

Li, Wenhui; Wilde, Christian

Working Paper

Belief formation and belief updating under ambiguity: Evidence from experiments

SAFE Working Paper, No. 251

Provided in Cooperation with:

Leibniz Institute for Financial Research SAFE

Suggested Citation: Li, Wenhui; Wilde, Christian (2020) : Belief formation and belief updating under ambiguity: Evidence from experiments, SAFE Working Paper, No. 251, Leibniz Institute for Financial Research SAFE, Frankfurt a. M.,
<https://doi.org/10.2139/ssrn.3399983>

This Version is available at:

<https://hdl.handle.net/10419/230682>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Wenhui Li – Christian Wilde

Belief Formation and Belief Updating under Ambiguity: Evidence from Experiments

SAFE Working Paper No. 251

SAFE | Sustainable Architecture for Finance in Europe

A cooperation of the Center for Financial Studies and Goethe University Frankfurt

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

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

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

Wenhui Li and Christian Wilde[†]

December 30, 2020

Abstract

This paper investigates beliefs in an ambiguous environment. In contrast to many previous studies, the beliefs regarding possible scenarios are measured independently from attitudes. We use laboratory experiments to estimate the entire distribution of subjective beliefs and examine how beliefs are updated, incorporating new information. We find that beliefs and updating rules are quite heterogeneous. For most subjects, we can reject the *objective equality hypothesis* that beliefs are uniformly distributed. The *unbiased belief hypothesis* cannot be rejected overall; Most subjects display no bias towards pessimism/optimism in beliefs. The *Bayesian updating hypothesis* can be rejected; Most subjects under-adjust beliefs in response to new information. Finally, we find that subjects adjust their beliefs symmetrically to good news and to bad news.

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 and seminar participants at the European Winter Meeting of the Econometric Society 2020 for helpful comments. We gratefully acknowledge research support from the 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 ambiguous environment formed and updated, and what is the attitude towards ambiguity. These two aspects, belief and attitude, work very differently in shaping final decisions. This indicates that beliefs should be studied independently from attitudes. However, a main stream of empirical literature does not separate beliefs from attitudes (Abdellaoui et al. 2011; Baillon et al. 2018a,b; Dimmock et al. 2016). Other literature which investigates beliefs independently from attitudes either ends up with some probably confounded measurements, or only provides a rough characterization of beliefs (eg. aggregate beliefs of all subjects participating in an experiment). This paper aims to fill these gaps: first, we empirically extract beliefs separated from attitudes. That is, we cleanly estimate subjects' beliefs without being confounded by attitudes. Second, this paper interprets a subject's belief as a distribution. This belief distribution captures how a subject assigns probability mass to each possible scenario of the ambiguous environment. Accordingly, we estimate the belief distributions of each subject. This provides a complete view of belief characterizations on an individual level, rather than some rough depictions of beliefs. Third, we investigate how each subject updates her belief distributions in response to new information. Since decisions are generally made on an individual level, a clean estimation of belief distributions of each subject after each update is fundamental to the study of decision making under ambiguity.

Ambiguity describes a situation in which some outcomes of an event occur with unknown probability (Becker and Brownson 1964; Ellsberg 1961; Epstein 1999; Knight 1921, to name a few). Among the ambiguity literature, the KMM model (Klibanoff et al. 2005) is a prominent theory which achieves the separation between beliefs and attitudes. It is the foundation for many other theoretical works which distinguish beliefs from attitudes (Gollier 2011; Izhakian 2020; Izhakian and Benninga 2011; Karni 2018; Maccheroni et al. 2013). The KMM model explicates an ambiguous environment with a two-order probability system, in which beliefs are characterized as distributions. Building on the KMM model, this paper manages to isolate beliefs from attitudes and characterizes beliefs as distributions in an empirical setting.

We use laboratory experiments to investigate subjects' beliefs. Ambiguity is operationalized by an urn. It is known to the subjects participating in the experiment that the urn contains in total 100 balls, and that each ball is either white or black. However, neither the number of white balls nor the number of black balls is known to any subject. In this way, a completely ambiguous environment is operationalized. Each possible composition of the urn, i.e. each possible combination of white balls and black balls, corresponds with a possible scenario. A possible scenario is equivalent to the first-order probability in the KMM model. In the experiments, we investigate the belief distribution of each subject. That is, we investigate how each subject assigns a probability to each possible scenario. This belief distribution is equivalent to the (conceived) second-order probability in the KMM model. In addition, we also investigate how each subject updates her belief distribution when new information about the urn is provided to her. With the recovery of each subject's belief distribution and its updating process,

we test three hypotheses: (1) *objective equality hypothesis*; (2) *unbiased belief hypothesis*; (3) *Bayesian updating hypothesis*. The hypothesis testing in turn helps us better understand how beliefs are formed and updated.

The first hypothesis, *objective equality hypothesis*, pertains to the shape of the initial prior of each subject. A subject’s initial prior is her belief distribution when no information regarding the composition of the urn is revealed. It is hypothesized that in an ambiguous environment, all possible, objectively-equally-likely scenarios are conceived equally likely by a subject. This hypothesis also implies that the initial prior of a subject is uniformly distributed under complete ambiguity if all scenarios are objectively equally likely. This is exactly the case in the experimental setting of this paper: before any draw has occurred and without any additional information about the composition of the urn is revealed, there is no reason to believe that one scenario is more likely than another scenario. This paper recovers the initial prior of each subject, and studies the shapes of the initial priors thoroughly.

