

Updating Ambiguous Beliefs in a Social Learning Experiment

Roberta De Filippis, Antonio Guarino,
Philippe Jehiel and Toru Kitagawa*

Abstract

In our social learning experiment, each subject predicts the value of a good twice: after observing his predecessors' prediction ("first belief"); and after observing a private signal too. We find that subjects update on their signal asymmetrically: they weigh it in a Bayesian fashion when it confirms their first belief and overweight it when it contradicts it. This way of updating, incompatible with Bayesianism, is explained by ambiguous beliefs (multiple priors on the predecessor's rationality) and a generalization of the ML updating rule. Our experiment offers a better understanding of the overweighting of private information found in previous studies.

1 Introduction

In many economic and social situations we make decisions having our own information about which action may be the best one and also observing the decisions of others who faced a similar problem in the past. An investment decision or the purchase of a new product or service are just among the many examples of these situations. It is in fact difficult to think of cases in which we are the first to make a decision and have no information about how others have decided in the past. Observing the decision of others is useful, since we learn what others thought the best action was on the basis of the information they had.

Given the pervasiveness of the phenomenon, it is of course of crucial importance to answer questions such as: how do people make inferences from the decision of other agents? How do they combine the information coming from that observation with their own private information?

*De Filippis: Department of Economics, University College London (e-mail: roberta.filippis.10@ucl.ac.uk); Guarino: Department of Economics, University College London (e-mail: a.guarino@ucl.ac.uk); Jehiel: Paris School of Economics and Department of Economics, University College London (e-mail: jehiel@enpc.fr); Kitagawa: Department of Economics, University College London (e-mail: t.kitagawa@ucl.ac.uk). We thank participants in seminars at University College London, European University Institute, New Economic School, Bocconi, Vienna, Stanford, Berkeley, San Diego, Michigan, Wisconsin and at the LBEG workshop for useful comments and suggestions. Guarino gratefully acknowledges the financial support of the ERC and of the INET, and the hospitality of the EUI and of EIEF.

A large literature, starting with the seminal work of Banerjee (1992) and Bikhchandani *et al.* (1992), has investigated this type of issues and pointed out that observational learning can lead to phenomena such as herding and informational cascades. Theories of observational learning, typically relying on the assumption of full rationality, have been extensively tested through laboratory experiments, since the experiment of Anderson and Holt (1997). A common finding of these experiments is that participants tend to put relatively more weight on their private information than on the public information contained in the choices of other participants, as compared with the rational benchmark as well as with what would have been optimal given the actual behavior of participants in the laboratory. Nöth and Weber (2003), for instance, using an ingenious design in which subjects observe signals of different precision, conclude that “participants put too much weight on their private signal compared to the public information which clearly indicates the existence of overconfidence.” Çelen and Kariv (2004), in a study aimed at distinguishing informational cascades from herd behavior, also find that “subjects give excessive weight to their private information relative to the public information revealed by the behavior of others.” Goeree *et al.* (2007), revisiting the original Anderson and Holt (1997) experimental design with longer sequences of decision makers, analyze the data through the lenses of the Quantal Response Equilibrium (QRE) and also conclude that there is strong evidence of overweighting of the private information.

These experiments were designed to analyze herding or informational cascades. The observation of the overweighting of the private signal was a by-product of this analysis. In this paper we propose a new experiment, designed to elicit the beliefs of a subject after he observes another subject’s decision, and the updated belief after he receives private information. With this novel design, we can study in detail how people combine the information coming from the observation of others’ decisions with their private information. Our purpose is to analyze how well human subjects’ behavior conforms to Bayesian updating when they have to make inferences from a private signal and from the decision of another human subject, and to understand the determinants of the overweighting of the own information documented in previous studies (to the extent it shows up in our experiment too).

To be specific, in our experiment, we ask subjects to predict whether a good is worth 0 or 100 units, two events that are, *a priori*, equally likely. A first subject receives a noisy symmetric binary signal about the true value realization: either a “good signal”, which is more likely if the value is 100; or a “bad signal”, which is more likely if the value is 0. After receiving his signal, the subject is asked to state his belief on the value being 100.¹ To elicit his belief we use a quadratic scoring rule. We then ask a second subject to make the same type of prediction based on the observation of the first subject’s decision only. Finally, we provide the second subject with another, conditionally independent, signal about the

¹Specifically, subjects are asked to choose a number between 0 and 100. The number is the probability (expressed as a percentage) that the value is 100.

value of the good and ask him to make a new prediction.

Whereas in previous experiments, subjects' beliefs are hidden under a binary decision, in our experiment we elicit them. The belief of the first subject tells us how he updates from the observation of a private signal. The first action of the subject at time 2 gives us the "first belief" that he forms upon observing the predecessor's action only. His second decision gives us the "posterior belief" that he forms by observing his private signal. Asking subjects to make decisions in a continuous action space and eliciting a subject's beliefs both before and after receiving the private signal are novel features in the experimental social learning literature, which allow us to separately observe how a subject updates from observing others' actions and how this is combined with the subject's own private information.

The main results of our investigation are the following. First, after observing a private signal only, subjects do not show a particular bias in updating their beliefs. In particular, there is no systematic overweight of the signal. While there is a lot of heterogeneity in updating, the median subject's update is in line with Bayesian updating. Second, at time 2, after observing the predecessor's decision at time 1 only, subjects "discount" the informativeness of the predecessor's action, attaching to it a lower weight (as if the action were less informative than the signal on which it is based). This is so despite the fact that an action at time 1 above (respectively, below) 50 is almost as informative as the original signal observed by the time 1 subject (since subjects update very rarely in the wrong direction). Third, and most importantly, when at time 2 subjects observe their private signal, they update their belief in an asymmetric way. When the signal is in agreement with their first belief (e.g., when they first state a belief higher than 50% and then receive a signal indicating that the more likely value is 100), they weigh the signal as a Bayesian agent would do. When, instead, they receive a signal contradicting their first belief, they put considerably more weight on it.

In previous experiments on social learning, this asymmetry could not be observed. When subjects had a signal in agreement with the previous history of actions, they typically followed it. This decision is essentially uninformative for the experimenter on how subjects update their private information. In fact, on the basis of previous experimental results, one could have thought that overweighing private information is a general feature of human subjects' updating in this type of experiments. Our work shows that this is not the case, since it only happens when the private information contradicts the first belief.

This asymmetric updating, not identified in the previous literature, is incompatible with standard Bayesianism. The subject's "first belief" (i.e., his belief after observing the predecessor but before receiving the private signal) may differ from the theoretical (Perfect Bayesian Equilibrium) one if the subject at time 2 has a misconception of the precision of the signals or if he conceives the possibility that his predecessor's action may not perfectly reveal the private information he received, e.g., because of mistakes or boundedly rational behavior (no matter how bounded rationality is modelled, e.g., including the approaches of level k , cursed equilibrium or analogy-based expectation equilibrium). Whatever this

“first belief”, however, the subject should simply update it on the basis of the new information, giving the same weight to the signal, independently of its realization (whether or not it goes against the first belief), as this is a mere implication of the signals at time 1 and time 2 being independent conditional on the realization of the value of the good.²

It should be stressed that the asymmetry we observe cannot be explained in terms of standard psychological biases such as the confirmation bias, according to which subjects have a tendency to discount (maybe ignore with positive probability) new information in disagreement with their original view. Indeed, confirmation bias would imply that upon receiving a private information that contradicts the first belief, that information should be discounted (as compared to when the private information agrees with the first belief), but we observe the opposite asymmetry. In order to rule out any form of psychological bias (including confirmation bias and base rate neglect) that would be purely based on errors in signal processing and not on the multi-player aspect of our social learning experiment, we ran a control (individual decision making) treatment in which the same subject received two signals in sequence drawn according to the same process as in our social learning treatment, and reported his belief on the value of the good after the first and after the second signal. In this treatment, we observed heterogeneity (thereby supporting the view that subjects may attach subjective and dispersed beliefs to the precision of signals) but not the asymmetry in updating (thereby supporting the view that subjects had a reasonable understanding of the conditional independence feature highlighted above). This control treatment thus establishes experimentally that the asymmetric updating is intrinsically related to the social learning aspect of our experiment.

To explain the asymmetric updating we observe in our social learning experiment, we propose that, like in models of ambiguity, the subject at time 2 may entertain several possible theories about time 1 subject’s “rationality” (where a subject is considered “rational” if he chooses an action higher than 50 when he receives a good signal and lower than 50 when he receives a bad signal; a “noise” subject, in contrast, chooses any action independently of the signal).³ Moreover, each time he has to make a decision, he selects the theory that maximizes the likelihood of the realized observations. Based on this selected theory about the rationality of the predecessor, subject 2 updates his belief in a standard Bayesian fashion (possibly using subjective representations of the precision of the signals).

²In other words, even if subject 2 may perceive subject 1 as being irrational, there should be no asymmetric updating as long as the cognitive type of subject 1 is determined independently of the realization of the state and of the signal drawn by subject 2. This is so since, conditional on the state, the first action and the second signal are independent events. The conditional independence should be clearly understood by the subjects given that in the instructions we indicate that the balls at times 1 and 2 are drawn (with replacement) from an urn whose composition depends only on the value of the good. We discuss this issue in more detail in the following pages, when we describe the control treatment and the results.

³As we will discuss in the next section, we use this minimal requirement in our definition of rationality, since it is enough to infer the signal from the subject’s action, which is the only thing that matters.

Intuitively, this explains the asymmetry we observe for the following reason. Imagine a subject observing the predecessor taking an action greater than 50 (i.e., an action that presumably comes from a good signal, indicating the value is 100). Suppose he considers that the event is most likely under the prior that the predecessor is rational and, therefore, chooses his own action (his “first belief”) accordingly. After he observes a private signal confirming his first belief (that the value is more likely to be 100), the subject remains confident that the predecessor was rational, that is, sticks to the same prior on the predecessor’s rationality. He updates on that prior belief and so the weight he puts on the signal seems identical to that of a Bayesian agent. Consider now the case in which he receives a signal contradicting his first belief (i.e., a bad signal, indicating that the more likely value is 0). In such a case he now deems it an unlikely event that the predecessor was rational. In other words, he selects another prior belief on the predecessor’s rationality, giving a much higher probability to his predecessor being noise. Once he has selected this new prior on the predecessor’s rationality, he updates on the basis of the signal realization. This time it will look like he puts much more weight on the signal, since the signal first has made him change the prior on the rationality of the predecessor (becoming more distrusting) and then update on the basis of that prior.

Going back to our original research question, the overweighting of private information documented by previous social learning experiments seems to be not just the result of distrust of predecessors’ rationality, but the result of ambiguity entertained by subjects about such rationality. This is true since a single prior of distrust would not explain the asymmetry in updating that we observe in our social learning treatment, an asymmetry that, instead, we do not observe in the individual decision making treatment. The idea that there is some distrust about others’ rationality and some uncertainty about it, seems intuitively appealing to us. We are not aware of any alternative explanation to the observed asymmetry in updating.

Given our proposed explanation in terms of multiple priors (on the predecessor’s rationality), our experiment is also relevant for the debate in decision theory on how to update multiple priors. There are two main models of updating that have been proposed and axiomatized in the multiple prior setting (see Gilboa and Marinacci, 2013 for a survey). One is the Maximum Likelihood Updating (MLU) rule (axiomatized by Gilboa and Schmeidler, 1993) on which our method of updating that we refer to as “Likelihood Ratio Test Updating” builds. The other is the Full Bayesian Updating (FBU) model, in which agents have multiple priors and update prior by prior. Typically, after updating all priors, an agent makes his decision by using Maxmin Expected Utility (see Pires, 2002 for an axiomatization).

In our structural econometric analysis—another methodological contribution of our paper—we compare the Bayesian Updating (BU) model to the Likelihood Ratio Test Updating (LRTU) model (a generalization of MLU) and to the FBU model.⁴ We show robust statistical evidence that the LRTU’s model fits our data

⁴When we designed the experiment, we were interested in testing the BU model. Since the

significantly better than the other two models. A novelty of our econometric analysis is that we make use of the observations in the control treatment to identify non-parametrically the distribution of unobserved heterogeneity in the subjective beliefs about the precision of the signal. We can perform a robust statistical comparison of the three models without relying on any distribution specification of heterogeneity. To the best of our knowledge, using a control treatment to account for heterogeneity in individual decisions in the laboratory has not been done in previous experimental work.

