/optimism in initial belief, and updating rule (Continued)

Panel C: updating rule				
	without significance test		$H_0 : \gamma_i^w = 1$ & $\gamma_i^b = 1$ can be rejected at 5%.	
	count	%	count	%
(1) perfect Bayesian			(16)[§]	15.69
(2) imperfect Bayesian	99	97.06	83	81.37
(3) never update	3	2.94	3	2.94
Total	102	100	102	100

[§] For 16 subjects, $H_0 : \gamma_i^w = 1$ & $\gamma_i^b = 1$ cannot be rejected. 18 subjects have no significance test results. These include 14 subjects with $\hat{\gamma}_i^w = 0$ or $\hat{\gamma}_i^b = 0$, three subjects who never update beliefs, and one subject estimated by $\hat{\gamma}_i^w = 4.41$ and $\hat{\gamma}_i^b = 1.60$. For these 18 subjects, the standard errors of $\hat{\gamma}_i^w$ and $\hat{\gamma}_i^b$ are close to zero, and thus no test results are available. These 18 subjects belong to the same category in all columns.

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 an integer between zero and 100 (inclusive) to announce her guess about the proportion of white balls in the ambiguous urn. In order to permit learning, draws with replacement are implemented from the ambiguous urn as the source of new information. G_n denotes the guess game played after n draws are implemented ($0 \leq n \leq T$, where $T = 14$ in Session I, otherwise $T = 15$). The guess games are played in the following sequence. In Session I: G_0, the first draw, G_1, the second draw, $\cdots$, G_{14}, the 15^{th} draw. In all other sessions: G_0, the first draw, G_1, the second draw, $\cdots$, G_{14}, the 15^{th} draw, *plus* G_{15}. 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 (G_0, the guess game before any draw is implemented); As an example, Figure (b) shows the screen display of the guess game after five draws are implemented (G_5).

(a) Guess game before any draw: G_0

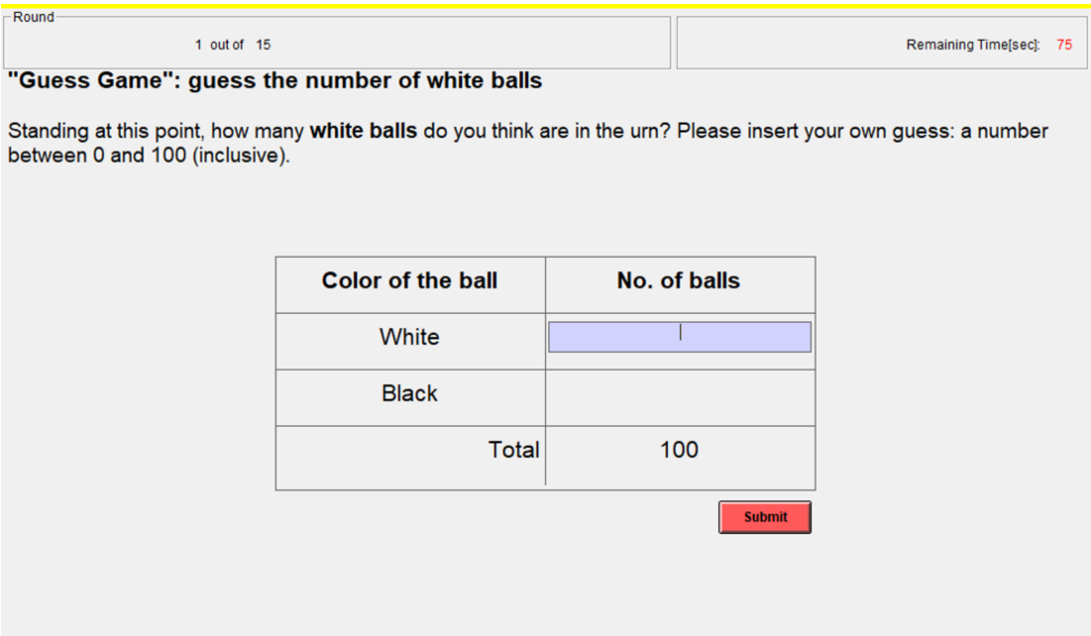

(b) Guess game after five draws: G_5

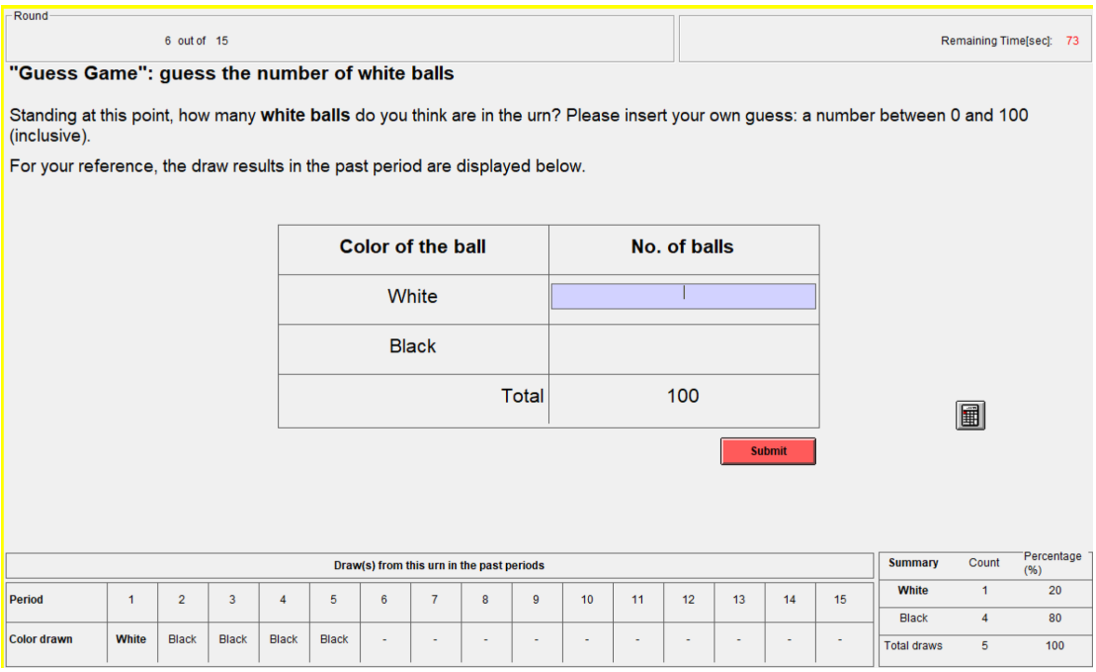

Figure 2: Response of the first guess game: $white_{i,0}$

This diagram illustrates frequency distribution of subjects' responses in the first guess game G_0: $white_{i,0}$. In discussions of pessimism/optimism/unbiased belief, the mid-point 50 is chosen as the neutral benchmark. $white_{i,0} = 50$ indicates that subject i has an unbiased initial belief, $white_{i,0} < 50$ indicates pessimism in i's initial belief, and $white_{i,0} > 50$ indicates optimism in i's initial belief. Sample size: N=102.