Most previous literature which investigates beliefs does not elicit the entire prior distribution of each subject. Some literature only recovers the aggregate belief distribution across subjects (Buser et al. 2018; Cubitt et al. 2018; Mobius et al. 2014), rather than the belief distribution of each individual subject as in our paper. Some literature does elicit the initial prior distribution of each subject. One example is Ertac (2011). The author, however, only investigates the probability mass assigned to some subsets of the scenarios, not each individual scenario. This only gives a rough picture of the initial prior. We recover the probability mass assigned to each individual scenario, providing a full characterization of each initial prior. Other examples which elicit the initial prior distribution of each subject include Coffman et al. (2019), Eil and Rao (2011). From a theoretical perspective, the elicitation methods in these two papers are not incentive compatible. In comparison, our experiments have sound incentive compatibility. More importantly, in any previous literature mentioned in this paragraph, the *objective equality hypothesis* cannot be tested. The reason is that these papers investigate beliefs regarding subjects’ performance in some self-related tasks. A subject naturally does not treat all scenarios equally likely after the task, since she already has a picture about her own performance in the task. In other words, scenarios are anyway not equally likely in the first place. Therefore, the *objective equality hypothesis* cannot be tested in any of these papers. Unlike them, our paper sets up an ambiguous environment in which all scenarios are equally likely from an objective point of view when eliciting initial priors. This enables the testing of the first hypothesis. In addition, since all subjects face the identical initial environment, the estimated initial priors and the test results of the first hypothesis are comparable across subjects.

In the course of the experiment, the urn is used to determine the payoff of some lottery: a ball is randomly drawn out from the urn; in case that a white ball is drawn out, the lottery pays out a positive financial reward; in case that a black ball is drawn out, the lottery pays out zero. Therefore, a scenario (regarding the composition of the urn) with a large number of white balls is interpreted as a *good scenario*, and a white draw is regarded as *good news*. On the contrary, a scenario with a large number of black balls is interpreted as a *bad scenario*, and a black draw is regarded as *bad news*. This is relevant to the second hypothesis.

The second hypothesis, *unbiased belief hypothesis*, tests whether there exists pessimism (i.e. bias towards bad scenarios) or optimism (i.e. bias towards good scenarios) in beliefs. It is hypothesized that a subject has an unbiased initial prior. That is, there exists no pessimism or optimism in a subject’s initial prior. It seems that pessimism/optimism has no unified definition in literature. Giraud and Thomas (2017) as well as Binmore et al. (2012) treat pessimism as a synonym of ambiguity aversion, and optimism as a synonym of ambiguity loving. On the other hand, Ahn et al. (2014) define pessimism as overweighting the probabilities of low payoffs and underweighting the probabilities of high payoffs, a belief characteristic. Our paper is in line with the latter way of defining pessimism/optimism. In our paper, pessimism is defined as a bias in belief such that more probability mass is assigned to bad scenarios (i.e. a large number of black balls in the urn) than a neutral benchmark belief. Analogously, optimism is defined as a bias in belief such that more probability mass is assigned to good scenarios (i.e. a large number of white balls in the urn) than a neutral benchmark belief. An unbiased belief is a belief displaying no bias in either way. These definitions also indicate that this paper only investigates pessimism/optimism in beliefs, not in attitudes. The clean estimation of beliefs in this paper makes the investigation possible.

To the best of our knowledge, pessimism/optimism strictly defined as a bias in beliefs in comparison with an objective benchmark is not yet researched. More often studied is overconfidence in beliefs, which is documented in a stream of current literature (Brenner et al. 2015; Grossman and Owens 2012). Although similar to optimism, overconfidence in fact is a different concept. Overconfidence is defined as a subject’s overestimation of her performance/ability. That is, overconfidence is a bias in beliefs using a subject’s actual performance/ability as the benchmark, i.e. a subjective benchmark. In comparison, an objective benchmark is used to define optimism. Some literature uses the terms “optimistic” and “pessimistic” to describe biased beliefs (Coffman et al. 2019; Ertac 2011). But in fact the authors refer to overconfidence and its antonym, since the benchmarks are still subjects’ performance/ability in this literature. Our paper investigates biases in beliefs in comparison with an objective benchmark. The investigation of pessimism/optimism in beliefs is original in the studies of beliefs under ambiguity.

The third hypothesis, *Bayesian updating hypothesis*, pertains to subjects’ belief updating (also called *learning*). It is hypothesized that a subject perfectly employs Bayes’ rule when updating beliefs. Our experiments permit learning in the way that new information about the composition of the urn, generated by random draws from the urn, is provided to subjects (for the detailed design, see Chapter 3.1.). Previous literature modeling learning strategies can be roughly categorized into two streams: literature assuming Bayesian updates (Branger et al. 2013; Gilboa and Schmeidler 1993; Hanany and Klibanoff 2007; Peijnenburg 2018; Pires 2002), where subjects update their beliefs employing Bayes’ rule, and literature assuming non-Bayesian updates (Epstein 2006; Epstein et al. 2008; Marinacci 2002), where subjects deviate from Bayes’ rule or completely abandon Bayes’ rule. To better understand the latent updating mechanism of each subject, this paper introduces a learning strategy which incorporates both, obedience to Bayes’ rule and departure from Bayes’ rule. For each subject, we estimate her belief updating rule (i.e. learning strategy). This allows us to test the third hypothesis. Moreover,

we study the updating rule of each subject in a more detailed manner. In particular, we focus on two characteristics of a belief updating rule: whether a subject under-/over-adjusts beliefs compared with what Bayes' rule would suggest, and whether a subject responds to good news (i.e. white draws) and bad news (i.e. black draws) symmetrically.

A sizable number of empirical literature studies whether subjects employ Bayes' rule when updating beliefs. Departure from Bayes' rule is often documented in this literature (Charness and Levin 2005; Filippis et al. 2017; Ouwersloot et al. 1998; Zizzo et al. 2000, to name a few). Among the departures from Bayes' rule, under-adjustment of beliefs is observed in Buser et al. (2018) and Mobius et al. (2014). As for symmetry of updating, Buser et al. (2018) find that subjects respond to good news and bad news symmetrically, whereas Eil and Rao (2011) as well as Mobius et al. (2014) discover that subjects are more responsive to good news than to bad news. These results indicate that how a subject deviates from Bayes' rule specifically is still inconclusive. Our paper recovers the updating rule of each individual subject and provides subject-level evidence regarding this question.