From a decision theory viewpoint, it is important to remark that our LRTU model differs from the MLU model axiomatized by Gilboa and Schmeidler (1993) in a crucial aspect. Whereas in Gilboa and Schmeidler (1993) it is assumed that, once the prior is selected, the agent sticks to it (as if ambiguity were totally eliminated after the agent receives a first piece of information), in our model we let the agent change the prior after receiving extra information (as a statistician would do, using new data to select the most likely prior — in the statistics literature this dates back, among others, to Good, 1965). The possibility that the set of priors does not collapse to a single prior is contemplated in Epstein and Schneider (2007)’s model of dynamic updating. Such a method of updating appears also in a more recent work by Ortoleva (2012), who obtains it as one possible representation theorem by relaxing dynamic consistency. We will further discuss Ortoleva (2012)’s model in Section 7.

We believe the results of our experiment should inform future work on the updating of multiple priors and belief dynamics. We also find it interesting that our work shows an experimental application of the ambiguous beliefs literature beyond the classical Ellsberg experiment.⁵

The paper is organized as follows. Section 2 describes the theoretical model of social learning and its (Perfect Bayesian) equilibrium predictions. Section 3 presents the experiment. Section 4 contains the results. Section 5 illustrates how multiple priors can, theoretically, lead to asymmetric updating. Section 6 illustrates the econometric analysis. Section 7 offers further discussion of our

results were at odds with the BU model (even allowing for subjective beliefs), we considered these alternative theories which dispense with the assumption of one prior and which pre-existed our experimental work.

⁵We cannot compare our results to other work, since the experimental literature on updating in a context of ambiguity is still to be developed. To our knowledge, there are only two related experiments, (Cohen *et al.*, 2000 and Dominiak *et al.*, 2012), but they are very different from ours and a comparison is difficult. Both studies consider Ellsberg’s original urn experiment in which the proportion of yellow balls is known but only the aggregate proportion of blue and green balls is known. Cohen *et al.* (2000) ask subjects to choose between acts (specifying rewards as a function of the drawn ball color) conditional on learning that the drawn ball is not green. Dominiak *et al.* (2012) conduct a similar experiment, although their focus is on whether subjects violate dynamic consistency and/or consequentialism (consequentialism is assumed in Cohen *et al.*, 2000). Both experiments find that the proportion of subjects whose behavior is compatible with FBU is higher than that compatible with MLU. Our findings derived in social learning environments do not support FBU. In our experiment, two pieces of information (predecessor’s action, then private signal) arrive over time, and this is a crucial ingredient of our design which has no counterpart in these other two experiments (in which the first choice does not require any inference). Since the action space is rich in our experiment, we can observe beliefs, which is impossible in the other experiments.

findings. Section 8 concludes. An Appendix contains additional material.

2 The Theoretical Model

We now describe the simple theoretical social learning model on which the experiment is based and then illustrate the experimental procedures.

In our economy there is a good that can take two values, $V \in \{0, 100\}$. The two values are equally likely. There are two agents who make a decision in sequence. The decision consists in choosing a number in the interval $[0, 100]$. Each agent t ($t = 1, 2$) receives a private signal $s_t \in \{0, 1\}$ correlated with the true value V . Specifically, each agents receives a symmetric binary signal distributed as follows:

$$\Pr(s_t = 1 \mid V = 100) = \Pr(s_t = 0 \mid V = 0) = 0.7.$$

This means that, conditional on the value of the good, the signals are identically and independently distributed over time, with precision 0.7. Since the signal $s_t = 1$ increases the probability that the value is 100, we will also refer to it as the good signal, and to $s_t = 0$ as the bad signal. Agent 1 (randomly chosen) observes the signal s_1 and takes an action a_1 . At time 2, agent 2 observes a_1 and takes a first action a_2^1 . He then observes the private signal s_2 and takes a second action a_2^2 . The agent's payoff from each choice depends on his choice and on the value of the good. The payoff is quadratic and, in particular, equal to $-(V - a_t^j)^2$.⁶ Given his information I_t^j , the agent chooses a_t^j to maximize his expected payoff $E^S[-(V - a_t)^2 | I_t^j]$ (where the superscript S stands for "subjective"). Therefore, his optimal action is $a_t^{j*} = E^S(V | I_t^j)$. When the subjective beliefs coincide with correct beliefs, what we are describing coincides with a Perfect Bayesian Equilibrium (PBE). In the PBE, agent 1 chooses 70 upon observing $s_1 = 1$ and 30 upon observing $s_1 = 0$. Since each action perfectly reveals the signal realization, observing the action is identical to observing the signal. Therefore, agent 2 chooses his first action such that $a_2^1 = a_1$. After observing the private signal s_2 , the agent updates his belief and chooses $a_2^2 = E(V | a_1 = 70, s_2 = 1) = 84.48$, $a_2^2 = E(V | a_1 = 30, s_2 = 0) = 15.52$ and $a_2^2 = E(V | a_1 = 30, s_2 = 1) = E(V | a_1 = 70, s_2 = 0) = 50$.

Allowing agents to have subjective views about the precisions $q_t^S \in (0.5, 1]$ of signals (which we will later use to explain the observed heterogeneity of behavior) does not change the analysis qualitatively. As long as agent 1 thinks the signal is informative ($q_1^S > 0.5$), he updates in the "correct direction" although not exactly as in the PBE. Agent 2 is still able to infer the signal realization from the action, and can form his own belief, $a_2^1/100$, which again may depend on his subjective belief on the precision of signal 1. Note that agent 2 may also have subjective beliefs on the "rationality" of agent 1, that is, on his ability to

⁶Here, and in the following analysis, we use the subscript $t = 1, 2$ to indicate the first and second period, and the superscript $j = 1, 2$ to indicate the first and second subperiod in which agent 2 makes a decision; the superscript j is immaterial when $t = 1$.

read the signal correctly and update in the correct direction, above or below 50. If the subject entertains the possibility that, with some probability, the predecessor updates incorrectly, he will form the expectation on the value of the good accordingly. We are qualifying a subject as “rational” as long as he updates in the correct direction, since the only thing that agent 2 has to learn from subject 1 is indeed the signal realization (since the objective precisions of the signals are known and do not have to be learned), and this is revealed under the minimal requirement that the agent updates in the right direction. In any case, whatever model of reasoning agent 2 uses to compute his expected value, that is what he expresses by stating $a_2^1/100$. Once he has stated this belief, he then updates it, possibly using a subjective precision q_2^5 .

One can, of course, go beyond the Bayesian paradigm just discussed (even allowing for subjective views both about the precisions of signals and about the rationality of the predecessor). Specifically, one can consider a model in which subject 2 may have ambiguous beliefs on subject 1’s rationality (along the line of the literature with multiple priors).⁷ In such a case, as we anticipated in the Introduction, it is important to specify how a subject updates beliefs. For exposition motives, we find it convenient to postpone this discussion to Section 5, after we will have presented the main results.

3 The Experiment and the Experimental Design

3.1 The Experiment

This work is part of a larger experimental project, designed to answer several research questions. The experiment was conducted with multiple, rather than with just two periods of decision making. We describe the entire experiment subjects participated in, even though we will then only focus on their decisions in the first two periods.

We ran the experiment in the ELSE Experimental Laboratory at the Department of Economics at University College London (UCL) in the fall 2009, winter 2010, fall 2011 and spring 2014. The subject pool mainly consisted of undergraduate students in all disciplines at UCL. They had no previous experience with this experiment. In total, we recruited 267 students. Each subject participated in one session only.

The sessions started with written instructions given to all subjects. We explained to participants that they were all receiving the same instructions. Subjects could ask clarifying questions, which we answered privately. The experiment was programmed and conducted with a built-on-purpose software.

Here we describe the baseline treatment (SL1). In the next section, we will explain the experimental design. We ran five sessions for this treatment. In each session we used 10 participants. The procedures were the following:

⁷Another departure from Bayesianism would be to allow for ambiguity on the precision of signals. This would be a more far fetched hypothesis, since the composition of the urns is known. In contrast, nothing is known about the rationality of subjects, which, potentially, leaves subject 2 in a situation of uncertainty about the rationality of the predecessor.

1. Each session consisted of fifteen rounds. At the beginning of each round, the computer program randomly chose the value of a good. The value was equal to 0 or 100 with the same probability, independently of previous realizations.
2. In each round we asked all subjects to make decisions in sequence, one after the other. For each round, the sequence was randomly chosen by the computer software. Each subject had an equal probability of being chosen in any position in the sequence.
3. Participants were not told the value of the good. They knew, however, that they would receive information about the value, in the form of a symmetric binary signal. If the value was equal to 100, a participant would receive a “green ball” with probability 0.7 and a “red ball” with probability 0.3; if the value was equal to 0, the probabilities were inverted. That is, the green signal corresponded to $s_t = 1$ and the red signal to $s_t = 0$, the signal precision q_t was equal to 0.7 at any time.
4. As we said, each round consisted of 10 periods. In the first period a subject was randomly chosen to make a decision. He received a signal and chose a number between 0 and 100, up to two decimal points.
5. The other subjects observed the decision made by the first subject on their screens. The identity of the subject was not revealed.
6. In the second period, a second subject was randomly selected. He was asked to choose a number between 0 and 100, having observed the first subject’s choice only.
7. After he had made that choice, he received a signal and had to make a second decision. This time, therefore, the decision was based on the observation of the predecessor’s action and of the private signal.
8. In the third period, a third subject was randomly selected and asked to make two decisions, similarly to the second subject: a first decision after observing the choice of the first subject and the second choice of the second subject; a second decision after observing the private signal too. The same procedure was repeated for all the remaining periods, until all subjects had acted. Hence, each subject, from the second to the tenth, made two decisions: one after observing the history of all (second) decisions made by the predecessors; the other after observing the private signal too.⁸
9. At the end of the round, after all 10 subjects had made their decisions, subjects observed a feedback screen, in which they observed the value of the good and their own payoff for that round. The payoffs were computed

⁸As we explained above, the experiment was designed to address many research questions. Here we describe the entire experiment subjects participated in, although we focus our analysis on periods 1 and 2 only.

as $100 - 0.01(V - a_t)^2$ of a fictitious experimental currency called “lira.” After participants had observed their payoffs and clicked on an OK button, the software moved to the next round.

Note that essentially we asked subjects to state their beliefs. To elicit the beliefs, we used a quadratic scoring function, a quite standard elicitation method. In the instructions, we followed Nyarko and Schotter (2002) and explained to subjects that to maximize the amount of money they could expect to gain, it was in their interest to state their true belief.⁹

As should be clear from this description, compared to the existing experimental literature on social learning / informational cascades / herd behavior, we made two important procedural changes. First, in previous experiments subjects were asked to make a decision in a discrete (typically binary) action space, whereas we ask subjects to choose actions in a very rich space which practically replicates the continuum. This allows us to elicit their beliefs, rather than just observing whether they prefer one action to another.¹⁰ Second, in previous experiments subjects made one decision after observing both the predecessors and the signal. In our experiment, instead, they made two decisions, one based on public information only and one based on the private information as well.¹¹

To compute the final payment, we randomly chose (with equal chance) one round among the first five, one among rounds 6 – 10 and one among the last five rounds. For each of these round we then chose either decision 1 or decision 2 with equal chance (with the exception of subject 1, who was paid according to the only decision he made in the round). We summed up the payoffs obtained in these decisions and, then, converted the sum into pounds at the exchange rate of 100 liras for 7 GBP. Moreover, we paid a participation fee of £5. Subjects were paid in cash, in private, at the end of the experiment. On average, in this treatment subjects earned £21 for a 2 hour experiment.

3.2 Experimental Design

Social Learning. In addition to the social learning treatment (SL1) just described, we ran a second treatment (SL2) which only differed from the first because the signal had a precision which was randomly drawn in the interval

⁹This explanation helps the subjects, since they do not have to solve the maximization problem by themselves (and to which extent they are able to do so is not the aim of this paper). For a discussion of methodological issues related to elicitation methods, see the recent survey by Schotter and Trevino (2014).