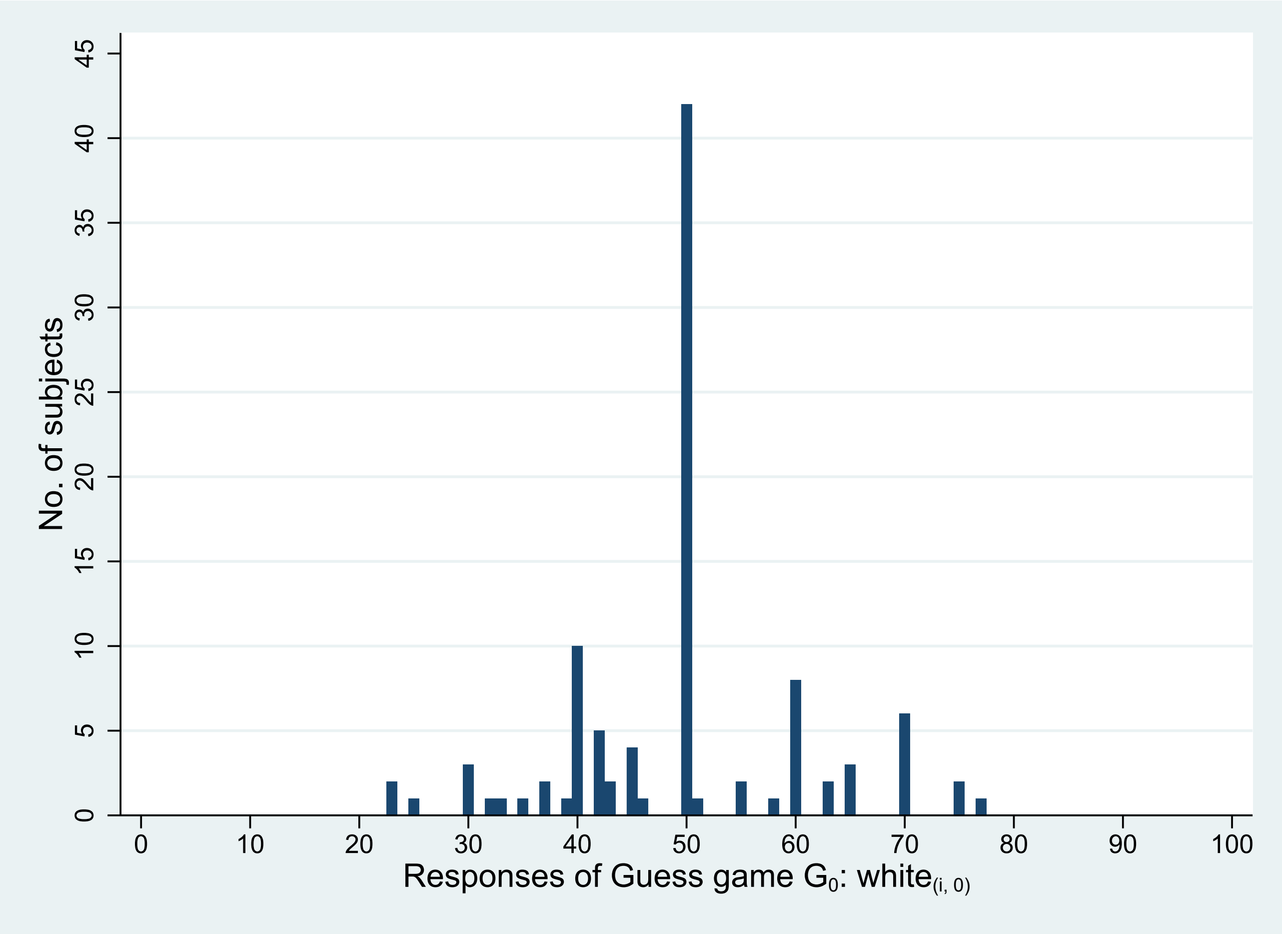

Figure 3: Guess game responses

This figure illustrates the draw histories and the guess game responses. In total, 30 paths of draw history are generated. Each sub-graph represents one draw history, independently implemented using an identical ambiguous urn. A white column represents that a white draw is observed; A gray column, a black draw. The blue dot-lines represent the guess game response in each guess game ($G_0, \cdots, G_{14}$ for sub-graph 1; $G_0, \cdots, G_{15}$ for others). Sub-graphs 1-11 represent the draw history in Markets 1-11, respectively. Markets 1-11 produce eleven different paths of draw history, and subjects in the same market observe the same draw history. The guess game responses (blue dot-line) in sub-graphs 1-11 are thus market-wide average values, computed by Equation (1). The sample size for sub-graphs 1-11 are reported in Table 2, i.e. No. of subjects in Markets 1-11, respectively. Sub-graphs 12-30 represent the 19 subjects in Market 12, respectively. In Market 12, draws are implemented per subject, i.e. Market 12 produces 19 different paths of draw history, one for each subject in Market 12. Each blue dot-line in sub-graphs 12-30 represents the guess game responses of one specific subject in Market 12. A red reference line of 50 is added.