This paper reaches the following findings: (1) The initial priors are mostly characterized by bell-shaped distributions, rather than a uniform distribution. This result speaks against the *objective equality hypothesis*. That is, most subjects do not conceive all scenarios under complete ambiguity equally likely. (2) The investigation of pessimism/optimism in initial beliefs supports the *unbiased belief hypothesis*: most subjects demonstrate no bias towards bad or good scenarios. Among the biased initial beliefs, optimism is more prevalent than pessimism. (3) The recovered belief updating rules show that the *Bayesian updating hypothesis* can be rejected for most subjects; Most subjects significantly deviate from Bayes' rule when updating beliefs. (4) Among the subjects who significantly deviate from Bayes' rule, most of them under-adjust beliefs in comparison with what Bayes' rule implies. (5) Most subjects respond to good news and bad news symmetrically.

The rest of the paper is organized as follows: Section 2 presents the background and underlying assumptions related to beliefs under ambiguity. Section 3 introduces the experiment design. Section 4 presents the descriptive analysis of the belief 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 in this paper.

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: $\pi : \sigma(\Omega) \rightarrow [0, 1]$. 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 the assumption 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: $\mu : \sigma(\Delta) \rightarrow [0, 1]$. The second source of uncertainty arises from the assumption that at least one likelihood measurement in μ is different from one and different from zero. 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 urn operationalizing ambiguity in this paper. There are only two possible states, white or black. This can be written as $\Omega = (w, b)$, where w denotes white and b denotes black. The proportion of white balls in the urn (in percent) is denoted as π_w . Likewise, we can define $\pi_b \equiv 1 - \pi_w$ as the proportion of black balls. $\pi = (\pi_w, \pi_b)$ constructs the first-order probability. There are multiple possible values for π . Given the design of the urn, π can be $(0, 1), (0.01, 0.99), \dots, (1, 0)$. Each possible value of π corresponds with a possible scenario regarding the composition of the urn. We formally define all possible values of π as $\Delta = \{\pi^j = (j/100, 1 - j/100) : j = 0, 1, \dots, 100\}$. The second-order probability is denoted as $\mu : \sigma(\Delta) \rightarrow [0, 1]$. In the context of the urn, μ is the probability mass assigned to each possible composition of the urn (i.e. the belief).

In our paper, μ is assumed to be subjective and non-degenerate for most of the subjects. Each individual can formulate a belief, μ , 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 μ evaluated at any single π^j is different from one, at least for most subjects. In other words, the belief, μ , is non-trivial.

Belief updating. The belief, μ , is not static over time. Subjects can update μ using the available relevant information. This process of updating μ is also called *learning*. A subject's belief, μ , can be updated every time a random draw is implemented and the realized state (white or black) is observed by the subject. The initial μ , i.e. the belief before any draw has occurred, is called the initial prior. An updated μ , i.e. a belief after some random draw(s) is/are observed, is called posterior. Each random draw generates one piece of new information, and thus corresponds with one specific posterior.

Pessimism, optimism, and unbiased beliefs. As explained above, the urn is also used to determine the payoffs of several lotteries: a white draw always results in a positive payoff, while a black draw always results in zero payoff. We define that good scenarios (regarding the composition of the urn) are those $\pi^j \in \Delta$ with a large proportion of white balls (i.e. large π_w), and bad scenarios are those $\pi^j \in \Delta$ with a large proportion of black balls (i.e. large π_b). Therefore, the best scenario is $\pi^{100} = (1, 0)$, i.e. $\pi_w = 1, \pi_b = 0$, while the worst scenario is $\pi^0 = (0, 1)$, i.e. $\pi_w = 0, \pi_b = 1$. The midpoint between the best and the worst scenario, $\pi^{50} = (0.5, 0.5)$, is a natural threshold discriminating between good and bad scenarios.

As a characteristic of belief, pessimism is defined as a bias in belief towards bad scenarios. It is characterized by a second-order probability μ which associates more probability mass to bad scenarios than a neutral benchmark belief. In contrast, optimism is defined as a bias in belief towards good scenarios. It is characterized by a second-order probability μ which associates more probability mass to good scenarios than a neutral benchmark belief. Unbiased

beliefs are beliefs which display no biases in either way.

3 Experiment design

3.1 Guess game

As explained above, we operationalize ambiguity with an urn. Subjects are told that there are in total 100 balls in the urn, and that each ball is either a white ball or a black ball. However, neither the number of white balls nor the number of black balls is known to any subject. No information is conveyed regarding how the urn is assembled. A physical urn with such ambiguous feature is simply displayed in front of the subjects as a final product. In this way, the urn operationalizes a completely ambiguous environment. This also implies that at the very beginning of the experiment, from an objective point of view, each possible urn composition is equally likely.

A so-called *guess game* is designed to track down subjects' beliefs regarding the composition of the urn, or simply the conceived number of white balls in the urn. In a guess game, a subject needs to answer the following question: standing at this point, how many white balls do you think are there in the urn? A subject enters an integer between zero and 100 (both ends are inclusive). 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 random draws from the 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, repeats 15 times. In addition, in all Sessions except Session I, an additional guess game, G_{15} , is played after the 15th draw¹. 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 ($n = 0, 1, \dots, 15$). 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 the 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.

Subjects are incentivized in the way that every time a subject enters the correct number of white balls 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 belief distribution in each

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

guess game. The earning a subject obtained from the guess games is announced to her only at the very end of the experiment, hence the ambiguous feature of the urn sustains throughout the entire experiment.

The guess games elicit the mode values of subjects' belief distributions, rather than each entire belief distribution. This design has several advantages compared with other alternatives. First, elicitation of mode values has sound incentive compatibility (more explanations can be found in Chapter 3.2). One alternative could be to elicit each entire belief distribution directly, with incentivization following the quadratic scoring rule (QSR) (Eil and Rao 2011; Ertac 2011). Yet, the incentive compatibility of QSR relies on the assumption of risk neutrality. This assumption is too strong since previous evidence shows that subjects are mostly risk averse in laboratory experiments (Holt and Laury 2002). Second, eliciting a singleton mode provides more precise information than eliciting, for instance, which subset of the possible scenarios contains the true scenario. One of the experiments in Coffman et al. (2019) elicits such a subset. This only identifies a (rough) range in which the true value is conceived to be located. Third, it is cognitively easier for a subject to guess the number of white balls than to report the entire distribution of her belief. Hence, by eliciting only mode values, we reduce the chance of noise in the data. Considering all these factors, we finally choose to elicit the mode of a subject's belief distribution, rather than other alternatives.