¹⁰Within the discrete action space experiments, exceptions to the binary action space are the financial market experiments of Cipriani and Guarino (2005, 2009) and Drehman *et al.* (2005) where subjects can choose to buy, to sell or not to trade. In the interesting experimental design of Celen and Kariv (2004), subjects choose a cut off value in a continuous signal space: depending on the realization of the signal, one of the two actions is implemented (as in a Becker, DeGroot and Marschak, 1964, mechanism). That design allows the authors to distinguish herd behavior from informational cascades.

¹¹Cipriani and Guarino (2009) use a quasi strategy method, asking subject to make decisions conditional on either signal they might receive. Still, at each time, a subject never makes a decision based only on the predecessors’ decisions.

[0.7, 0.71] (instead of having a precision always exactly equal to 0.7). Of course, each subject observed not only the ball color but also the exact precision of his own signal. A third treatment (SL3) was identical to SL2, with the exception that instead of having sequences of 10 subjects, we had sequences of 4 subjects. Given the smaller number of subjects, each round lasted less time, obviously; for this reason, we decided to run 30 rounds per session, rather than 15. The results we obtained for times 1 and 2 for these three treatments are not statistically different (as we show in the next section and in the Appendix). For the purposes of this paper, we consider the three treatments as just one experimental condition. We will refer to it as the SL treatment. Drawing the precision from the tiny interval [0.7, 0.71], instead of having the simpler set up with fixed precision equal to 0.7, was only due to a research question motivated by the theory of Guarino and Jehiel (2013), where the precision is indeed supposed to differ agent by agent; this research question, however, is not the object of this paper. Reducing the length of the sequence to 4 subjects was instead motivated by the opportuneness to collect more data for the first periods of the sequence.

Individual Decision Making. In the social learning treatments subjects make decisions after observing private signals and the actions of others. Clearly, we may expect departures from the PBE even independently of the social learning aspect if subjects do not update in a Bayesian fashion. To control for this, we ran a treatment in which subjects observed a sequence of signals and made more than one decision.¹² Specifically, a subject received a signal (as subject 1 in the SL treatments) and had to make a choice in the interval [0, 100]. Then, with a 50% probability, he received another signal and had to make a second decision (similarly to the second decision of subject 2 in the SL treatments). Note that, at the cost of collecting less data, we decided not to ask subjects to make a second decision in all rounds. Our purpose was to make the task of the subject as close possible as possible to that of a subject in the SL treatments. In other words, we wanted the subject to make his first decision not knowing whether he would be asked to make a second one; this way, his first decision was in a condition very similar to that of subject 1 in the other treatments; once the subject was given another signal and was asked to make another decision, he was in a situation comparable to that of subject 2 in the SL treatments.

4 Results

Our main interest is in understanding how human subjects weigh private and public information. To this aim, we will focus on subjects' second decisions at time 2, that is, after they have observed both their predecessor's action and their private signal. Before doing so, however, we will briefly discuss the decisions of subjects at time 1 (when they have only observed a private signal) and the first

¹²This treatment was programmed and conducted with the software z-Tree (Fischbacher, 2007) in the fall 2014. The payment followed the same rules. The exchange rate was appropriately modified before each treatment so that, in expectation, subjects could receive a similar amount of money per hour spent in the laboratory.

Treatments	Signal Precision	Sequence	Subjects in a group	Groups	Participants	Rounds
SL1	0.7	10	10	5	50	15
SL2	[0.7,0.71]	10	10	5	49	15
SL3	[0.7,0.71]	4	4	5	20	30
IDM	0.7	1 or 2	-	-	36	30

Table 1: Treatments’ features. SL: Social Learning; IDM: Individual Decision Making. Note that in SL2 there are 49 subjects since onse session was run with 9 participants rather than 10 due to a last minute unavailability of one subject.

decisions of subjects at time 2, based on the observation of their predecessor’s choice only.

4.1 How do subjects make inference from their own signal only?

At time 1, a subject makes his decision on the basis of his signal only. His task—to infer the value of the good from a signal drawn from an urn—is the same in the SL and in the IDM treatments; for this reason we pool all data together (for a total of 1380 observations).¹³

Figure 1 shows the frequency of decisions at time 1, separately for the cases in which the signal the subject received was good or bad. The top panel refers to the case of a good signal. A high percentage of decisions (34.5%) are in line with Bayesian updating, deviating from it by less than 5 units; 19.5% of actions are smaller than the Bayesian one and 43.3% of actions are larger. Note, in particular, that in 9.4% of the cases subjects did not update their belief at all after seeing the signal, choosing an action exactly equal to 50. On the other hand, in 13% of the cases, subjects went to the boundary of the support, choosing the action 100. Finally, there is a small proportion (2.8%) of actions in the wrong direction (i.e., updating down rather than up).

The bottom panel refers to the bad signal. The picture looks almost like the mirror image of the previous one, with the mode around 30, masses of 12.8% in 50 and of 12.4% in 0, and other actions distributed similarly to what explained above.

One interpretation of these results is that subjects put different weights on the signal they receive (which is equivalent to subjects attaching to signals different, subjective precisions). A simple model that allows to quantify this

¹³We ran a Mann-Witney U test (Wilcoxon rank-sum test) on the medians of each session (the most conservative option to guarantee independence of observations) for the SL treatment and on the medians of each individual’s decisions in the IDM treatment; we cannot reject the null hypothesis that they come from the same distribution (p-value = 0.47). Note that we also ran the same test to compare the three SL treatments and we cannot reject the same hypothesis (at the 5% significance level) when we compare SL1 with SL2 (p-value = 0.5), SL1 with SL3 (p-value = 0.08), or SL2 with SL3 (p-value = 0.22).




Figure 1: Distribution of actions at time 1. The top (bottom) panel refers to actions upon receiving $s_1 = 1$ ($s_1 = 0$).

phenomenon is the following:

$$a_{1i} = 100 \left(s_{1i} \frac{q^{\alpha_{1i}}}{q^{\alpha_{1i}} + (1-q)^{\alpha_{1i}}} + (1-s_{1i}) \frac{(1-q)^{\alpha_{1i}}}{q^{\alpha_{1i}} + (1-q)^{\alpha_{1i}}} \right), \quad (1)$$

where $\alpha_{1i} \in \mathbb{R}$ is the weight put on the signal in observation i and the precision of the signal q is considered to be always 0.7.¹⁴ Note that for $\alpha_{1i} = 1$ expression (1) gives the Bayesian updating formula, and so $\alpha_{1i} = 1$ is the weight that a Bayesian agent would put on the signal. A value higher (lower) than 1 indicates that the subject overweights (underweights) the signal. For instance, for $\alpha_{1i} = 2$, the expression is equivalent to Bayesian updating after receiving two conditionally independent signals and can, therefore, be interpreted as the action of a Bayesian agent acting upon receiving two signals (with the same realization). A subject that does not put any weight on the signal ($\alpha_{1i} = 0$) of course does not update at all upon observing it ($a_{1i} = 50$), whereas a subject who puts an infinite weight on it chooses an extreme action ($a_{1i} = 0$ or

¹⁴Recall that a subject made many choices in the same experiment, since he participated in several rounds; the index i refers to the observation i at time 1, and not to the subject acting at that time. Of course the same subject could have chosen different weights in different decisions. Moreover, recall that in some sessions the exact precision of the signal was randomly drawn from $[0.7, 0.71]$ rather than being identical to 0.7. By using the exact precision we obtain, of course, almost identical results, with differences at most at the decimal point. We prefer to present the results for $q = 0.7$ for consistency with our analysis at time 2.

$a_{1i} = 100$), as if he were convinced that the signal fully reveals the value of the good. Finally, a negative value of α_{1i} indicates that the subject misreads the signal, e.g., interpreting a good signal as a bad one.

Table 2 reports the quartiles of the distribution of the computed α_{1i} .¹⁵ Note that the median α_{1i} is 1, indicating that the median subject is actually Bayesian.

	1st Quartile	Median	3rd Quartile
α_{1i}	0.73	1.00	2.05

Table 2: Distribution of weights on private signal for actions at time 1.

The table shows the quartiles of the distribution of weights on private signal for actions at time 1.

In this analysis, we have allowed for heterogeneous weights on the signal and assumed that subjects did state their beliefs correctly. Of course, another approach would be to take into account that subjects could have made mistakes while reporting their beliefs, as in the following model:

$$a_{1i} = 100 \left(s_{1i} \frac{q^{\alpha_1}}{q^{\alpha_1} + (1-q)^{\alpha_1}} + (1 - s_{1i}) \frac{(1-q)^{\alpha_1}}{q^{\alpha_1} + (1-q)^{\alpha_1}} \right) + \varepsilon_{1i}, \quad (2)$$

where the weight on the signal is the same for all subjects but each subject makes a random mistake ε_{1i} . It is easy to show that, as long as the error term has zero median, the estimated median α_1 in this model coincides with the median α_{1i} computed above.

Of course, other interpretations are possible. One may, for instance, argue that the fact that a subject chooses 70, while compatible with Bayesian updating, is not necessarily indication that he is a proper Bayesian: he may be choosing 70 simply because that is the precision of his signal. The fact that the median subject is Bayesian for a bad signal too, however, lends some credibility to the fact that the subjects are doing more than just inputting their signal precision. Action 50 may also be the result of different heuristics. A subject may feel that one signal alone is not enough for him to make any update; or perhaps he is happy to choose the least risky action. The extreme actions, on the other hand, may be the expression of a “guessing type” who, despite the incentives given in the laboratory, simply tries to guess the most likely outcome. It should be noticed, though, that of all subjects who acted at time 1 more than once, only one chose an extreme action (0 or 100) every time; similarly, only 5.7% of them chose the action 50 every time.¹⁶

As we said in the Introduction, in previous social learning experiments, deviations from equilibrium have been interpreted sometimes as subjects being overconfident in their own signal. Our analysis shows that there is much heterogeneity in the way subjects update their beliefs after receiving a signal. Despite

¹⁵When $a_{1i} = 0$ or 100, we compute α_{1i} by approximating $a_{1i} = 0$ with ε and $a_{1i} = 100$ with $100 - \varepsilon$ (with $\varepsilon = 0.01$). We prefer to report the quartiles rather than the mean or other statistics whose computations are affected by the approximation of α_{1i} .

¹⁶We will comment more on risk preferences in Section 4.3.

these subjective beliefs, there is no systematic bias to overweight or underweight the signal. As a matter of fact, the median belief is perfectly in line with Bayesian updating.

4.2 How do subjects make inference from their predecessor's action?

We now turn to the question of whether and how subjects infer the value of the good from the predecessor's action. We focus on the first decision at time 2 (denoted by a_2^1) since it is based on the observation of that action only. Of course, here we only consider the data from the SL treatment.

A subject at time 2 has to infer which signal his predecessor received on the basis of the action he took. We know from the previous analysis that only rarely (in 3.5% of the cases), subjects at time 1 updated in the "wrong direction" (i.e., chose an action greater (lower) than 50 after observing a bad (good) signal). Therefore, subjects at time 2 could have simply considered an action strictly greater (or lower) than 50 as a good (or bad) signal.

We have pooled together all cases in which the observed choice at time 1 was greater than 50 and, similarly, all cases in which it was lower than 50 (see Figure 2). Compared to Figure 1, Figure 2 shows a higher mass for $a_2^1 = 50$ and a lower one around 70 or 30 (for the case of $a_1 > 50$ and $a_1 < 50$, respectively). When the subject at time 1 had chosen $a_1 = 50$, perhaps not surprisingly, the distribution has a large mass at 50.

Figure 3 shows the difference between the actions a_2^1 and the corresponding action a_1 that a subject has observed (excluding the cases in which $a_1 = 50$). If subjects simply imitated the predecessor's decision, all the mass would be concentrated around zero. While there are approximately 30% of cases in which this happens, we observe that the distribution has in fact a larger mass below 0, indicating that subjects had the tendency to choose lower values than the predecessors'.¹⁷

We replicated the model discussed in the previous section, by replacing the case in which the subject observed a good signal with the case in which the subject observed $a_1 > 50$, and so chose a_{2i}^1 such that

$$a_{2i}^1 = 100 \frac{q^{\alpha_{2i}^1}}{q^{\alpha_{2i}^1} + (1-q)^{\alpha_{2i}^1}}; \quad (3)$$

analogously, for the case in which he observed $a_1 < 50$, he chose a_{2i}^1 such that

$$a_{2i}^1 = 100 \frac{(1-q)^{\alpha_{2i}^1}}{q^{\alpha_{2i}^1} + (1-q)^{\alpha_{2i}^1}}. \quad (4)$$

Essentially, in this model we are assuming that a subject considers actions higher (or lower) than 50 as good (bad) signals with the same precision $q = 0.7$. By

¹⁷In 30% of the observed cases, imitation coincides with the Bayesian action. There is no specific pattern in the remaining cases.