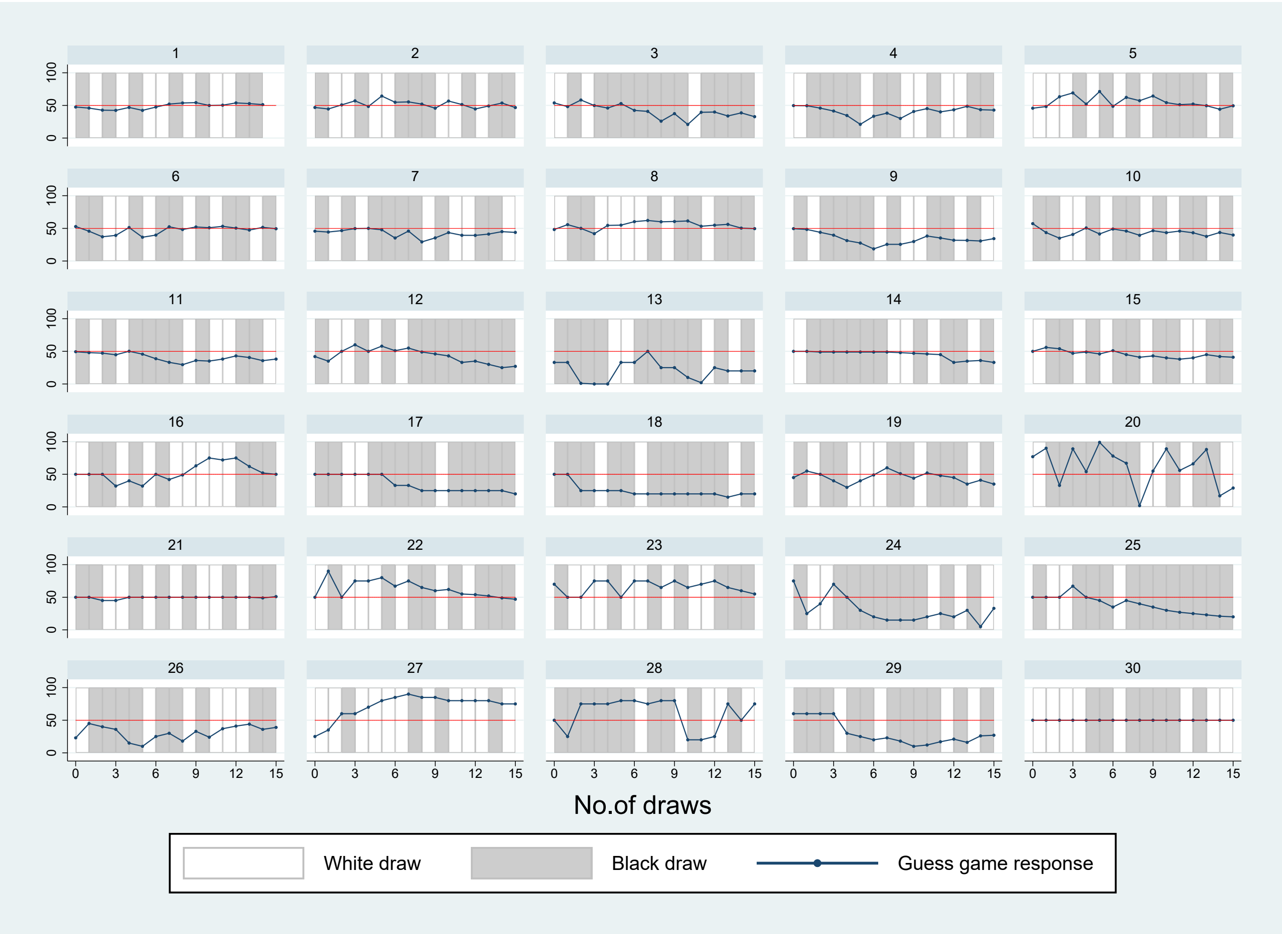

Figure 4: Rate of rational belief updates

This diagram illustrates the rate of rational belief updates (i.e. guess game response) by subject. A rational belief update is defined as such: if a white draw is observed, a subject does not downwards adjust her belief; if a black draw is observed, a subject does not upwards adjust her belief. The rate of rational belief updates of a given subject is the ratio of the number of her rational belief update(s) to the total number of guess games she plays, except the first guess game. Sample size: N=102.

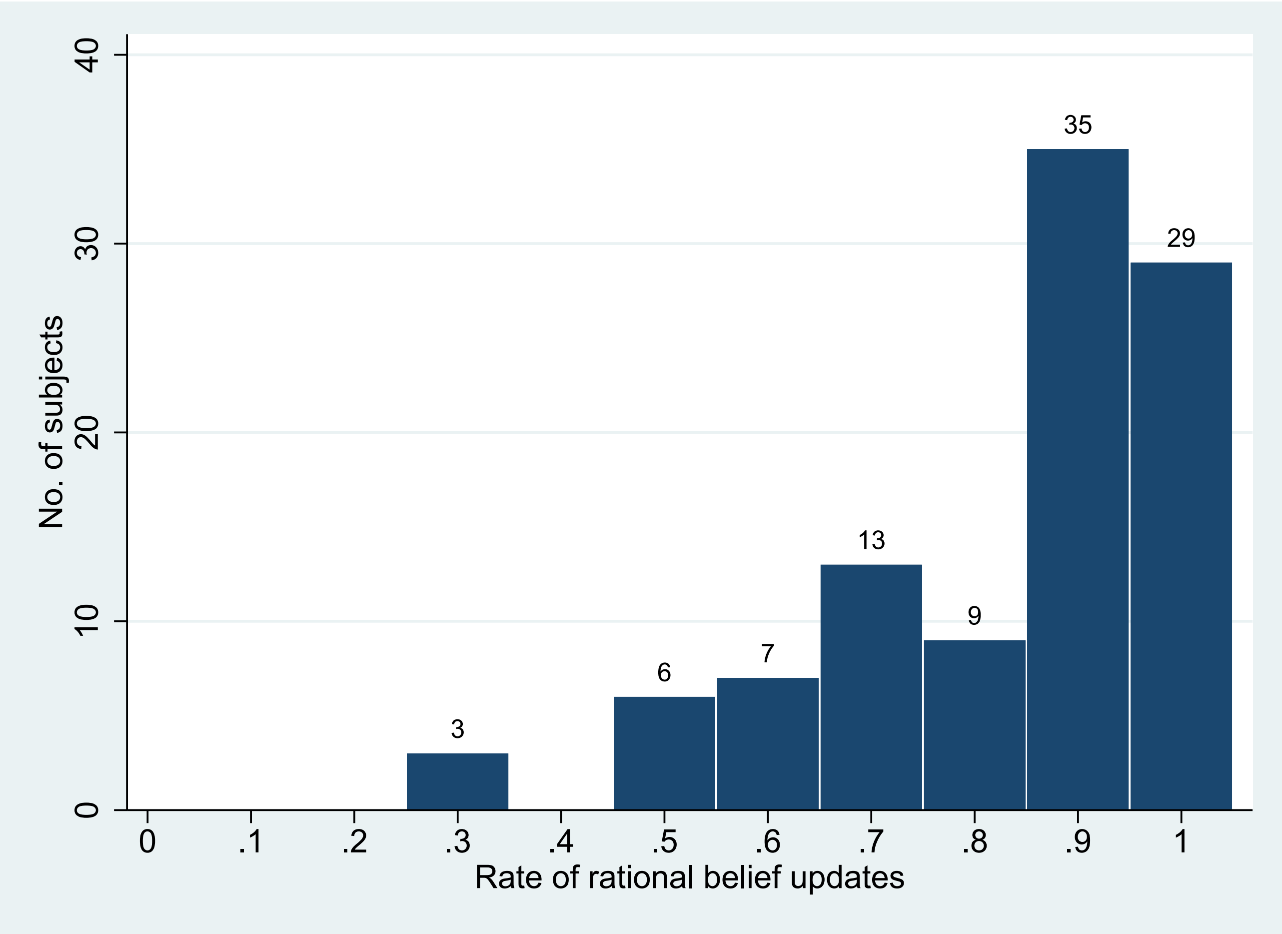

Figure 5: Examples of *beta*-distribution: PDF

This diagram illustrates the PDF 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 $(\alpha, \beta) = (1,1)$