The true proportion of white balls in the ambiguous urn is fixed at 40 for all subjects. Basically, it is beneficial to set the true value close to 50, since it generates more balanced proportions of white draws and black draws along a draw history. In contrast, setting the 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, setting the value close to 50 also increases the variation of draw history, conducive to the belief estimation. However, the value 50 itself is a very prominent number. 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 set 40 to be 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.

It is worth mentioning that the guess game design guarantees that the data obtained from the games are purely related to the beliefs regarding the ambiguous environment, independent from the attitude towards ambiguity (i.e. preference). Attitude plays a role in decision and becomes observable only when subjects compare different alternatives. Since the guess game does not imply any preference-related decisions between two alternatives, and it merely requests that subjects report their beliefs regarding 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 μ_j is the conceived probability mass assigned to the scenario in which the urn contains j units of white balls and $100-j$ units of black balls. Under the incentivization method, the expected earning of a subject who enters j' in a guess game is $2\mu_{j'}$ Euro. To maximize the expected earning, a subject 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 $G_n : n = 0, 1, \dots, 15$. Therefore, our experiment induces each subject 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 a subject'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 a subject 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, a subject's elicited beliefs may be different from her 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, subjects are asked to price several lotteries whose payoffs are determined by random draws from the urn used in the guess games (i.e. subjects need to report their conceived reservation values of these lotteries). 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.

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.

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 does not have 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), some quizzes as well as 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. A demonstration of a physical ambiguous 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 set-up and the tasks. 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 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 in Session I-VI 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:

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

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

where N_m denotes the number of subjects in market m , and T denotes the highest number of draws². For the 19 subjects in Session VII, each subject observes her own draw history, thus generating 19 independent paths of draw history. We can directly report $white_{i,n}$ for each of these 19 subjects. 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, corresponding with the 19 subjects in Session VII. A white bar represents a white draw, while a gray bar represents 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 Session VII, 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

²In Market 1, 14 draws (with replacement) are implemented, and accordingly 15 guess games (G_0, G_1, \dots, 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 ($G_0, G_1, \dots, G_{14}, G_{15}$) 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 $\theta = 0, 0.01, 0.02, \dots, 1$. 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 shape³. 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:

$$\begin{aligned} \text{Prior}(\theta|\alpha, \beta) &= \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} \theta^{\alpha-1} (1 - \theta)^{\beta-1} \\ \theta &\in [0, 1]; \quad i \in \{1, 2, \dots, N\}; \quad \alpha, \beta \geq 1 \end{aligned} \quad (2)$$

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:

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

where

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

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

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 $(\alpha + k, \beta + n - k)$.*

³For 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:

$$\alpha_{i,n} = \alpha_{i,0} + k_{i,n} \quad (6)$$

$$\beta_{i,n} = \beta_{i,0} + n - k_{i,n} \quad (7)$$

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

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 ($n = 1, 2, \dots$) of the same draw history. The proposition below formally presents this relation:

Proposition 2.

$$\begin{aligned} \operatorname{argmax}_{\theta} \text{Posterior}(\theta|n, k_{i,n}, \alpha_{i,0}, \beta_{i,0}) &= \frac{k_{i,n}}{n} \\ \alpha_{i,0} = \beta_{i,0} &= 1; \quad n = 1, 2, \dots, T; \quad i \in \{1, 2, \dots, N\} \end{aligned} \quad (8)$$

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-adjust or over-adjust beliefs in response 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:

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

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

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

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-adjusts beliefs in comparison with 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 over-adjustment of

beliefs in comparison with 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-adjustment of beliefs compared to Bayes' rule, and $\gamma_i^b > 1$ implies over-adjustment of beliefs 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 $\{(\alpha_{i,n}, \beta_{i,n}) : n = 0, 1, \dots, T\}$ 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:

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

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

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, α_0 , β_0 , γ^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:

$$\alpha_0 = 1 + \exp(a) \quad (13)$$

$$\beta_0 = 1 + \exp(b) \quad (14)$$

$$\gamma^w = \exp(r^w) \quad (15)$$

$$\gamma^b = \exp(r^b) \quad (16)$$

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. 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, $\hat{\alpha}_{i,0} = \hat{\beta}_{i,0} \rightarrow +\infty$ and $\hat{\gamma}_i^w = \hat{\gamma}_i^b = 0$. For the subject constantly responding 70, $\hat{\alpha}_{i,0}, \hat{\beta}_{i,0} \rightarrow +\infty$ 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 her 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 $\{\theta | \theta = 0, 0.01, 0.02, \dots, 1\}$ and compute the probability mass (PMF) evaluated at each element of the set:

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

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 ($n = 0, 1, \dots, T$).

5.4 Estimation results

In this chapter, we discuss the estimation results of the parameters as computed based on Equation (17)-(20). 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 6 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 7-8. Table 3 Panel A summarizes the results of $(\hat{\alpha}_{i,0}, \hat{\beta}_{i,0})$, i.e. subjects' initial priors. Figure 7a and 7b 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 7a represents the sub-samples with $\hat{\alpha}_{i,0}, \hat{\beta}_{i,0} \leq 70$ (N=94). For better visibility, Figure 7b 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 7c (in form of PMF) and in Figure 7d (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 8a and 8b 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 8a and 8b 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, not surprisingly, no subject perfectly follows Bayes' rule in belief updating. A large proportion of subjects display under-adjustment of beliefs in comparison with what Bayes' rule implies. In

fact, 75% of the subjects have both $\hat{\gamma}_i^w < 1$ and $\hat{\gamma}_i^b < 1$. In addition, subjects tend to respond to white draws and black draws symmetrically. This can be seen from both diagrams, in which most bubbles are located close to the 45-degree line.