Figure 2: Distribution of first actions at time 2 (the top panel refers to $a_1 > 50$, the middle to $a_1 = 50$ and the bottom to $a_1 < 50$).




Figure 3: Distribution of the difference between a_2^1 and the corresponding a_1 .

applying this model, we obtain the results reported in Table 3. The median weight is (slightly) lower than 1 and the first and third quartiles are 0.13 and 1.4 (versus 0.81 and 2.05 at time 1) reflecting the fact that subjects in these treatments seem to “discount” to some extent the information contained in the predecessor’s action.^{18 19}

It should be noticed that we could expect to observe the same distribution at time 1 and at time 2 under two different models. One model is that subjects at time 2 perfectly infer the signal from the observed action at time 1 and weigh the signal in the same heterogeneous ways at time 1 and time 2. The other is that subjects simply imitate the predecessors’ actions. Clearly both models are rejected by our data. To explain the data we need a model in which a subject acting at time 2 has subjective beliefs on how trustworthy the predecessor is (i.e., on how frequently the predecessor decision to update up or down from 50 reflects a good or bad signal).

To investigate this issue further, we computed the weights separately for different classes of a_1 , as illustrated in Table 3.²⁰

	1st Quartile	Median	3rd Quartile
α_2^1	0.13	0.94	1.4
α_2^1 (upon observing $50 < a_{1i} \leq 66.7$)	0	0.48	0.9
α_2^1 (upon observing $66.7 < a_{1i} \leq 83.4$)	0	0.89	1.33
α_2^1 (upon observing $a_{1i} > 83.4$)	0.9	1.31	2.8

Table 3: Distribution of weights for first actions at time 2.

The table shows the quartiles of the distribution of weights for first actions at time 2. The action at time 1 is considered as a signal (of precision 0.7) for the subject at time 2.

As one can see, subjects have the tendency to “discount” the actions close to 50 ($50 < a_{1i} \leq 66.7$) and, although less, those in a neighborhood of the Bayesian one ($66.7 < a_{1i} \leq 83.4$). They do not discount, instead, more extreme actions. This behavior is in line with a model of subjective beliefs in which subjects expect error rates to be inversely proportional to the cost of the error, since the expected cost of an action against the signal is of course increasing in the distance from 50. A well known model in which errors are inversely related to their costs is the Quantal Response Equilibrium (which also assumes expectations are rational). Our results are, however, not compatible with such

¹⁸We considered the medians of each session for the SL treatment and of each individual’s decisions in the IDM treatment for a_1 ; and the medians of each session for the SL treatment for a_2^1 ; we reject the null hypothesis that they come from the same distribution (p-value = 0.014). We repeated the same test considering only the IDM treatment for a_1 ; again, we reject the null hypothesis (p-value = 0.015).

¹⁹Discounting the predecessor’s action is found, in a stronger way, in the experiment by Çelen and Kariv (2004). They ask subjects at time 2 to report a threshold value that depends on what they learn from the first subject’s choice. Çelen and Kariv (2004, p.493) find that “subjects tend to undervalue sharply the first subjects’ decisions.”

²⁰We have chosen the cut-off points 66.7 and 83.4 simply to obtain intervals of equal length. We tried alternative cut-off points and did not find significant differences in the results.

a theory in that expectations about time 1 error rates are not correct. Indeed, the error rate at time 1 is very small. With subjects at time 1 choosing an action against their signal in 3.5% of the cases only, a Bayesian agent would have a belief on the value of the good being 100 equal to $\Pr(V = 100|a_1 > 50) = \frac{(0.7)(0.965)+(0.3)(0.035)}{(0.7)(0.965)+(0.3)(0.035)+(0.7)(0.035)+(0.3)(0.965)} = 69$, which barely changes from the case of no mistakes. Essentially, to explain our data, we need a model of incorrect subjective beliefs. In one such model, a subject at time 1 can be either rational (always updating in the correct direction) or noise (choosing any number independently of the signal). If a noise type chooses more frequently actions close to 50 (e.g., because he chooses actions as in a Normal distribution centered around 50) and a rational type chooses more frequently more extreme actions, letting a subject at time 2 having (incorrect) subjective beliefs on the proportion of these two types can lead to the observed results. We will illustrate this model in Section 6.

4.3 How do subjects weigh their signal relative to their predecessor’s action?

As we said in the Introduction, in the experimental social learning literature there is a long debate about how subjects weigh their own signal with respect to the public information contained in the predecessors’ actions. Several studies (e.g., Nöth and Weber, 2003) conclude that subjects are “overconfident” in that they put more weight on their signal than they should (according to Bayes’ rule). Our previous analysis shows that subjects do not have a systematic bias in overweighting their signal when it is the only source of information. We now study how they weigh it at time 2, after having observed their predecessor’s action. Time 2 offers the possibility of studying this issue in a very neat way. In the subsequent periods, the analysis becomes inevitably more confounded, since subjects may take the sequence of previous actions into account in a variety of ways (since their higher order beliefs on the predecessors’ type matter too). At time 2, instead, the only source of information for the subject is the predecessor’s action and the own signal.

As we already mentioned in the Introduction, we will refer to the first action that subjects take at time 2 as their “first belief” and to the second as their “posterior belief.” Figure 4 shows the frequency of posterior beliefs conditional on whether the subject received a signal confirming his first belief (i.e., $s_{2i} = 1$ after an action $a_{2i}^1 > 50$ or $s_{2i} = 0$ after an action $a_{2i}^1 < 50$) or contradicting it (i.e., $s_{2i} = 1$ after an action $a_{2i}^1 < 50$ or $s_{2i} = 0$ after an action $a_{2i}^1 > 50$).²¹ The figure is obtained after transforming an action $a_{2i}^1 < 50$ into $100 - a_{2i}^1$ and the

²¹In this analysis we exclude the cases in which the action at time 1 was uninformative ($a_{1i} = 50$). We do study the case in which a subject at time 2 observed an informative action at time 1 and chose $a_{2i}^1 = 50$; in this case we distinguish whether the action observed at time 1 confirmed or contradicted the realization of the signal s_{2i} . Note that an alternative definition of confirming and contradicting signal would be in reference to a_1 rather than to a_2^1 . This would not affect our results, since the difference is in one observation only (in which $a_2^1 > 50$ and $a_1 < 50$).




Figure 4: Distribution of a_2^2 given $a_2^1 > 50$ and a confirming (top panel) or contradicting (bottom panel) s_2 .

corresponding signal s_{1i} into $1 - s_{1i}$.

If subjects acted as in the PBE, in the case of confirming signal we would observe the entire distribution concentrated on 84. The empirical distribution shows much more heterogeneity, of course. Nevertheless, the median action as well as the mode are indeed close to the PBE. For the contradicting signal, the picture is rather different. Whereas in the PBE we would observe the entire distribution concentrated on 50, the empirical distribution looks very asymmetric around 50, with more than 70% of the mass below 50. To understand these results, we compute the weight that the subject puts on his signal by using our usual model of updating:

$$a_{2i}^2 = 100 \frac{q^{\alpha_{2i}^2} \frac{a_{2i}^1}{100}}{q^{\alpha_{2i}^2} \frac{a_{2i}^1}{100} + (1-q)^{\alpha_{2i}^2} \left(1 - \frac{a_{2i}^1}{100}\right)}, \quad (5)$$

when he observed $s_{2i} = 1$ and, analogously,

$$a_{2i}^2 = 100 \frac{(1-q)^{\alpha_{2i}^2} \frac{a_{2i}^1}{100}}{(1-q)^{\alpha_{2i}^2} \frac{a_{2i}^1}{100} + q^{\alpha_{2i}^2} \left(1 - \frac{a_{2i}^1}{100}\right)}, \quad (6)$$

when he observed $s_{2i} = 0$.

Table 4 reports the results.²² While in the case of a confirming signal the

²²The value of α_{2i}^2 is undetermined when $a_{2i}^1 = 100$, therefore we exclude these cases. When $a_{2i}^2 = 100$ we use the same approximation as previously discussed.

median subject puts only a slightly lower weight on the signal than a Bayesian agent would do, in the case of a contradicting signal, the weight is considerably higher, 1.70.²³ The different weight is observed also for the first and third quartiles. Essentially, subjects update in an asymmetric way, depending on whether the signal confirms or not their first beliefs: contradicting signals are overweighted with respect to Bayesian updating.²⁴

	1st Quartile	Median	3rd Quartile
α_2^2	0.68	1.16	2.04
α_2^2 (upon observing confirming signal)	0.54	0.96	1.35
α_2^2 (upon observing contradicting signal)	1.00	1.70	2.73

Table 4: Distribution of weights on the own signal in the SL treatment.

The table shows the quartiles of the distribution of the weight on the own signal for the second action at time 2 in the SL treatment. The data refer to all cases in which the first action at time 2 was different from 50.

	1st Quartile	Median	3rd Quartile
α_2^2	0.64	1.08	2.07
α_2^2 (upon observing confirming signal)	0.64	1.30	2.48
α_2^2 (upon observing contradicting signal)	0.93	1.00	1.76

Table 5: Distribution of weights on the own signal in the IDM treatment.

The table shows the quartiles of the distribution of the weight on the own signal for the action at time 2 in the IDM treatment. The data refer to all cases in which the action at time 1 was different from 50.

Of course, one may wonder whether this result is due to the social learning aspect of our experiment or, instead, is just the way human subjects update upon receiving two consecutive signals. To tackle this issue, we consider subjects' behavior in the IDM treatment, as reported in Table 5. As one can see, the asymmetry and the overweight of the contradicting signal disappear in this case: the median weight is equal to 1 for the contradicting signal and a bit higher for the confirming signal (it should be observed, though, that the order for the first quartile is reversed). We can conclude that the asymmetric updating we observe in the SL treatment does not just come from the way subjects update on a signal after having observed a first piece of information.

²³We ran a Mann-Witney U test (Wilcoxon rank-sum test) on the median weight for the confirming and contradicting signal; we can reject the null hypothesis that their distribution is the same (p-value = 0.000003).

²⁴As we said, our results do not change if we define the signal as contradicting or confirming with respect to the action a_1 rather than with respect to the first belief a_2^1 , since the difference is for one observation only. Moreover, we cannot reject the hypothesis that the results, both for confirming and contradicting signals, are the same for the three treatments SL1, SL2 and SL3. (see the Appendix for details).

	1st Quartile	Median	3rd Quartile
α_2^2	0.00	1.02	2.38
α_2^2 (upon observing confirming signal)	0.25	1.06	2.41
α_2^2 (upon observing contradicting signal)	0.00	0.98	2.06

Table 6: Distribution of weights on the own signal in the SL treatment.

The table shows the quartiles of the distribution of the weight on the own signal for the second action at time 2 in the SL treatment. The data refer to all cases in which the first action at time 2 was equal to 50.

	1st Quartile	Median	3rd Quartile
α_2^2	0.00	0.00	1.22
α_2^2 (upon observing confirming signal)	0.00	1.00	1.84
α_2^2 (upon observing contradicting signal)	0.00	0.00	0.00

Table 7: Distribution of weights on the own signal in the IDM treatment.

The table shows the quartiles of the distribution of the weight on the own signal for the action at time 2 in the IDM treatment. The data refer to all cases in which the action at time 1 was equal to 50.

It is also interesting to see the difference in behavior when subjects have first stated a first belief of 50 (after observing an informative action or signal). In the SL experiment (Table 6), the median subject puts approximately the same weight on the signal, independently of whether it is confirming or contradicting. In the IDM treatment (Table 7), instead, he updates as a Bayesian agent would do (after receiving just one signal) if the signal is confirming and puts no weight at all on it if it is contradicting. The latter result has a simple interpretation. A subject choosing $a_1 = 50$ in the IDM treatment is not confident in one piece of information (e.g., ball color) only, he needs two to update. When the second ball color is in disagreement with the first, the subject states again a belief of 50, which is quite natural, since he has received contradictory information; when instead, the second ball has the same color, he updates as if it were the first signal he has received.