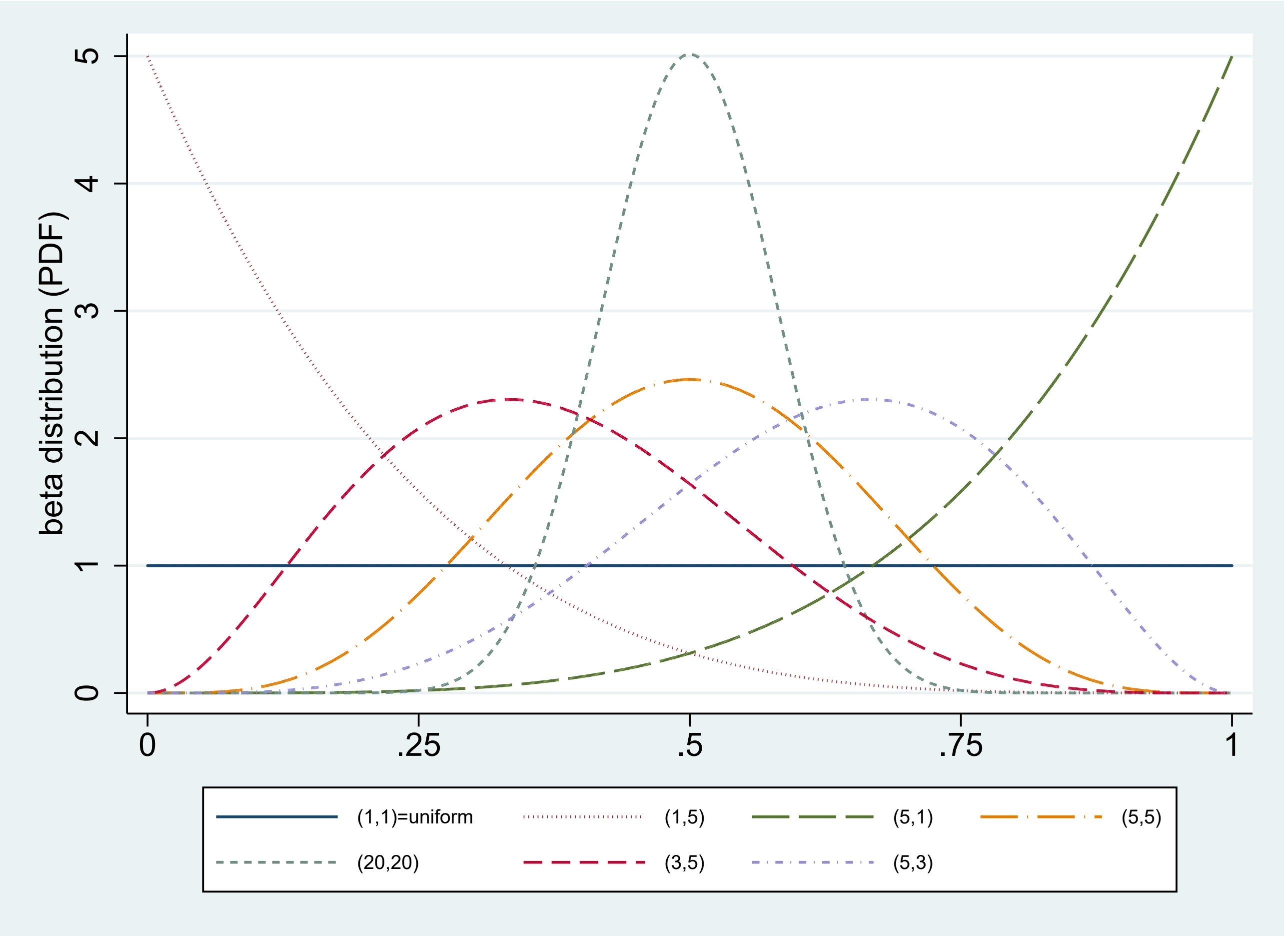

Figure 6: Goodness of fit: the learning strategy

This diagram reports the frequency distribution of the R^2 values (the goodness of fit of the regression based on Equation 11) across subjects. R^2 ranges from zero to one. The higher the value is, the better the learning strategy fits the observed data. Sample size: N=102.