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-adjust beliefs in response 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.

In general, the nonlinear regressions fit the observed data very well: most regressions achieve a high R^2 (mean=0.979; s.d.=0.060; max=0.999; min=0.683; N=102). The high R^2 does not seem to be driven by overfitting. First, from a theoretical perspective, overfitting is more likely to be a serious concern if the estimation model lacks theory support (Claeskens and Hjort 2008). This is not the case in this paper, since our belief estimation model has a sound theoretical structure. It is an extension of the conjugate prior theory, which is empirically tested in a sizable number of previous studies. Second, to rule out from an empirical perspective that overfitting is a serious problem in our regressions, we re-run the regressions using subsets of the data: for each subject, we randomly delete one data point (i.e. one guess game response); With this smaller data set, we re-run the belief estimation for each subject and test to which extent the estimations provide different results. Stable estimation results speak against overfitting. We find that the re-run estimation leads to very similar results to those reported in Table 3. For example, in Panel A, the re-run estimation skipping one data per subject fills Row (1)-(6) with 16, 1, 1, 3, 44, 37, respectively, similar to the results based on the complete data (Row 1-6: 10, 0, 2, 2, 50, 38, respectively). Thus, the re-run estimation generates stable results regarding initial priors as in the original estimation: most subjects still have bell-shaped initial prior distributions. As for $\hat{\gamma}_i^w$, the re-run estimation fills Row (7)-(10) of Panel B with 16, 65, 0, 21, respectively (the results based on the complete data are 12, 69, 0, 21, respectively); For $\hat{\gamma}_i^b$, the re-run estimation fills Row (7)-(10) with 19, 64, 0, 19, respectively (the results based on the complete data are 13, 68, 0, 21, respectively). The re-run estimation generates stable results regarding updating rules as in the original estimation: most subjects still under-adjust beliefs compared to Bayes' rule. The results from the re-run regressions indicate that overfitting does not seem to be a big problem in our estimation.

Using the estimated parameters, we recover the belief distribution dynamics of each subject. This allows us to formally test three hypotheses related to belief formation and belief updating.

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 an ambiguous environment, all possible, objectively equally likely scenarios are conceived equally likely;
- (b) **Unbiased belief hypothesis:** in a completely ambiguous environment, a subject’s belief demonstrates no bias towards bad or good scenarios; 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, objectively equal scenarios are conceived 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 7a and 7b: 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 possible, objectively equal 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 conceived 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 a subject'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 a subject's initial belief is to compare her 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 a subject'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 subject i equally weights bad scenarios and good scenarios. 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 subject i assigns higher likelihood to bad scenarios than to good scenarios. This is categorized as pessimism. In case of $\hat{\alpha}_{i,0} > \hat{\beta}_{i,0}$, the distribution is left-skewed, meaning that subject i assigns higher likelihood to good scenarios than to bad scenarios. 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 $H_0 : \alpha_{i,0} = \beta_{i,0}$ 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 $\hat{\alpha}_{i,0} < \hat{\beta}_{i,0}$, whereas 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 bad scenarios: the proportion of white balls is less than 0.5) with the average probability mass assigned to $\theta > 0.5$ (all good scenarios: the proportion of white balls is greater than 0.5) 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 differ from the current literature as follows. First, we empirically identify pessimism/optimism without any confounding effects from attitudes. This is achieved (in the second and the third diagnosis) by the estimation of the belief parameters in a regression absent from attitude parameters. In comparison, Ahn et al. (2014) jointly estimate parameters governing pessimism/optimism and parameters governing ambiguity attitude. This may lead to a biased estimation of pessimism/optimism. Second, a stream of literature studies a subject’s beliefs regarding her performance in some cognitive or

ego-related tasks. The subject’s true performance (e.g. her score in a test, or her ranking in a group) is set as the benchmark for the diagnosis of pessimism/optimism. Such benchmark is obviously subjective. Unlike this literature, we define and measure pessimism/optimism based on an objective neutral benchmark. This is original in the studies of beliefs under ambiguity.

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-adjust beliefs 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 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 8a and 8b: the blue solid bubbles, which represent the imperfect Bayesian updaters, are the majorities in both diagrams.

Furthermore, we scrutinize the belief updating of the 83 subjects who significantly deviate from Bayes’ rule. In particular, we focus on two belief updating features: under-/over-adjustment of belief, and symmetric/asymmetric updating. Among the 83 imperfect Bayesian updaters, most subjects under-adjust beliefs in response to new information compared to what Bayes’ rule implies. This can be directly seen from Figure 8a: most blue solid bubbles (i.e. 74 out of 83 subjects) lie in the area where $\hat{\gamma}_i^w < 1$ and $\hat{\gamma}_i^b < 1$. This finding confirms the results from several papers: Mobius et al. (2014) and Buser et al. (2018) also find that subjects display conservatism in belief updating (i.e. under-adjustment). Coffman et al. (2019) find that both men and women hold relatively lower beliefs than what Bayes’ rule would predict in gender incongruent tasks.

As for whether beliefs are updated symmetrically, most of the 83 imperfect Bayesian updaters respond to good news (white draws) and bad news (black draws) symmetrically. This is supported by the fact that for 50 out of the 83 subjects (i.e. 60%), the null hypothesis $H_0 : \gamma_i^w = \gamma_i^b$ cannot be rejected at 5%. This confirms the finding in Buser et al. (2018) that

asymmetry is not significant. But it contrasts the finding in Mobius et al. (2014) that asymmetry is significant, and the finding in Coffman et al. (2019) that women are more responsive to bad news than to good news. For the remaining 33 subjects who respond asymmetrically, 20 of them are more responsive to bad news than to good news (i.e. $\hat{\gamma}_i^w < \hat{\gamma}_i^b$). This result to some extent contrasts the findings in Mobius et al. (2014) as well as in Eil and Rao (2011). These authors conclude that subjects are more responsive to good news than to bad news. In general, the dominating symmetry in belief updating and higher responsiveness to bad news than to good news among asymmetric updating differentiate our paper from other works. These differences in results may be due to the fact that beliefs in our paper are beliefs regarding some objective, ego-unrelated event (guessing the number of white balls), while beliefs in all of the literature mentioned in this paragraph are beliefs regarding some ego-related events (eg. IQ test score/ranking, personal attractiveness). This implies that, as argued by Ertac (2011), belief updating differs across ego-related and ego-unrelated contexts. In addition, our paper also differs from other literature in estimation method: other papers mainly generate results on a sample-population level. We are able to recover belief distributions of each subject and conclude on an individual level, which generates a more precise zoom-in picture.