To understand the behavior in the SL treatment, we now look at how the weight on the signal changes with the first belief. Table 8 reports the quartiles for α_2^2 for three different classes of a_{2i}^1 . As one can immediately observe, the asymmetry occurs for the last two classes, but not for the first.²⁵

As we know from the previous analysis, the median subject chose an action $a_2^1 > 67$ mainly when he observed an action at time 1 greater than the theoretical Bayesian decision. These are cases in which the subject “trusted” the predecessor. These are also the cases in which subjects update in an asymmetric way. Table 9 reports the same analysis, but based on classes of predecessor’s

²⁵The 3rd quartile of 4.87 when $a_{2i}^1 > 83$ and the signal is confirming is of course influenced by subjects choosing 100 after having already chosen a number greater than 83.

	1st Quartile	Median	3rd Quartile
α_2^2 (upon observing confirming signal)			
Conditional on $50 \leq a_{2i}^1 \leq 66.7$	0.57	1.04	1.35
Conditional on $66.7 < a_{2i}^1 \leq 83.4$	0.18	0.91	1.57
Conditional on $a_{2i}^1 > 83$	0.43	2.10	4.87
α_2^2 (upon observing contradicting signal)			
Conditional on $50 \leq a_{2i}^1 \leq 66.7$	0.07	0.98	1.97
Conditional on $66.7 < a_{2i}^1 \leq 83.4$	1.02	1.68	2.11
Conditional on $a_{2i}^1 > 83.4$	2.53	3.34	4.26

Table 8: Distribution of weights for second actions at time 2 in the SL treatment.

The table shows the quartiles of the distribution of weights for second actions at time 2, conditional on different values of the first belief.

action, a_{1i} . Again, there is no asymmetry for the class $50 \leq a_{1i} \leq 66.7$, whereas there is for the extreme class. The middle class offers a less clear interpretation.

	1st Quartile	Median	3rd Quartile
α_2^2 (upon observing confirming signal)			
Conditional on $50 < a_{1i} \leq 66.7$	0.70	0.97	1.28
Conditional on $66.7 < a_{1i} \leq 83.4$	0.43	1.06	1.37
Conditional on $a_{1i} > 83.4$	0.50	1.01	2.36
α_2^2 (upon observing contradicting signal)			
Conditional on $50 < a_{1i} \leq 67$	0.96	1.06	2.72
Conditional on $66.7 < a_{1i} \leq 83.4$	0.49	1.20	2.11
Conditional on $a_{1i} > 83.4$	1.18	2.00	3.88

Table 9: Distribution of weights for second actions at time 2 in the SL treatment.

The table shows the quartiles of the distribution of weights for second actions at time 2, conditional on different values of the action at time 1.

In the next section we will offer an explanation for this phenomenon. We will show that introducing subjective beliefs (i.e., allowing for the possibility that a subject has incorrect beliefs) on the predecessor's rationality is not enough. We will need an extra ingredient.

Before we do so, let us make some observations.

First, our result cannot be explained in terms of risk preferences. As a matter of fact, risk aversion would push subjects receiving two contradicting pieces of information towards choosing 50, which makes our result even more striking. Moreover, the IDM treatment serves to control for risk preferences too, and we do see a striking difference of behavior between SL and IDM. Finally, a model in which subjects choose actions according to their risk preferences would not be able to predict asymmetric updating, unless risk preferences were correlated with the signal subjects receive, which is of course implausible.²⁶

²⁶The proof is simple and available upon request.

Second, if one thinks that the only inference subjects had to make from the predecessor’s action was the predecessor’s signal realization (and not the precision, since it was known), it is even more surprising that subjects simply did not choose 50 after a contradicting signal, since the fact that a good and a bad piece of information “cancel out” does not require sophisticated understanding of Bayes’s rule.

Third, and relatedly, one could observe that if a subject chose, e.g., $a_{2i}^1 = 84$ and then, after receiving a bad signal, $a_{2i}^2 = 50$, the corresponding α_{2i}^2 would be 2, which is compatible with the overweight we documented. It must be noticed, though, that if we exclude the cases in which $a_{2i}^2 = 50$, the asymmetry remains and is actually even stronger (see Table 11 in the Appendix). In other words, the asymmetry is not driven by subjects choosing $a_{2i}^2 = 50$.²⁷

Fourth, as we stressed in the Introduction, our result cannot be explained by and does not fall into categories of psychological biases sometimes invoked in decision making under uncertainty such as the base rate neglect or the confirmatory bias. Base rate neglect in our experiment would mean neglecting the first belief once the new piece of information (the private signal) is received. With such a bias, we should expect that the median choice of subjects first observing an action $a_1 > 50$ and then a signal $s_2 = 1$ should be equal to that at time 1 after observing a signal $s_1 = 1$, which is not the case (this would be equivalent to α_2^2 lower than or equal to 0, whereas it is slightly greater than 1). Moreover, such a bias should appear in the IDM treatment too, since it is not related to how the base rate is formed in the first place. As for the confirmatory bias, if subjects had the tendency to discard new information in disagreement with their original view, and only accept information confirming their original opinion (the definition of confirmatory bias) they should ignore (i.e., not update upon receiving) a contradicting signal, in sharp contrast with our results. Note that had we inverted the order in which information is presented (i.e., first the private signal and then the predecessor’s action) we would have not been able to rule out this possibility.²⁸

5 Explaining asymmetric updating

5.1 No asymmetry in Bayesian Updating

The asymmetric updating we observe in the laboratory is incompatible with Bayesianism. Whatever theory subject 2 has about subject 1’s behavior, once

²⁷In the econometric analysis that will follow, this type of concern is well taken into account, since we allow and elicit (from a control treatment) subjective beliefs.

²⁸Finally, it is worth mentioning that whereas in the social learning literature, as in much psychological literature, researchers have talked about “overconfidence,” in other experimental studies subjects show “underconfidence.” In particular, in experiments on decision making with naive advice, it has been observed that “when given a choice between getting advice or the information upon which the advice is based, subjects tend to opt for the advice, indicating a kind of underconfidence in their decision making abilities [...]” (Schotter, 2003). Our result is again not explained by this type of bias.

he has stated his first belief, he should simply put the same weight on the signal, independently of his realization.

One could be tempted to think that after observing a signal contradicting the predecessor's action, a subject could update down his belief on the rationality of the predecessor, revise the belief previously stated and, as a result, put more weight on his own private signal. This is, however, not in agreement with Bayesian updating. To see this, it suffices to notice that the posterior likelihood ratio on the value of the good is related to the prior likelihood ratio through this simple expression:

$$\frac{\Pr(V = 1|a_1, s_2)}{\Pr(V = 0|a_1, s_2)} = \frac{\Pr(s_2|V = 1, a_1) \Pr(V = 1|a_1)}{\Pr(s_2|V = 0, a_1) \Pr(V = 0|a_1)}. \quad (7)$$

Given the conditional independence of the signals, the expression simplifies to

$$\frac{\Pr(V = 1|a_1, s_2)}{\Pr(V = 0|a_1, s_2)} = \frac{\Pr(s_2|V = 1) \Pr(V = 1|a_1)}{\Pr(s_2|V = 0) \Pr(V = 0|a_1)}, \quad (8)$$

that is, to

$$\frac{\Pr(V = 1|a_1, s_2)}{\Pr(V = 0|a_1, s_2)} = \left(\frac{q_2^S}{1 - q_2^S} \right)^{2s_2 - 1} \frac{\Pr(V = 1|a_1)}{\Pr(V = 0|a_1)}. \quad (9)$$

where, using the same notation as in Section 2, q_2^S is the subjective precision attached to the signal by subject 2 (equivalent to a subjective weight α_2^2 , in the terminology of the previous section).

In the experiment, the subject states his belief $\Pr(V = 1|a_1)$ by making his first decision at time 2, a_2^1 . Therefore, we have,

$$\frac{\Pr(V = 1|a_1, s_2)}{\Pr(V = 0|a_1, s_2)} = \left(\frac{q_2^S}{1 - q_2^S} \right)^{2s_2 - 1} \frac{a_2^1/100}{1 - a_2^1/100}. \quad (10)$$

Whatever this first belief and whatever the model used to form it, if the subject were Bayesian, he should put the same weight on the signal, independently of its realization. Note that in this approach we have not imposed that the subject has correct expectations on the signal precision: indeed, we have allowed for subjective precisions. Nevertheless, for any precision the subject attaches to the signal, the weight must be the same for both realizations. The only requirement for this simple implication of Bayesian updating is that the signal realization (a draw from an urn) is independent of the rationality of the previous decision maker, which is logically undisputable.²⁹

5.2 Multiple priors and asymmetric updating

The intuition that observing a signal contradicting the first belief makes an agent update down on the predecessor's rationality and put more weight on his

²⁹Subjects know from the experimental design that signals are conditionally independent. The results of the IDM treatment are perfectly in line with subjects understanding it.

own signal, while in contradiction with Bayesianism, is, however, compatible with a model of updating in which an economic agent has multiple priors on the predecessor's rationality. In such a model, the own signal serves two purposes: it makes the agent select the prior on the predecessor's rationality; and, once this is done, to update on the first belief.

Specifically, suppose a subject at time 2 believes that the predecessor is of two types: either "rational" or "noise." A rational type always chooses an action greater than 50 after observing a good signal and an action lower than 50 after observing a bad signal. A noise type, instead, chooses any action between 0 and 100 independently of the signal. Let us denote these types by $T \in \{t_r, t_n\}$ and the probability that the subject is noise by $\Pr(T = t_n) \equiv \mu$. Whereas a Bayesian agent has a unique prior μ , a subject at time 2 has ambiguous belief on μ , that is, multiple priors belonging to the set $[\mu_*, \mu^*] \subseteq [0, 1]$.

To update his belief upon observing an event E , first of all the subject selects one of the priors in the set. If he is sufficiently confident that the event could occur conditional on the predecessor being rational, he will pick up the lowest prior μ_* , in the complementary case, he will pick up μ^* . In other words,

$$\begin{aligned} \text{if } \frac{\Pr(E|T = t_r)}{\Pr(E|T = t_n)} &\geq c, \text{ then } \mu = \mu_*, \text{ and} & (11) \\ \text{if } \frac{\Pr(E|T = t_r)}{\Pr(E|T = t_n)} &< c, \text{ then } \mu = \mu^*, \end{aligned}$$

where $c \in [0, \infty)$.

Note that in our experiment the subject makes this decision twice, first after observing the event $E \equiv \{a_1\}$ and then after observing the event $E \equiv \{a_1, s_2\}$.³⁰ Note also that after observing $\{a_1, s_2\}$ the subject, of course, also uses the signal realization s_2 to update on the first belief.

As we said in the Introduction, we refer to this model of updating based on the likelihood ratio $\frac{\Pr(E|T=t_r)}{\Pr(E|T=t_n)}$ as Likelihood Ratio Test Updating (LRTU) rule. It can be seen as a simple generalization of the Maximum Likelihood Updating (MLU) model (axiomatized by Gilboa and Schmeidler, 1993), in which the time 2 subject estimates μ to be the value in $[\mu_*, \mu^*]$ that maximizes the likelihood of observing the event E . Indeed, since

$$\Pr(E) = \Pr(E|T = t_r) \Pr(T = t_r) + \Pr(E|T = t_n) \Pr(T = t_n),$$

that is,

$$\Pr(E) = \Pr(E|T = t_r)(1 - \mu) + \Pr(E|T = t_n)\mu,$$

according to the MLU rule, the subject chooses either μ_* or μ^* , depending on whether the event is more likely conditional on the predecessor being rational or noise. That is,

³⁰Since a_1 is a continuous variable, $\Pr(\{a_1\}|T = t_r)$ should be read as a conditional density function.

$$\begin{aligned}
&\text{if } \frac{\Pr(E|T = t_r)}{\Pr(E|T = t_n)} \geq 1, \text{ then } \mu = \mu_*, \text{ and} & (12) \\
&\text{if } \frac{\Pr(E|T = t_r)}{\Pr(E|T = t_n)} < 1, \text{ then } \mu = \mu^*.
\end{aligned}$$

The LRTU model generalizes the MLU model to take into account that subjects may need stronger or weaker evidence in favor of one type in order to select a specific prior. This is equivalent to assuming that the subject acts as if he received another signal σ about the predecessor's type (and uncorrelated with the event). In this case, he would choose the prior to maximize the following probability:

$$\Pr(E, \sigma) = \Pr(E, \sigma|T = t_r) \Pr(T = t_r) + \Pr(E, \sigma|T = t_n) \Pr(T = t_n).$$

That is, he would select $\mu = \mu_*$ (or $\mu = \mu^*$) if the following inequality is (or is not) satisfied:

$$\frac{\Pr(E, \sigma|T = t_r)}{\Pr(E, \sigma|T = t_n)} \geq 1,$$

that is,

$$\frac{\Pr(E|T = t_r)}{\Pr(E|T = t_n)} \frac{\Pr(\sigma|T = t_r)}{\Pr(\sigma|T = t_n)} \geq 1,$$

or

$$\frac{\Pr(E|T = t_r)}{\Pr(E|T = t_n)} \geq \frac{\Pr(\sigma|T = t_n)}{\Pr(\sigma|T = t_r)}. \quad (14)$$

By setting $\frac{\Pr(\sigma|T=t_n)}{\Pr(\sigma|T=t_r)} \equiv c$, one obtains the LRTU model.