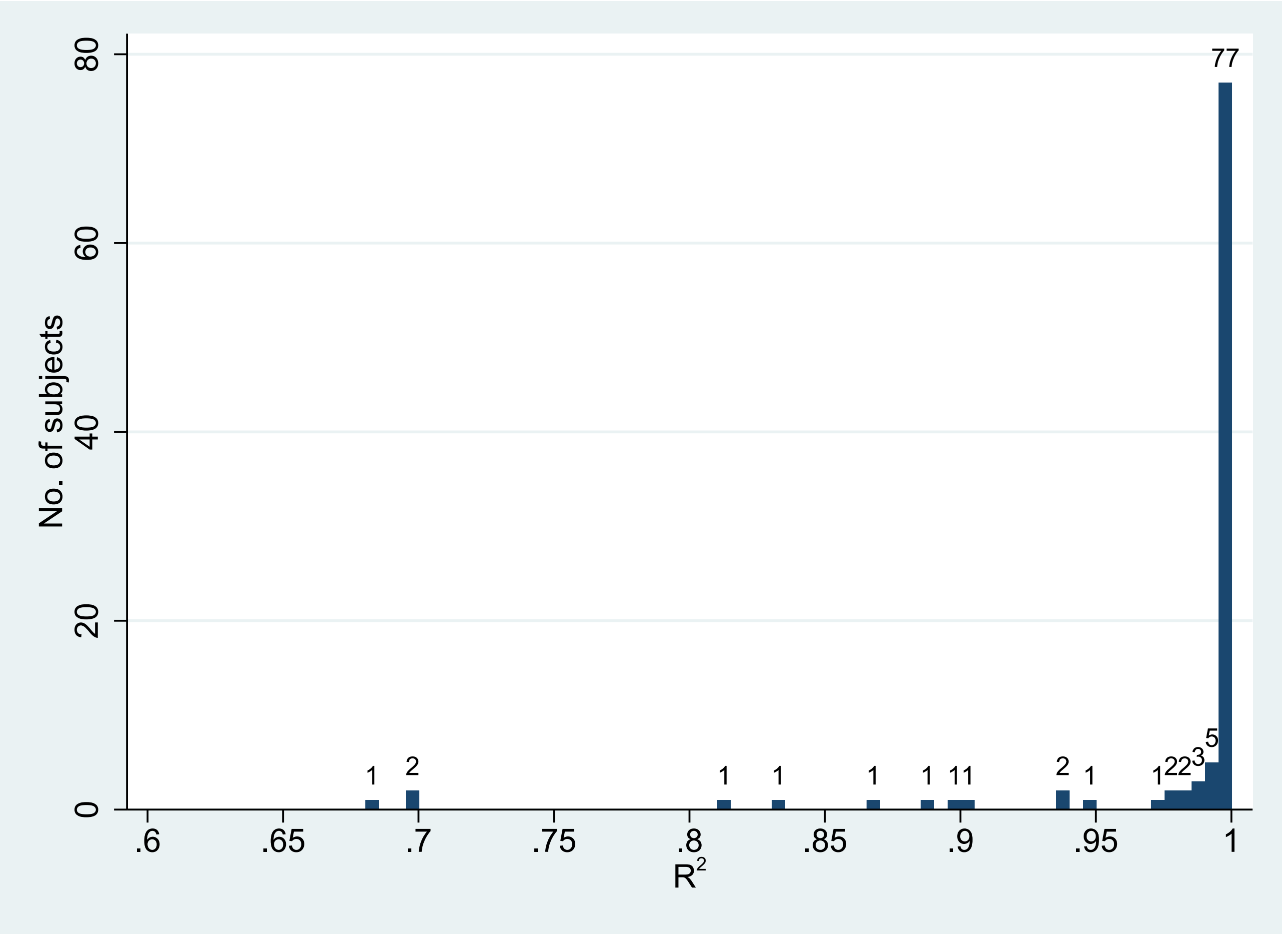

Figure 7: Estimated belief dynamic of Subject 46: an example

This diagram illustrates the belief updating dynamic recovered from the estimated parameters $\hat{\alpha}_i$, $\hat{\beta}_i$, $\hat{\gamma}_i^w$ and $\hat{\gamma}_i^b$ of Subject 46, as an example. For reference, the red vertical line represents the elicited modes (i.e. guess game responses), read from the X-axis. Each sub-graph represents the belief distribution after a certain number of draws (n).

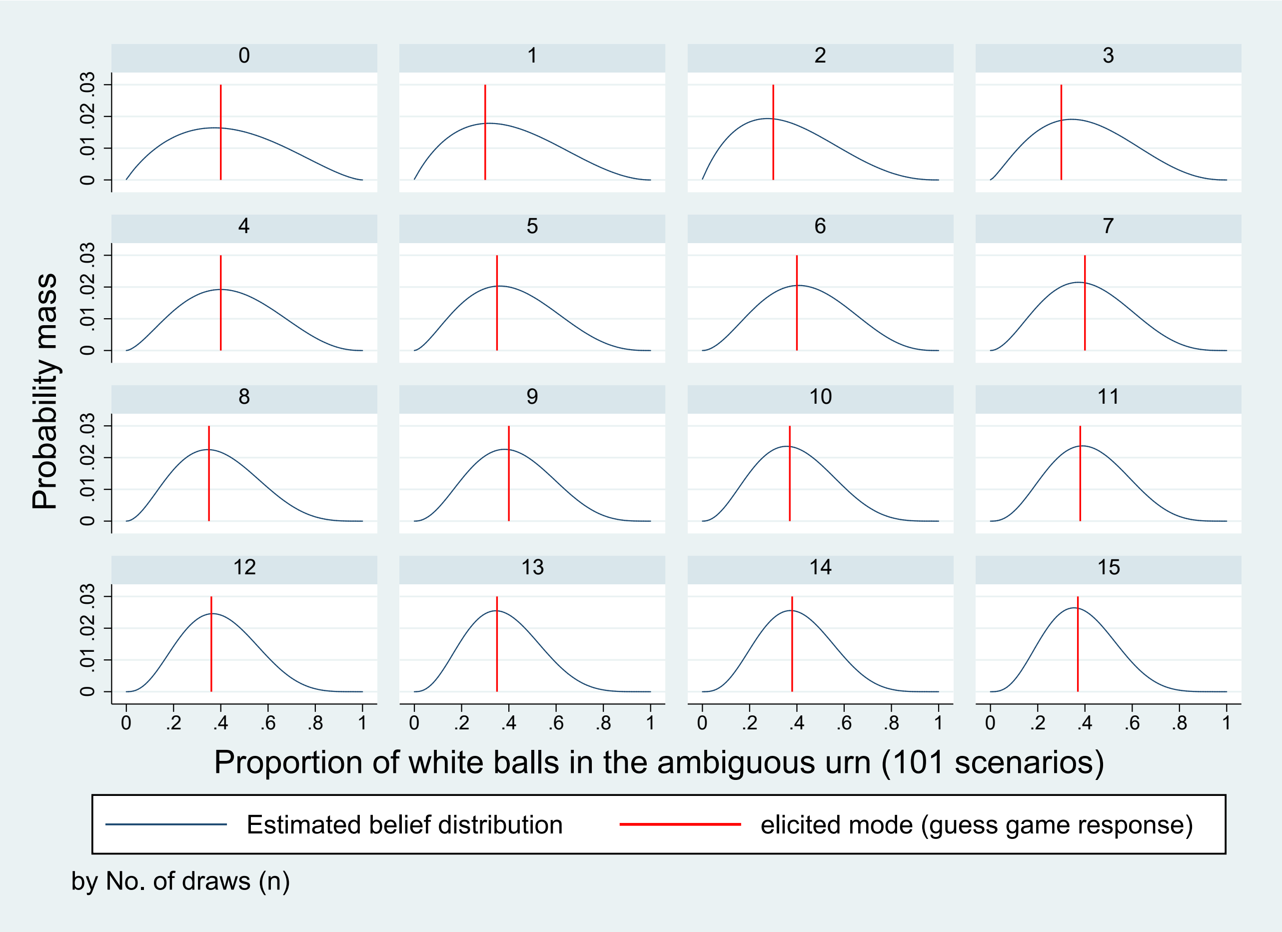

Figure 8: Estimated initial prior: $\alpha_{i,0}$ and $\beta_{i,0}$

These diagrams illustrate the estimation results of subjects' initial priors. The learning strategy assumes that a subject starts with a *beta*-distributed initial prior, and updates her beliefs by fully or partially reacting to what Bayes' rule implies. The shape of the initial prior is characterized by $(\alpha_{i,0},\ \beta_{i,0})$. They are estimated based on Equation (11). Diagram (a) plots $\hat{\beta}_{i,0}$ against $\hat{\alpha}_{i,0}$, for $\hat{\alpha}_{i,0}, \hat{\beta}_{i,0} \leq 70$. For better visibility, Diagram (b) restricts to the subsamples with $\hat{\alpha}_{i,0}, \hat{\beta}_{i,0} \leq 2$. The bubble size represents the number of subjects (the largest bubble represents ten subjects; the smallest bubble represents one subject). In Diagram (a)(b), the blue solid bubbles represent subjects for whom $H_0 : \alpha_{i,0} = 1 \ \&\ \beta_{i,0} = 1$ can be jointly rejected at 5% (significantly different from uniform initial prior). Otherwise, gray hollow bubbles. A 45-degree reference line is added in (a)(b). Diagram (c) illustrates the initial prior distributions in form of PMF, recovered from $(\hat{\alpha}_{i,0}, \hat{\beta}_{i,0})$ based on Equation (17)-(20). For visibility, the subject with $\hat{\alpha}_{i,0} = 65.5,\ \hat{\beta}_{i,0} = 1$ is excluded. Vertical lines represent subjects with extremely large $\hat{\alpha}_{i,0}$ and/or $\hat{\beta}_{i,0}$, i.e. assigning all probability mass to one scenario. Diagram (d) illustrates the initial prior distributions in form of CDF, recovered from $(\hat{\alpha}_{i,0}, \hat{\beta}_{i,0})$.