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-adjusting beliefs in response 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 a subject's initial belief follows a uniform distribution, i.e. that a subject conceives each objectively equal 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. These subjects who significantly deviate from Bayes' rule mostly under-adjust their beliefs compared to what Bayes' rule suggests, and they respond to good news and bad news symmetrically.

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 subjects' 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.
- Baillon, A., Bleichrodt, H., Keskin, U., l'Haridon, O., and Li, C. (2018a). The Effect of Learning on Ambiguity Attitudes. *Management Science*, 64(5):2181–2198.
- Baillon, A., Huang, Z., Selim, A., and P. Wakker, P. (2018b). Measuring ambiguity attitudes for all (natural) events. *Econometrica*, 86:1839–1858.
- Becker, S. W. and Brownson, F. O. (1964). What price ambiguity? or the role of ambiguity in decision-making. *Journal of Political Economy*, 72(1):62–73.

- 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.
- Brenner, M., Izhakian, Y., and Sade, O. (2015). Ambiguity and Overconfidence. Working Papers 11-06, New York University, Leonard N. Stern School of Business, Department of Economics.
- 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.
- Charness, G. and Levin, D. (2005). When optimal choices feel wrong: A laboratory study of bayesian updating, complexity, and affect. *American Economic Review*, 95(4):1300–1309.
- Claeskens, G. and Hjort, N. L. (2008). *Model selection and model averaging*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge Univ. Press, Leiden.
- 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. *Theory and Decision*, 83(3):275–302.
- Diaconis, P. and Ylvisaker, D. (1979). Conjugate priors for exponential families. *The Annals of Statistics*, 7(2):269–281.
- Dimmock, S., Kouwenberg, R., and Wakker, P. (2016). Ambiguity attitudes in a large representative sample. *Management Science*, 62(5):1363–1380.
- 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.
- Ellsberg, D. (1961). Risk, ambiguity, and the savage axioms. *The Quarterly Journal of Economics*, 75(4):643–669.
- 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.
- 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.
- Gelman, A., Carlin, J. B., Stern, H. S., and Rubin, D. B. (2004). *Bayesian Data Analysis*. Chapman and Hall/CRC, 2nd ed. edition.
- 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.
- Grossman, Z. and Owens, D. (2012). An unlucky feeling: Overconfidence and noisy feedback. *Journal of Economic Behavior & Organization*, 84(2):510–524.
- 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.
- Izhakian, Y. (2020). A theoretical foundation of ambiguity measurement. *Journal of Economic Theory*, 187:105001.
- Izhakian, Y. and Benninga, S. (2011). The uncertainty premium in an ambiguous economy. *Quarterly Journal of Finance (QJF)*, 01(02):323–354.
- 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.
- Maccheroni, F., Marinacci, M., and Ruffino, D. (2013). Alpha as ambiguity: Robust mean variance portfolio analysis. *Econometrica*, 81(3):1075–1113.
- 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. Working paper.
- Ouwersloot, H., Nijkamp, P., and Rietveld, P. (1998). Errors in probability updating behaviour: Measurement and impact analysis. *Journal of Economic Psychology*, 19(5):535 – 563.
- Peijnenburg, K. (2018). Life-cycle asset allocation with ambiguity aversion and learning. *Journal of Financial and Quantitative Analysis*, 53(5):1963–1994.
- Pires, C. (2002). A rule for updating ambiguous beliefs. *Theory and Decision*, 53(2):137–152.
- Schlaifer, R. and Raiffa, H. (1961). Applied statistical decision theory. *Journal of the American Statistical Association*, 57.
- Zizzo, D. J., Stolarz-Fantino, S., Wen, J., and Fantino, E. (2000). A violation of the monotonicity axiom: experimental evidence on the conjunction fallacy. *Journal of Economic Behavior & Organization*, 41(3):263 – 276.

Table 1: Experiment procedure

This table reports the main procedure of a complete experiment session. Some games may show up only in some sessions, instead of all sessions, but the time-line applies to 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. 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 observes its own draw. The ambiguous urn design is identical universally in all sessions.

n=0 →	n=1 →	...	n=14 →	n=15
Choice games				
↓				
G_0	G_1	...	G_{14}	G_{15}
↓	↓		↓	↓
Asset Trading	Asset Trading	...	Asset Trading	Choice games
↓	↓		↓	↓
1 st draw	2 nd draw	...	15 th draw	Earning announced; Questionnaire; Final payment.

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, \dots, G_{14}	Path 1	choice games, asset trading
II	2	14-20	7	G_0, G_1, \dots, G_{15}	Path 2	choice games, asset trading
	3	21-27	7		Path 3	
III	4	28-34	7	G_0, G_1, \dots, G_{15}	Path 4	choice games, asset trading
	5	35-41	7		Path 5	
IV	6	42-48	7	G_0, G_1, \dots, G_{15}	Path 6	choice games, asset trading
	7	49-55	7		Path 7	
V	8	56-62	7	G_0, G_1, \dots, G_{15}	Path 8	choice games, asset trading
	9	63-69	7		Path 9	
VI	10	70-76	7	G_0, G_1, \dots, G_{15}	Path 10	choice games, asset trading
	11	77-83	7		Path 11	
VII	-	84-102	19	G_0, G_1, \dots, G_{15}	Path 12-30	choice games
<i>Total=102</i>					<i>Total=30</i>	

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 perfectly/imperfectly following Bayes' rule. 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 $\hat{\alpha}_{i,0} = \hat{\beta}_{i,0} \rightarrow +\infty$ (entering Row 4). The subject who sticks to 70 in all guess games is estimated by $\hat{\alpha}_{i,0}, \hat{\beta}_{i,0} \rightarrow +\infty$ 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-adjustment compared to Bayes' rule	69	67.65	68	66.67
(9) $\hat{\gamma} = 1$: employment of Bayes' rule				
(10) $\hat{\gamma} > 1$: over-adjustment compared 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 is grouped into the same category in the last two columns as in the “without significance test” columns.