As we explained in the Introduction, updating by first selecting one prior and then applying Bayes's rule is one way in which the decision theory literature has solved the problem of updating beliefs when there are multiple priors. A second paradigm, referred to as Full Bayesian Updating (FBU) consists in updating all priors, by using Bayes's rule for each of them. The choice then depends on the agent's preferences. We will consider the most common case, axiomatized by Pires (2002) in which the agent has maxmin preferences.

Before showing these different models and their structural estimation in detail (in the next Section) we first illustrate them through a simple example. The example will give the main intuition as to why the LRTU model can generate the type of asymmetric updating we observe in our data, whereas the FBU model cannot.

5.3 An example

Suppose that subject 2 has multiple priors $[\mu_*, \mu^*] = [0, 1]$ on the predecessor's type. Suppose that he observes $a_1 = 70$ and then the signal $s_2 = 0$. Let us consider first the LRTU model and suppose the threshold is $c = 1$, so that the model is equivalent to the MLU model.

Suppose that subject 2 has expectations on the rational and noise types' actions at time 1 such that $\frac{\Pr(a_1=70|T=t_r)}{\Pr(a_1=70|T=t_n)} \geq 1$. In this case, the subject selects the prior $\mu_* = 0$. The subject is confident on the predecessor's rationality, and, therefore, chooses $a_2^1 = 70$. After receiving the signal $s_2 = 0$, the subject now reassesses the predecessor's rationality. The probability of observing an action greater than 50 and a negative signal conditional on the predecessor being rational is now lower. If, in particular, $\frac{\Pr(a_1=70, s_2=0|T=t_r)}{\Pr(a_1=70, s_2=0, |T=t_n)} < 1$, then the subject chooses $\mu^* = 1$. Being now confident that the predecessor was a noise type, the subject considers $a_1 = 70$ completely uninformative, which would imply a belief of 0.5 on $V = 100$. On top of this, the subject has observed a bad signal: by applying Bayes's rule to a prior of 0.5, the subject obtains a posterior belief of 0.3 on the value being 100 and, as a result, chooses $a_2^2 = 30$. In terms of our previous analysis, this is equivalent to a subject overweighting the signal, with $\alpha_2^2 = 2$, since $30 = 100 \frac{(1-q)^2 \frac{70}{100}}{(1-q)^2 \frac{70}{100} + q^2 (1 - \frac{70}{100})}$. A similar analysis applies to the case in which the subject observes a signal $s_2 = 1$. It is easy to see that if $\frac{\Pr(a_1=70|T=t_r)}{\Pr(a_1=70|T=t_n)} \geq 1$, then a fortiori $\frac{\Pr(a_1=70, s_2=1|T=t_r)}{\Pr(a_1=70, s_2=1, |T=t_n)} \geq 1$. Therefore, in this case the subject sticks to the prior $\mu_* = 0$. Since the subject is still confident that the predecessor was rational, he does not change his first belief on $V = 100$, which remains 0.7. Since the subject has observed a good signal, by applying Bayes's rule to a belief of 0.7, he obtains a posterior belief of 0.84 on the value being 100 and, as a result, chooses $a_2^2 = 84$. This is equivalent to a subject weighting the signal as a Bayesian agent would do, with $\alpha_2^2 = 1$. This way of updating, thus, generates the asymmetry we observe in our data.

Let us consider now the FBU model, in which the subject, with maxmin preferences, updates all priors. After observing $a_1 = 70$, the subject updates his belief on the value of the good using each prior $\mu \in [0, 1]$. This means that his posterior beliefs on $V = 100$ lie in $[0.5, 0.7]$. Therefore, he chooses $a_2^1 = 50$, the action that maximizes the minimum payoff he can obtain. After receiving the signal $s_2 = 0$, the subject updates his set of beliefs to $[0.3, 0.5]$. Of course, this implies that again he chooses $a_2^2 = 50$, which is equivalent to $\alpha_2^2 = 0$. After receiving the signal $s_2 = 1$, instead, the subject updates his set of beliefs to $[0.7, 0.84]$. He will then maximize his utility by choosing $a_2^2 = 70$, which is equivalent to $\alpha_2^2 = 1$. This updating rule, therefore, would imply no updating at all (rather than overweighting the signal) after receiving a contradicting signal, and updating as a Bayesian after observing a confirming signal (an asymmetric way of updating that sharply differs from that we observe).

6 Econometric analysis

So far we have illustrated the three models of updating in a heuristic way. In this section, we perform a formal statistical comparison to quantify the evidence in favor of the LRTU model against the other two models of updating. For each updating rule (BU, LRTU, FBU), in our econometrics models, we explicitly consider the individual heterogeneity observed in the data. As we have seen

in the previous sections, the reported beliefs both at time 1 and at time 2 are quite heterogeneous, with non-regular features (e.g., multi-modal, asymmetric distributions, mixtures of probability masses and continuous distributions). To take this into account, we use the IDM treatment observations to obtain a nonparametric estimator for the distribution of the unobservable heterogeneity, and develop a model comparison procedure that does not rely on parametric specifications. Our purpose is to understand which model explains the behavior of subjects at time 2 best. The three models will have two common ingredients:

- i*) subjective beliefs on the informativeness (precision) of the private signal;
- ii*) subjective beliefs on the rationality of the subject acting at time 1.

The models will instead differ in the way a subject at time 2 updates his beliefs (and in the way he behaves as a function of the beliefs).

Let us start discussing point *i* above. We know that there is heterogeneity in how subjects update their beliefs on the basis of their private signal. To take this into account, in our analysis we let the subjective precisions $q_{1i}^S = \Pr(s_{1i} = 1|V = 1) = \Pr(s_{1i} = 0|V = 0)$ and $q_{2i}^S = \Pr(s_{2i} = 1|V = 1) = \Pr(s_{2i} = 0|V = 0)$ vary for each observation i (recall that the superscript S stands for subjective). Recall that in both the SL and the IDM treatments, we observe the distribution of stated beliefs at time 1, which are based on the observation of one signal only. Furthermore, in the IDM treatment, in 50% of the rounds, we observe the joint distribution of stated beliefs at times 1 and 2. From these stated beliefs, we can recover q_{1i}^S , and q_{2i}^S , since there is a one-to-one map between beliefs and precisions (e.g., $a_{1i} = 73$ after observing $s_{1i} = 1$ is equivalent to $q_{1i}^S = 0.73$; in the IDM treatment, $a_{2i} = 80$ after $a_{1i} = 73$ and $s_{2i} = 1$ is equivalent to $q_{2i}^S = 0.60$). We will use the empirical distribution of q_{1i}^S so recovered, as representing the distribution of the subjective precision of a signal at time 1. When, for estimation, we will need the joint distribution of precisions, we will use the empirical distribution obtained by considering the sample of observations i 's for which both (q_{1i}^S, q_{2i}^S) can be recovered in the IDM treatment.³¹

Let us move to point *ii*. In line with the above discussion, we assume that a subject at time 2 believes that the predecessor is of two types: either “rational” (t_r) or “noise” (t_n), with $\Pr(t_n) \equiv \mu$. A rational type is defined as someone who always chooses an action strictly greater than 50 after observing a good signal and an action lower than 50 after observing a bad signal. A noise type, instead, chooses any action between 0 and 100 independently of the signal.³² The BU model assumes a unique μ ; in the LRTU and FBU model, instead, the ambiguous beliefs on the predecessor’s rationality consists in a set of priors $[\mu_*, \mu^*]$. We will estimate the unique μ or the lower and upper bounds μ_* and μ^* by fitting the models to the data.

As we know from Section 4, the empirical distribution of actions at time 1

³¹In our estimations, we assume that the distribution of subjective signal precisions be independent of the signal realization. In another specification, we also considered the distribution conditional on the realization: the results do not change.

³²As we explained in Section 2, we use this definition of rationality since the only thing that subject 2 has to learn from subject 1 is, indeed, the signal realization, and this is revealed under the minimal requirement that the subject updates in the right direction.

conditional on a good signal is almost the mirror image (with respect to 50) of the distribution conditional on a bad signal. For this reason, we now pool all the observations by transforming a_{1i} into $100 - a_{1i}$ whenever $s_{1i} = 0$. We can then focus our analysis on actions strictly greater than 50. In particular, given this transformation, a rational subject always chooses an action greater than 50.

In the spirit of the descriptive analysis, we divide the interval $(50, 100]$ into three “bins” $B_1 = (50, 66.7]$, $B_2 = (66.7, 83.4]$ and $B_3 = (83.4, 100]$. As highlighted by the previous analysis, subjects react differently to a predecessor’s choice of an action below the Bayesian one, in the neighborhood of the Bayesian one, or more extreme than it. We want to understand this behavior more in depth in our econometric analysis. Of course by pooling the data together for these intervals of actions, we also have enough data to estimate our models.

For the noise type, we assume that (subject 2 believes that) his actions follow a distribution $g(a_1)$ symmetric around 50. We construct a histogram density in the following way. Let $\Phi_\sigma(B)$ be the probability assigned to an interval B by a normal distribution with mean 50 and variance σ^2 . Then,

$$g_\sigma(a_1) = \frac{1}{\Phi_\sigma([0, 100])} \sum_{l=1}^3 \frac{\Phi_\sigma(B_l)}{|B_l|} \cdot 1\{a_1 \in B_l\}, \text{ for } a_1 > 50, \quad (15)$$

where $|B_l|$ denotes the width of B_l . In words, we construct the histogram by considering a truncated normal distribution, and computing the resulting density for the three chosen bins.

To estimate the parameter σ we use the cases in which subjects at time 1 updated their beliefs in the wrong direction. Indeed we estimate it by the empirical standard deviation $\hat{\sigma} = \sqrt{\frac{1}{\#\{i: a_{1i} \in \Theta\}} \sum_{i \in \Theta} (a_{1i} - 50)^2}$, where Θ is the set of actions $a_{1i} < 50$ (> 50) taken after the observation of a good (bad) signal.³³ We obtain the estimate $\hat{\sigma} = 0.273$ (with a standard error —computed by delta method— of 0.006). Given this estimated value of σ , we re-denote the distribution $g_\sigma(a_1)$ by $g(a_1)$. Note that, since $g(a_1)$ is symmetric, the probability of observing a mistake (i.e., updating in the wrong direction) from the point of view of subject 2 is given by $\Pr(a_1 > 50 | s_1 = 0) = \Pr(a_1 < 50 | s_1 = 1) = \frac{\mu}{2}$.

As for the rational type, we assume that subjects at time 2 have correct expectations on the distribution of actions at time 1 by rational subjects. Consider the empirical distribution of time 1 subject’s actions. The histogram density for the actions greater than 50 is

$$h(a_1) = \sum_{l=1}^3 \hat{b}_l 1\{a_1 \in B_l\} \quad \text{for } a_1 > 50, \quad (16)$$

where $\hat{b}_l = \frac{1}{|B_l|} \frac{\sum_i 1\{a_{1i} \in B_l\}}{\sum_i 1\{a_{1i} > 50\}}$. This means that $\hat{b}_1, \hat{b}_2, \hat{b}_3$ are the histogram density estimates for the three intervals we are considering.³⁴ Note, however, than not

³³Of course, given the above transformation of data, all incorrect actions are below 50.