(a) $\hat{\alpha}_{i,0}, \hat{\beta}_{i,0} \leq 70$ (N=94)

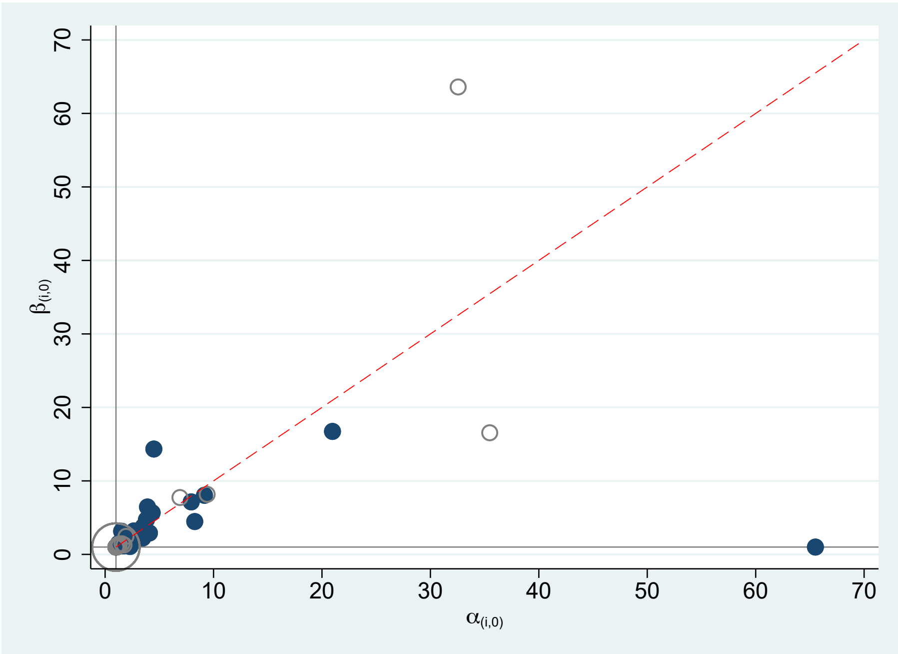

(b) $\hat{\alpha}_{i,0}, \hat{\beta}_{i,0} \leq 2$ (N=59)

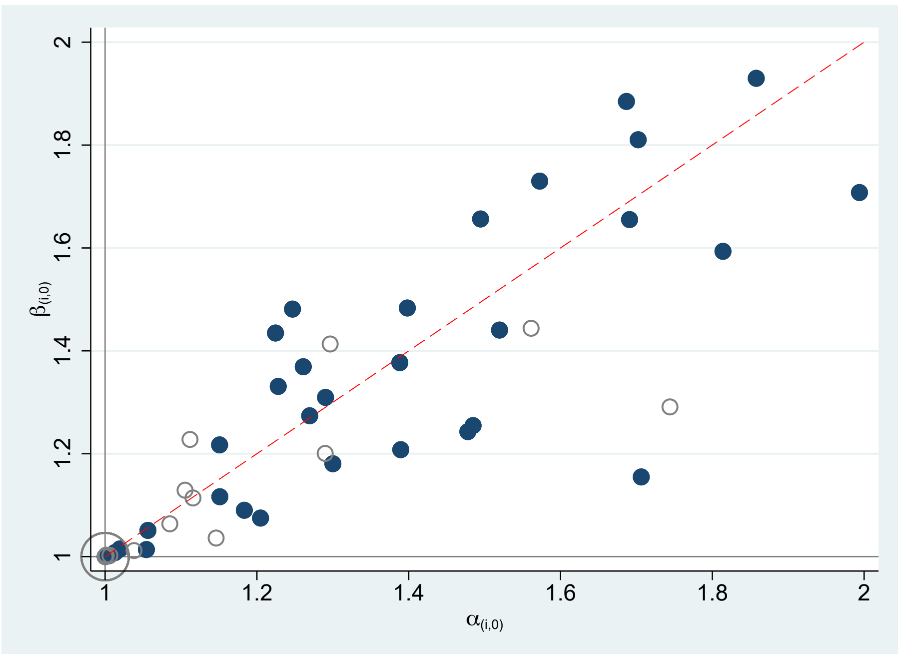

(c) Initial prior distributions: PMF representations (N=101)

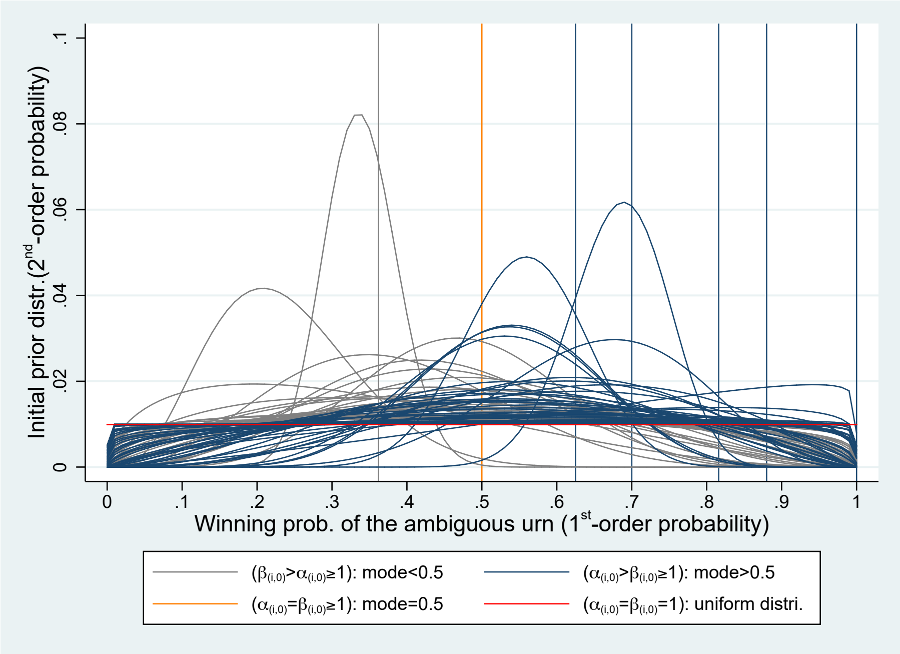

(d) Initial prior distributions: CDF representations (N=102)