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. In the last two columns, each of them is grouped into the same category as they are in the “without significance test” 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. In the last two columns, each of them is grouped into the same category as they are in the “without significance test” 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. In the last two columns, each of them is grouped into the same category as they are in the "without significance test" 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, \dots , G_{14} , the 15th draw. In all other sessions: G_0 , the first draw, G_1 , the second draw, \dots , G_{14} , the 15th 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

Round 1 out of 15	Remaining Time[sec]: 75								
"Guess Game": guess the number of white balls Standing at this point, how many white balls do you think are in the urn? Please insert your own guess: a number between 0 and 100 (inclusive).									
<table border="1" style="margin: auto; border-collapse: collapse;"> <thead> <tr> <th style="padding: 5px;">Color of the ball</th> <th style="padding: 5px;">No. of balls</th> </tr> </thead> <tbody> <tr> <td style="padding: 5px; text-align: center;">White</td> <td style="padding: 5px;"><input style="width: 100%;" type="text" value="1"/></td> </tr> <tr> <td style="padding: 5px; text-align: center;">Black</td> <td style="padding: 5px;"><input style="width: 100%;" type="text" value=""/></td> </tr> <tr> <td style="padding: 5px; text-align: center;">Total</td> <td style="padding: 5px; text-align: center;">100</td> </tr> </tbody> </table>		Color of the ball	No. of balls	White	<input style="width: 100%;" type="text" value="1"/>	Black	<input style="width: 100%;" type="text" value=""/>	Total	100
Color of the ball	No. of balls								
White	<input style="width: 100%;" type="text" value="1"/>								
Black	<input style="width: 100%;" type="text" value=""/>								
Total	100								
<input style="background-color: red; color: white; padding: 5px 15px; border: none;" type="button" value="Submit"/>									

(b) Guess game after five draws: G_5

Round 6 out of 15	Remaining Time[sec]: 73																																																																											
"Guess Game": guess the number of white balls Standing at this point, how many white balls do you think are in the urn? Please insert your own guess: a number between 0 and 100 (inclusive). For your reference, the draw results in the past period are displayed below.																																																																												
<table border="1" style="margin: auto; border-collapse: collapse;"> <thead> <tr> <th style="padding: 5px;">Color of the ball</th> <th style="padding: 5px;">No. of balls</th> </tr> </thead> <tbody> <tr> <td style="padding: 5px; text-align: center;">White</td> <td style="padding: 5px;"><input style="width: 100%;" type="text" value="1"/></td> </tr> <tr> <td style="padding: 5px; text-align: center;">Black</td> <td style="padding: 5px;"><input style="width: 100%;" type="text" value=""/></td> </tr> <tr> <td style="padding: 5px; text-align: center;">Total</td> <td style="padding: 5px; text-align: center;">100</td> </tr> </tbody> </table>		Color of the ball	No. of balls	White	<input style="width: 100%;" type="text" value="1"/>	Black	<input style="width: 100%;" type="text" value=""/>	Total	100																																																																			
Color of the ball	No. of balls																																																																											
White	<input style="width: 100%;" type="text" value="1"/>																																																																											
Black	<input style="width: 100%;" type="text" value=""/>																																																																											
Total	100																																																																											
<input style="background-color: red; color: white; padding: 5px 15px; border: none;" type="button" value="Submit"/>																																																																												
<table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th colspan="15" style="text-align: left; padding: 5px;">Draw(s) from this urn in the past periods</th> <th style="padding: 5px;">Summary</th> <th style="padding: 5px;">Count</th> <th style="padding: 5px;">Percentage (%)</th> </tr> </thead> <tbody> <tr> <td style="padding: 5px;">Period</td> <td style="padding: 5px;">1</td> <td style="padding: 5px;">2</td> <td style="padding: 5px;">3</td> <td style="padding: 5px;">4</td> <td style="padding: 5px;">5</td> <td style="padding: 5px;">6</td> <td style="padding: 5px;">7</td> <td style="padding: 5px;">8</td> <td style="padding: 5px;">9</td> <td style="padding: 5px;">10</td> <td style="padding: 5px;">11</td> <td style="padding: 5px;">12</td> <td style="padding: 5px;">13</td> <td style="padding: 5px;">14</td> <td style="padding: 5px;">15</td> <td style="padding: 5px;">White</td> <td style="padding: 5px;">1</td> <td style="padding: 5px;">20</td> </tr> <tr> <td style="padding: 5px;">Color drawn</td> <td style="padding: 5px;">White</td> <td style="padding: 5px;">Black</td> <td style="padding: 5px;">Black</td> <td style="padding: 5px;">Black</td> <td style="padding: 5px;">Black</td> <td style="padding: 5px;">-</td> <td style="padding: 5px;">-</td> <td style="padding: 5px;">-</td> <td style="padding: 5px;">-</td> <td style="padding: 5px;">-</td> <td style="padding: 5px;">-</td> <td style="padding: 5px;">-</td> <td style="padding: 5px;">-</td> <td style="padding: 5px;">-</td> <td style="padding: 5px;">-</td> <td style="padding: 5px;">Black</td> <td style="padding: 5px;">4</td> <td style="padding: 5px;">80</td> </tr> <tr> <td colspan="16"></td> <td style="padding: 5px;">Total draws</td> <td style="padding: 5px;">5</td> <td style="padding: 5px;">100</td> </tr> </tbody> </table>		Draw(s) from this urn in the past periods															Summary	Count	Percentage (%)	Period	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	White	1	20	Color drawn	White	Black	Black	Black	Black	-	-	-	-	-	-	-	-	-	-	Black	4	80																	Total draws	5	100
Draw(s) from this urn in the past periods															Summary	Count	Percentage (%)																																																											
Period	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	White	1	20																																																										
Color drawn	White	Black	Black	Black	Black	-	-	-	-	-	-	-	-	-	-	Black	4	80																																																										
																Total draws	5	100																																																										

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, \dots, G_{14} for sub-graph 1; G_0, \dots, 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 Session VII, respectively. In Session VII, draws are implemented per subject, i.e. Session VII produces 19 different paths of draw history, one for each subject in Session VII. Each blue dot-line in sub-graphs 12-30 represents the guess game responses of one specific subject in Session VII. A red reference line of 50 is added.