³⁴Note that, of course, we exclude $a_{1i} = 50$. This action is uninformative and, therefore, has a different status from any other action.




Figure 5: Histograms $f(a_1)$ (solid line) and $g(a_1)$ (dotted line) for rational and noise actions at time 1.

all observed actions greater than 50 can be considered as coming from rational subjects, since noise type subjects choose correct decisions half of the time. To correct for the proportion of irrational actions, we consider the distribution of rational actions to be³⁵

$$f(a_1) = \frac{h(a_1) - (0.07)g(a_1)}{0.93}.$$

Figure 5 shows the estimated histograms.

Given these histograms, a (rational) subject i at time 2, observing an action $a_{1i} > 50$, has the following conditional beliefs (density functions):

$$\begin{aligned} p(a_{1i}|V = 1, t_r) &= p(a_{1i}|s_{1i} = 1, V = 1, t_r)q_{1i}^S + p(a_{1i}|s_{1i} = 0, V = 1, t_r)(1 - q_{1i}^S) = q_{1i}^S f(a_{1i}), \\ p(a_{1i}|V = 0, t_r) &= (1 - q_{1i}^S)f(a_{1i}), \\ p(a_{1i}|V = 1, t_n) &= p(a_{1i}|V = 0, t_n) = g(a_{1i}). \end{aligned} \quad (17a)$$

While subjects are constrained to have correct expectations on the distribution of rational actions (and on the standard deviation of the noise actions), they have subjective beliefs on the precisions of signals as well as on the proportion of the noise type (μ) and of the rational type ($1 - \mu$).

Given these common ingredients, we can now describe how a subject forms his beliefs on the value of the good depending on the updating model.

³⁵Recall that we observed 3.5% of incorrect updating at time 1. Given the symmetry of $g(a_1)$, they must result from a 7% of noise type's actions.

The BU model

According to the BU model, given a prior belief μ on the proportion of noise type subjects at time 1, a subject applies Bayes's rule to determine his first action,

$$\begin{aligned} a_{2i,B}^1(\mu, q_{1i}^S) &\equiv 100 \Pr(V = 1|a_{1i}) = 100 \frac{(1-\mu)q_{1i}^S f(a_{1i}) + g(a_{1i})\mu}{(1-\mu)f(a_{1i}) + 2g(a_{1i})\mu} \quad (18) \\ &= 100 \frac{(1-\mu)q_{1i}^S \frac{f(a_{1i})}{g(a_{1i})} + \mu}{(1-\mu)\frac{f(a_{1i})}{g(a_{1i})} + 2\mu}. \end{aligned}$$

To simplify notation, let us denote the log-likelihood ratio by $l(\cdot)$, that is, $l(x) =: \ln \frac{x}{1-x}$. Then, after receiving a confirming signal ($s_{2i} = 1$), a subject chooses an action $a_{2i,B}^2$ such that the following equality holds:

$$l\left(\frac{a_{2i,B}^2(\mu, q_{1i}^S, q_{2i}^S)}{100}\right) = l\left(\frac{a_{2i,B}^1(\mu, q_{1i}^S)}{100}\right) + l(q_{2i}^S); \quad (19)$$

similarly, after a contradicting signal, action $a_{2i,B}^2$ will satisfy

$$l\left(\frac{a_{2i,B}^2(\mu, q_{1i}^S, q_{2i}^S)}{100}\right) = l\left(\frac{a_{2i,B}^1(\mu, q_{1i}^S)}{100}\right) + l(1 - q_{2i}^S). \quad (20)$$

Note that $a_{2i,B}^2$ is fully determined by $a_{2i,B}^1$ and q_{2i}^S given that the dependence on μ is summarized in $a_{2i,B}^1(\mu, q_{1i}^S)$.

The LRTU model

In this model, subject 2 starts with a set of priors $[\mu_*, \mu^*]$ on the proportion of noise type subjects. He selects one prior in $[\mu_*, \mu^*]$ on the basis of the likelihood ratio

$$\frac{p(a_{1i}|T = t_r)}{p(a_{1i}|T = t_n)} = \frac{\frac{1}{2}q_{1i}^S f(a_{1i}) + \frac{1}{2}(1 - q_{1i}^S)f(a_{1i})}{g(a_{1i})} = \frac{f(a_{1i})}{2g(a_{1i})}. \quad (21)$$

In particular, he selects μ_{2i}^1 as follows:

$$\mu_{2i}^1 = \begin{cases} \mu_* & \text{if } \frac{f(a_{1i})}{g(a_{1i})} \geq 2c, \\ \mu^* & \text{if } \frac{f(a_{1i})}{g(a_{1i})} < 2c. \end{cases} \quad (22)$$

He then applies Bayes's rule to determine his first action, $a_{2i,L}^1(\mu_{2i}^1, q_{1i}^S)$, which is identical to expression (18), after substituting μ_{2i}^1 to μ . Note that $a_{2i,L}^1(\mu_{2i}^1, q_{1i}^S)$ varies from $100q_{1i}^S$ to 50 as μ_{2i}^1 varies from 0 to 1. Moreover, note that although the same q_{1i}^S was used both in (21) and in (18), (21) does not depend on q_{1i}^S .

Now, consider the second action at time 2 and suppose the subject receives a confirming signal ($s_{2i} = 1$). Then,

$$\begin{aligned} p(a_{1i}, s_{2i} = 1|t_r) &= \frac{1}{2} [q_{1i}^S q_{2i}^S + (1 - q_{1i}^S)(1 - q_{2i}^S)] f(a_{1i}), \\ p(a_{1i}, s_{2i} = 1|t_{ir}) &= \frac{1}{2} g(a_{1i}). \end{aligned}$$

Therefore,

$$\mu_{2i,confirm}^2 = \begin{cases} \mu_* & \text{if } \frac{f(a_{1i})}{g(a_{1i})} \geq \frac{c}{q_{1i}^S q_{2i}^S + (1 - q_{1i}^S)(1 - q_{2i}^S)}, \\ \mu^* & \text{if } \frac{f(a_{1i})}{g(a_{1i})} < \frac{c}{q_{1i}^S q_{2i}^S + (1 - q_{1i}^S)(1 - q_{2i}^S)}. \end{cases} \quad (23)$$

Given $\mu_{2i,confirm}^2$ and q_{2i}^S , $a_{2i,L}^2$ satisfies

$$l\left(\frac{a_{2i,L}^2(\mu_{2i,confirm}^2, q_{1i}^S, q_{2i}^S)}{100}\right) \equiv l\left(\frac{a_{2i,L}^1(\mu_{2i,confirm}^2, q_{1i}^S)}{100}\right) + l(q_{2i}^S), \quad (24)$$

where $a_{2i,L}^1(\mu_{2i,confirm}^2, q_{1i}^S)$ is equal to (18) with the exception that μ_{2i}^1 is replaced by $\mu_{2i,confirm}^2$.

Note that the threshold in (23) is lower than that in (22).

For the contradicting signal case, the analysis is analogous; we have

$$\begin{aligned} \Pr(a_{1i}, s_{2i} = 0|t_r) &= \frac{1}{2} [q_{1i}^S(1 - q_{2i}^S) + (1 - q_{1i}^S)q_{2i}^S] f(a_{1i}), \\ \Pr(a_{1i}, s_{2i} = 0|t_{ir}) &= \frac{1}{2} g(a_{1i}), \end{aligned}$$

and, therefore,

$$\mu_{2i,contradict}^2 = \begin{cases} \mu_* & \text{if } \frac{f(a_{1i})}{g(a_{1i})} \geq \frac{c}{q_{1i}^S(1 - q_{2i}^S) + (1 - q_{1i}^S)q_{2i}^S}, \\ \mu^* & \text{if } \frac{f(a_{1i})}{g(a_{1i})} < \frac{c}{q_{1i}^S(1 - q_{2i}^S) + (1 - q_{1i}^S)q_{2i}^S}. \end{cases} \quad (25)$$

Given $\mu_{2i,contradict}^2$ and q_{2i}^S , $a_{2i,L}^1$ satisfies

$$l\left(\frac{a_{2i,L}^1(\mu_{2i,contradict}^2, q_{1i}^S, q_{2i}^S)}{100}\right) \equiv l\left(\frac{a_{2i,L}^1(\mu_{2i,contradict}^2, q_{1i}^S)}{100}\right) + l(1 - q_{2i}^S). \quad (26)$$

Note that the threshold in (25) is higher than that in (22): a confirming signal lowers the threshold to trust the predecessor's rationality, whereas a contradicting signal raises it.

The FBU model

In this model too a subject at time 2 starts with a set of priors $[\mu_*, \mu^*]$ on the proportion of noise type subjects at time 1. The subject applies Bayes's rule for each prior μ_{2i}^1 in $[\mu_*, \mu^*]$ and obtains a belief

$$p_{2i}^1(\mu_{2i}^1, q_{1i}^S) \equiv \Pr(V = 1 | a_{1i}; \mu_{2i}^1, q_{1i}^S) = \frac{(1 - \mu_{2i}^1)q_{1i}^S f(a_{1i}) + g(a_{1i})\mu_{2i}^1}{(1 - \mu_{2i}^1)f(a_{1i}) + 2g(a_{1i})\mu_{2i}^1}. \quad (27)$$

As a result, he has a range of beliefs on the value of the good being 100: $[p_{2i}^1(\mu^*, q_{1i}^S), p_{2i}^1(\mu_*, q_{1i}^S)]$.

After receiving a confirming signal case, the subject updates his range of beliefs so that

$$[l(p_{2i}^2(\mu^*, q_{1i}^S, q_{2i}^S)), l(p_{2i}^2(\mu_*, q_{1i}^S, q_{2i}^S))] = [l(p_{2i}^1(\mu^*, q_{1i}^S)), l(p_{2i}^1(\mu_*, q_{1i}^S))] + l(q_{2i}^S), \quad (28)$$

where $b + [c, d]$ means $[b + c, b + d]$. Similarly, in the contradicting signal case,

$$[l(p_{2i}^2(\mu^*, q_{1i}^S, q_{2i}^S)), l(p_{2i}^2(\mu_*, q_{1i}^S, q_{2i}^S))] = [l(p_{2i}^1(\mu^*, q_{1i}^S)), l(p_{2i}^1(\mu_*, q_{1i}^S))] + l(1 - q_{2i}^S). \quad (29)$$

Recall that a maxmin expected utility agent with a set of beliefs $[\underline{p}_i, \bar{p}_i]$ chooses the optimal action $a_{i, \max \min}$ such that

$$a_{i, \max \min} = \arg \max_a \min_{p \in [\underline{p}_i, \bar{p}_i]} E_p(100 - 0.01(V - a)^2),$$

that is,

$$a_{i, \max \min} = \begin{cases} 100\underline{p}_i, & \text{if } \underline{p}_i > \frac{1}{2}, \\ 50, & \text{if } \underline{p}_i \leq \frac{1}{2} \text{ and } \bar{p}_i \geq \frac{1}{2}, \\ 100\bar{p}_i, & \text{if } \bar{p}_i < \frac{1}{2}. \end{cases}$$

Therefore, in the FBU model, since $p_{2i}^1(\mu^*, q_{1i}^S) \geq \frac{1}{2}$, the subject's first action is based on the most pessimistic prior, $\mu = \mu^*$:

$$a_{2i, F}^1 = a_{2i}^1(\mu^*, q_{1i}^S) = 100p_{2i}^1(\mu^*, q_{1i}^S).$$

Similarly, the second action is

$$a_{2i, F}^2 = \begin{cases} 100p_{2i}^2(\mu^*, q_{1i}^S, q_{2i}^S), & \text{if } p_{2i}^2(\mu^*, q_{1i}^S, q_{2i}^S) > \frac{1}{2}, \\ 50, & \text{if } p_{2i}^2(\mu^*, q_{1i}^S) < \frac{1}{2}, \text{ and } p_{2i}^2(\mu_*, q_{1i}^S) > \frac{1}{2}, \\ 100p_{2i}^2(\mu_*, q_{1i}^S, q_{2i}^S), & \text{if } p_{2i}^2(\mu_*, q_{1i}^S, q_{2i}^S) < \frac{1}{2}. \end{cases}$$

6.1 Estimation methodology and results

We estimate the three models by the Generalized Method of Moments (GMM). In each of our models, the heterogeneity in the subjective precision of signals induces a distribution of actions at time 2 or any fixed value of the parameters. The estimation strategy consists in finding the parameter values such that the distribution of actions predicted by a model is closest to the actual distribution.