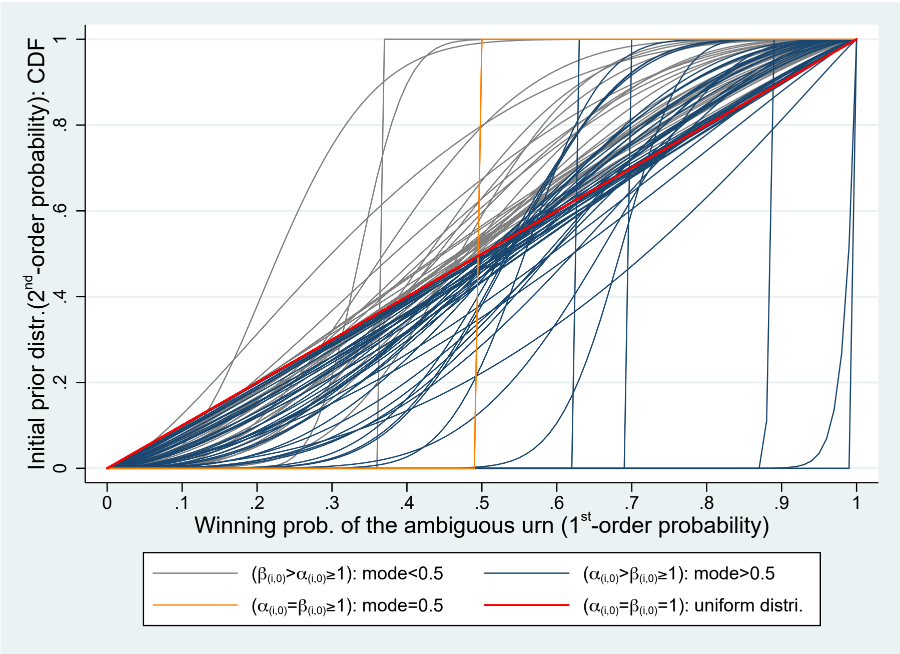

Figure 9: Estimated updating rule: $\hat{\gamma}_i^w$ and $\hat{\gamma}_i^b$

This diagram illustrates the estimation results of subjects' updating rule. The learning strategy assumes that a subject starts with a *beta*-distributed initial prior, and updates her beliefs by fully or partially reacting to what Bayes' rule implies. γ_i^w (γ_i^b) governs how subject i updates her belief parameters in response to a white (black) draw, in comparison with what Bayes' rule suggests. Parameters are estimated based on Equation (11). The diagrams plot $\hat{\gamma}_i^b$ against $\hat{\gamma}_i^w$ for each subject. For better visibility, Diagram (a) restricts to the subsamples with $\hat{\gamma}_i^w, \hat{\gamma}_i^b \leq 5$, and Diagram (b) restricts to the subsamples with $\hat{\gamma}_i^w, \hat{\gamma}_i^b \leq 1$. The bubble size represents the number of subjects (the largest bubble represents ten subjects; the smallest bubble represents one subject). Three reference lines are added: $\hat{\gamma}_i^w = 1$, $\hat{\gamma}_i^b = 1$, and a 45-degree line. $\hat{\gamma}_i^w = 1$ ($\hat{\gamma}_i^b = 1$) indicates that subject i perfectly follows Bayes' rule in belief updating when observing a white (black) draw. In both diagrams, the blue solid bubbles represent subjects for whom $H_0 : \gamma_i^w = 1$ & $\gamma_i^b = 1$ can be jointly rejected at 5% (significantly different from Bayesian updating). Otherwise, gray hollow bubbles.

(a) $\hat{\gamma}_i^w, \hat{\gamma}_i^b \leq 5$ (N=93)

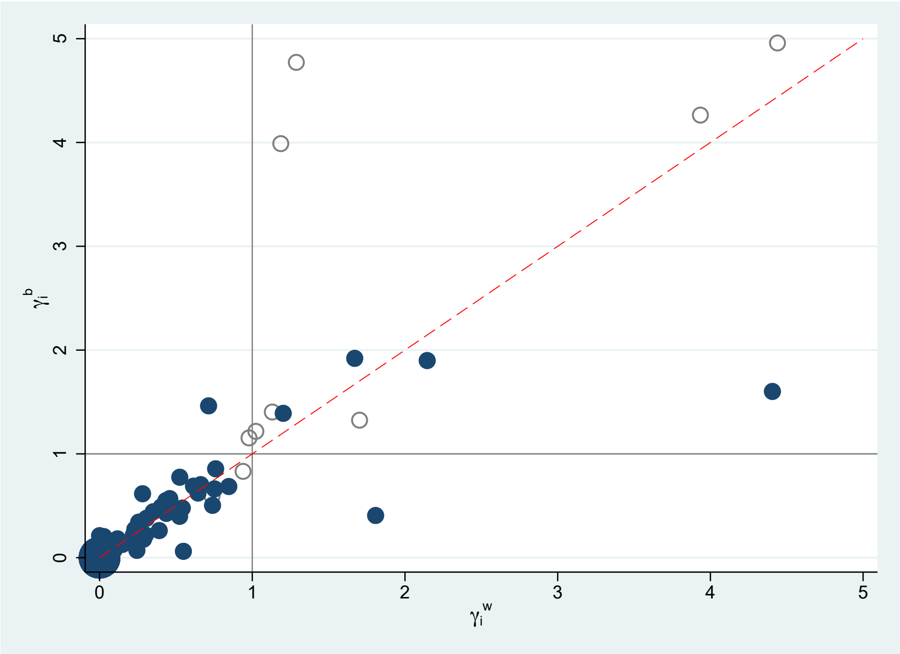

(b) $\hat{\gamma}_i^w, \hat{\gamma}_i^b \leq 1$ (N=79)

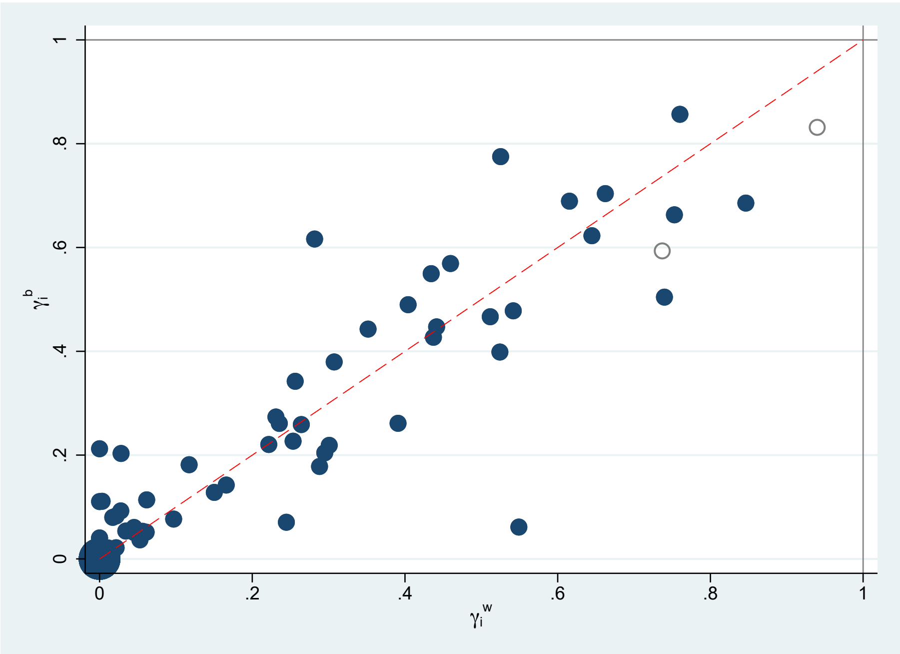

8 Appendix

Appendix A: Proposition 1

Proof. Plugging Equation (2)(4)(5) into (3) yields

$$\begin{aligned} Posterior(\theta|n,k;\alpha,\beta) &= \frac{\theta^{\alpha-1}(1-\theta)^{\beta-1}\theta^{k}(1-\theta)^{n-k}}{\int_{\theta'}\theta'^{\alpha-1}(1-\theta')^{\beta-1}\theta'^{k}(1-\theta')^{n-k}d\theta'} \\ &= \frac{\theta^{\alpha+k-1}(1-\theta)^{\beta+(n-k)-1}}{\int_{\theta'}\theta'^{\alpha+k-1}(1-\theta')^{\beta+(n-k)-1}d\theta'} \end{aligned} \tag{21}$$