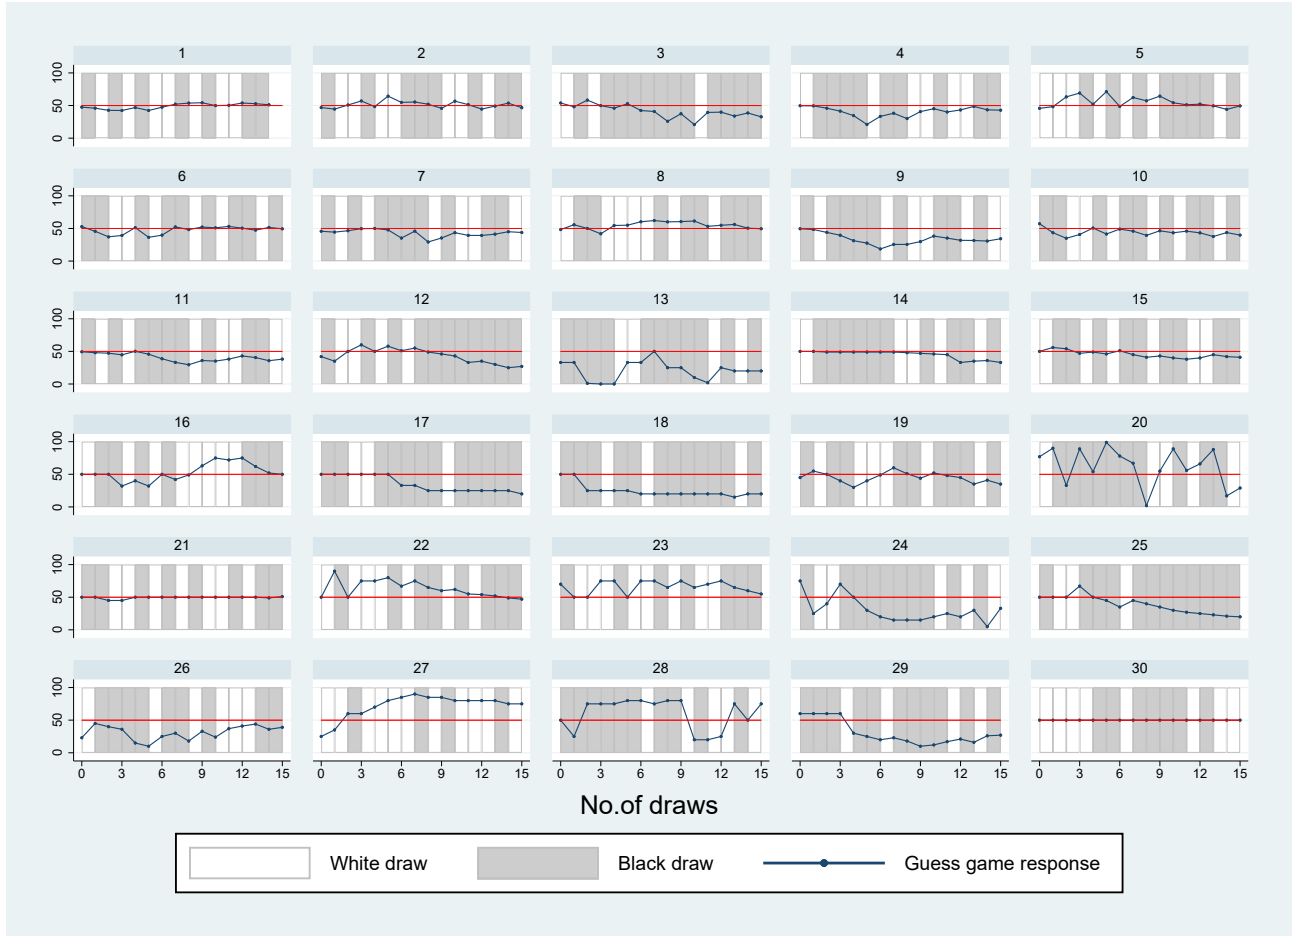


Figure 4: Rate of rational belief updates

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: 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 7: 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 perfectly/imperfectly following Bayes' rule. 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 8: 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 perfectly/imperfectly following Bayes' rule. 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

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

Since

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

hence

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

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

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

Plugging (24) into (21) yields:

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

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

Appendix B: Proposition 2

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

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

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:

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

FOC leads to:

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

SOC leads to:

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

Hence, the internal solution is:

$$\theta^* = \frac{k}{n} \quad (30)$$

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

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

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

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

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

Appendix C: Belief estimation: Equation (11)

Proof. For a given subject after a given number of draws, her prior/posterior distribution recovered from the learning strategy 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:

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

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

$$\left[\frac{\alpha_n - 1}{\alpha_n + \beta_n - 2} \right]^{\alpha_n - 1} \left[1 - \frac{\alpha_n - 1}{\alpha_n + \beta_n - 2} \right]^{\beta_n - 1} = \left[\frac{\text{white}_n}{100} \right]^{\alpha_n - 1} \left[1 - \frac{\text{white}_n}{100} \right]^{\beta_n - 1} \quad (34)$$

Plugging in Equation (9)(10) yields:

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

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

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

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

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, Lorian 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 Environmental Responsibility
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 Massenet	Macroprudential Policy in the Lab