Since

$$\int_{\theta} Prior(\theta|\alpha,\beta)d\theta = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)}\int_{\theta}\theta^{\alpha-1}(1-\theta)^{\beta-1}d\theta = 1 \tag{22}$$

hence

$$\int_{\theta}\theta^{\alpha-1}(1-\theta)^{\beta-1}d\theta = \frac{\Gamma(\alpha)\Gamma(\beta)}{\Gamma(\alpha+\beta)} \tag{23}$$

Replacing α with $\alpha+k$ and β with $\beta+(n-k)$ in equation (23) yields:

$$\int_{\theta}\theta^{\alpha+k-1}(1-\theta)^{\beta+(n-k)-1}d\theta = \frac{\Gamma(\alpha+k)\Gamma(\beta+n-k)}{\Gamma(\alpha+k+\beta+n-k)} \tag{24}$$

Plugging (24) into (21) yields:

$$Posterior(\theta|n,k;\alpha,\beta) = \frac{\Gamma(\alpha+k+\beta+n-k)}{\Gamma(\alpha+k)\Gamma(\beta+n-k)}\theta^{\alpha+k-1}(1-\theta)^{\beta+(n-k)-1} \tag{25}$$

Therefore, the posterior inherits the *beta*-distribution, characterized by the updated parameter bundle $(\alpha+k,\ \beta+n-k)$. □

Appendix B: Proposition 2

Proof. Using Equation (3) and $\alpha_{i,0}=\beta_{i,0}=1$ to rewrite the LHS of Equation (8) yields :

$$\underset{\theta}{\operatorname{argmax}}\, Posterior(\theta|n,k,\alpha,\beta) = \underset{\theta}{\operatorname{argmax}}[\theta^{k}(1-\theta)^{n-k}] \tag{26}$$

For readability, we suppress the subscripts of $\alpha_{i,0}$, $\beta_{i,0}$ and $k_{i,n}$. Take note that all other terms in the posterior expression can be taken out of the "argmax" operator, since they are independent of θ. First, consider an internal solution such that $\theta^*\neq 0$ and $\theta^*\neq 1$. The maximization problem can be rewritten as:

$$\max_{\theta\in(0,1)}[k\ln(\theta)+(n-k)\ln(1-\theta)] \tag{27}$$

FOC leads to:

$$\frac{k}{\theta^*} = \frac{(n-k)}{1-\theta^*} \tag{28}$$

SOC leads to:

$$-\frac{k}{\theta^{*2}} - \frac{(n-k)}{(1-\theta^*)^2} \leq 0 \tag{29}$$

Hence, the internal solution is:

$$\theta^* = \frac{k}{n} \tag{30}$$

where the RHS is the maximum likelihood update after observing k units of white draws out of n draws ($n \geq 1$).

Now consider if $\theta = 0$ is the solution to (26). If $k \neq 0$, the objective function is always positive with any $\theta \in (0,1)$. Therefore, $\theta = 0$ cannot be the solution. However if $k = 0$, the objective function reduces to $(1-\theta)^n$. Given the domains of the variables, $\theta^* = 0$ is the solution. In other words, only if $k = 0$, $\theta^* = 0$ is the solution to (26). Hence

$$\theta^* = 0 = \frac{k}{n}; \quad \text{if } k = 0. \tag{31}$$

Analogously, $\theta^* = 1$ is the solution to (26) only if $k = n$, therefore

$$\theta^* = 1 = \frac{k}{n}; \quad \text{if } k = n. \tag{32}$$

Combining Equation (30)-(32) yields Equation (8) □

Appendix C: Estimation equation of LS3

Proof. For a given subject after a given number of draws, her prior/posterior distribution recovered from LS3 should satisfy that the mode of this prior/posterior is equal to the elicited mode (the response in the guess game). Thus, the probability evaluated at the mode of this prior/posterior should be equal to the probability evaluated at the elicited mode:

$$\text{pdf}\left(\frac{\alpha_n - 1}{\alpha_n + \beta_n - 2}\middle| \alpha_n, \beta_n\right) = \text{pdf}\left(\frac{white_n}{100}\middle| \alpha_n, \beta_n\right) \tag{33}$$

The LHS denotes the second-order probability (in PDF) induced by the parameter bundle (α_n, β_n), evaluated at the true mode of the distribution. The RHS denotes the second-order probability (in PDF) induced by the parameter bundle (α_n, β_n), evaluated at the elicited mode. Using the PDF formula of *beta*-distribution (Equation 2) to rewrite Equation (33) yields

$$\left[\frac{\alpha_n - 1}{\alpha_n + \beta_n - 2}\right]^{\alpha_n - 1}\left[1 - \frac{\alpha_n - 1}{\alpha_n + \beta_n - 2}\right]^{\beta_n - 1} = \left[\frac{white_n}{100}\right]^{\alpha_n - 1}\left[1 - \frac{white_n}{100}\right]^{\beta_n - 1} \tag{34}$$

Plugging in Equation (9)(10) yields:

$$(M_n)^{\alpha_0 + \gamma^w k_n - 1}(1 - M_n)^{\beta_0 + \gamma^b (n - k_n) - 1} = \left[\frac{white_n}{100}\right]^{\alpha_0 + \gamma^w k_n - 1}\left[1 - \frac{white_n}{100}\right]^{\beta_0 + \gamma^b (n - k_n) - 1} \tag{35}$$

$$\text{where} \quad M_n \equiv \frac{\alpha_0 + \gamma^w k_n - 1}{\alpha_0 + \gamma^w k_n + \beta_0 + \gamma^b (n - k_n) - 2} \tag{36}$$

Re-arranging Equation (35) yields the regression equation:

$$\frac{white_n}{100} = M_n \left[\frac{1 - M_n}{1 - white_n/100}\right]^{\frac{\beta_0 + \gamma^b (n - k_n) - 1}{\alpha_0 + \gamma^w k_n - 1}} + \epsilon_n \tag{37}$$

where ϵ_n denotes the error term. □

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