With maximum likelihood, we would need to specify a parametric distribution for (q_{1i}^S, q_{2i}^S) . In our experiment, however, we do observe the empirical distribution of (q_{1i}^S, q_{2i}^S) . With GMM, we can use it without parametric assumptions. We have a gain in terms of robustness of the estimates, with a potential sacrifice in terms of efficiency.

Specifically, in the descriptive analysis, we have reported the three quartiles of the empirical distribution of the weights α 's for a) the first action at time 2; b) the second action at time 2, conditional upon receiving a confirming signal; c) the second action at time 2, conditional upon receiving a contradicting signal. For each model, we now match the value of the cumulative distribution functions of α 's at each of these quartiles, for all these three cases (for a total of nine moment conditions). We do so separately for each of the three intervals in which we have divided $(50, 100]$. In other words, we estimate the parameters that make a model generate data whose distribution is as close as possible to the true dataset's in terms of the three observed quartiles, conditional on a subject at time 2 having observed a_{1i} belonging to either $B_1 = (50, 66.7]$, or $B_2 = (66.7, 83.4]$ or $B_3 = (83.4, 100]$. The estimate will, therefore, result from 27 moment conditions (nine for each type of action).³⁶

Since our models predict the behavior of a rational type, we restrict our analysis to the dataset consisting of rational actions only. In other words, we eliminate the (few) cases in which a subject updated in the "wrong direction" after receiving a piece of information (e.g., updating down after receiving a good signal). Consistently, we also restrict the sample of q_{1i} and q_{2i} to those that are weakly greater than 0.5.

We refer the readers to the Appendix for a detailed illustration of the estimation procedure. Here we simply observe that for the BU model we must estimate one parameter, that is, the proportion of noise type subjects, μ . For the LRTU model, we must estimate three parameters: the bounds of the support for the prior on the proportion of noise type subjects, μ_* and μ^* , as well as the threshold c . Finally, for the FBU model, we must only estimate μ_* and μ^* .

Table 10 reports the results of the second stage GMM estimation (non-parametric bootstrapped standard errors in parenthesis).

Model	μ	μ_*	μ^*	c
BU	0.30 (0.053)			
LRTU		0 (0.019)	0.30 (0.045)	[1.65, 1.73] (0.073)
FBU		0.30 (0.070)	0.30 (0.069)	

Table 10: Parameter Estimates

The table shows the parameter estimates of the three models. The standard errors in parenthesis are computed by non-parametric bootstrap with 1000 bootstrap samples. The standard error for c refers to 1.65.

³⁶For the BU model, as observed above, given $a_{2i,B}^1(\mu, q_{1i}^S)$, action $a_{2i,B}^2(\mu, q_{1i}^S, q_{2i}^S)$ only depends on q_{2i} . For this reason, the estimate of μ is only based on the first action at time 2 (i.e., on 9 moment conditions).

The estimated proportion of noise type subjects in the BU model is $\mu = 0.3$. This of course reflects the tendency of subjects at time 2 to “discount” the actions a_{1i} , in particular those in bins B_1 and B_2 , when choosing a_{2i}^1 , as documented in Section 4. Given the densities $f(a_1)$ and $g(a_1)$ clearly they did not discount more extreme actions too much.

Of course, $\mu = 0.3$ implies a belief that in 15% of the cases a subject at time 1 updated in the wrong direction, which is higher than the actual (3.5%) proportion of mistakes we observed at time 1, thus showing that subjects at time 2 did not have rational expectations on the proportion of noise and rational predecessors.

Let us now move to the FBU model. Such a model can in principle explain the observed behavior better, given that there is an extra degree of freedom. It turns out, however, that the FBU model’s estimates coincide with the BU’s, since the support for the multiple priors is estimated to be just the point 0.3. In other words, adding multiple priors in this case does not provide a different and better fit of the data, compared to the BU model.

Let us now look at the LRTU model. First of all note that the GMM objective function does not have a unique minimizer for the parameter c : $c \in [1.65, 1.73]$. Nevertheless, the other parameters have the same estimate for any $c \in [1.65, 1.73]$. This parameter c co-determines the thresholds to trust or not the predecessor. It is clear that the inequalities in (22), (24), (26) may be satisfied for a set of parameter values. The estimates shows that to “trust” a predecessor’s action, a subject needs the likelihood ratios to be greater than a threshold equal to 1.65, that is, he requires stronger evidence of rationality than what assumed in the MLU model (in which $c = 1$). When this threshold is reached, the subject considers the observed action as fully rational (since the estimated lower bound for proportion of a noise type is $\mu_* = 0$). When, instead, the threshold is not reached, he updates as if the probability of a noise predecessor were $\mu^* = 0.3$. Note that this is actually the estimate for the single prior in the BU model. Essentially, according to our estimates, when the subject observes an action that he trusts, he fully does so; when, he does not trust it, he attaches a probability of 0.30 to it coming from a noise type. It is interesting to see the implications of these parameter estimates for subjects’s behavior. Let us consider first a_{2i}^1 . Given the parameter estimates, when choosing a_{2i}^1 , subjects do not trust an action $a_{1i} \in (50, 66.7]$ or $a_{1i} \in (66.7, 83.4]$ (that is, they pick the prior $\mu^* = 0.3$); they do trust an action $a_{1i} \in (83.4, 100]$. Let us consider now a_{2i}^2 . The decision to trust or not the predecessor depends on the subjective precisions of signals, in this case, as one can notice from (23) and (25). After receiving a confirming signal, they keep not trusting an action $a_{1i} \in (50, 66.7]$, whereas in 72.7% of the cases they become trusting of an action $a_{1i} \in (66.7, 83.4]$.³⁷ Of course they keep trusting $a_{1i} \in (83.4, 100]$. After receiving a contradicting signal, they keep not trusting an action $a_{1i} \in (50, 66.7]$ or $a_{1i} \in (66.7, 83.4]$, of course, and in 68.9% of the cases they stop trusting an

³⁷This is in fact a feature we did not observe in our descriptive analysis, an instance in which this model does not fit the data well. Despite this, the model is the best predictor of the distribution of actual actions, as we will show.

action $a_{1i} \in (83.4, 100]$.

The final question is whether the LRTU model provides a better explanation for the observed behavior than the BU model (and the FBU model, since they happen to coincide). A simple comparison of the minimized GMM objective functions for the two models would not be an appropriate way of measuring their relative fitness, since one model allows for more degree of freedom (has more parameters) than the other. There is a large literature on model specification test that accounts for over-fitting of the models with extra parameters within the framework of GMM (see Newey and McFadden, 1994). No existing test, however, can be readily applied to our case, due to the non-standard features of our moment conditions. In particular, note that (i) the GMM objective function for the LRTU model is discontinuous and non-differentiable; (ii) one parameter of the LRTU model can only be set identified; and (iii) the LRTU nests the BU model at the boundary of the parameter space (e.g., $\mu_* = \mu^*$). Instead of developing a new asymptotically valid model selection test that can overcome all these issues, we consider a model comparison test based on the idea of resampling p -value, which heuristically quantifies the strength of evidence against a null model without relying on an asymptotic theory (at the cost of being computationally intensive). We refer the reader to the Appendix for the details. Here we note that in the model comparison test, we set up the null hypothesis “the BU model with parameter value $\mu = 0.3$ is the true data generating process.” We simulate 1000 datasets from the BU model with $\mu = 0.3$, of course resampling (q_{1i}^S, q_{2i}^S) from the empirical distribution, as discussed above. For each of these data sets, we then estimate the BU and LRTU models by GMM and let \hat{L}_{BU}^j and \hat{L}_{LRTU}^j be the resulting minimized values of the GMM objective function for sample $j = 1, 2, \dots, 1000$. Note that $\Delta \hat{L}^j = \hat{L}_{BU}^j - \hat{L}_{LRTU}^j$ is non-negative since the LRTU model nests the BU model, and hence represents a gain in model fitness solely due to “over-parametrization” of the LRTU model relative to the BU model. We take the empirical distribution of $\Delta \hat{L}^j$ ($j = 1, \dots, 1000$) as the null distribution of the model fitness criterion. We compute $\Delta \hat{L} = \hat{L}_{BU} - \hat{L}_{LRTU}$ as the difference between the minimized GMM objective functions of the BU and LRTU models for our dataset. To measure how unlikely $\Delta \hat{L}$ is in terms of the null distribution, we compute the p -value by

$$\frac{1}{1000} \sum_{j=1}^{1000} 1 \left\{ \Delta \hat{L}^j \geq \Delta \hat{L} \right\},$$

where $1\{\}$ is the indicator function. The p -value, so computed, is 0.008, that is, we can reject the null hypothesis and consider our evidence in support of the LRTU model. The LRTU model fits the data significantly better than the BU model after we have properly taken into account the gain of over-parametrization. Moreover, in our approach we did not impose any parametric restriction on the heterogeneity of subjective precisions: the evidence in favor of the LRTU model is robust to individual heterogeneity (i.e., it does not depend on a parametric assumption on heterogeneity).

7 Discussion

We now want to discuss some features of our LRTU and FBU models, and consider some alternative approaches, to highlight how our experimental work could inform future theoretical developments.

A crucial aspect of our LRTU model is that we let the subject pick a different prior from the same set of priors every time he receives new information. This is in line with the tradition of the statistics literature, and dates back to the Type-II maximum likelihood of Good (1965), in which new observations are used to estimate a prior for an unknown parameter (see, e.g., Berger 1985). In this methodology, the set of priors (from which one prior is estimated) is invariant to the new arrival of information. This approach is, however, less well established in the decision theory literature. In their axiomatization of the MLU model, Gilboa and Schmeidler (1993) do not consider a multi-period problem. In their MLU framework an agent only updates once, therefore the problem of how to update once new information arrives is not immediately relevant. Nevertheless, in their analysis, implicitly the choice of the prior is once and for all. This would be equivalent, in our experiment, to the subject having to stick to the prior he has selected after observing the predecessor's action only. Pires (2002) observes that in the spirit of ambiguity aversion it is sensible to assume that the agent keeps all possible priors alive and for this reason she advocates the FBU model. Gilboa and Marinacci (2013) describe the MLU and FBU models as two extremes: one in which only one prior is used and one in which all are. We view our model as somehow in between these two extremes. In the LRTU model, the subject does pick one prior, but this does not eliminate ambiguity for ever, since the subject can pick another prior after new information arrives.³⁸ Of course a model in which the agent picks different priors every time new information arrives exhibits a form of time inconsistency.³⁹ In such a model preferences are not stable, which may be problematic from a normative view point (similar objections apply to Epstein and Schneider, 2007). Nevertheless, from a descriptive viewpoint, the model that best fits the data lets the subjects choose the prior every time (from a set that we estimate).

Our way of modelling updating multiple priors is closely related to the model proposed by Ortoleva (2012). In his model, the decision maker has a prior over possible priors (referred to as theories to avoid confusion). Initially, the theory with the highest prior probability is selected. When a new event occurs, if the likelihood of that event is higher than a specific threshold, then the theory is maintained; otherwise, the prior over priors is updated on the basis of the likelihood of the event, and the theory with the highest posterior probability is selected. In our approach, we have not considered a threshold, which is equivalent to set it equal to 1 (so that there is always a reconsideration of all

³⁸Epstein and Schneider (2007) consider an intertemporal economy. They do not impose that once a prior is chosen, it is chosen forever, letting the agent re-choose the prior in a neighborhood of the prior previously chosen.

³⁹For a theoretical investigation of dynamically consistent updating of ambiguous beliefs see Hanany *et al.* (2007).

theories). Moreover, our parameter c can be viewed as playing a similar role as the prior over priors in Ortoleva’s model (since in both set ups, an agent can be a priori biased in favour of a particular theory). We have not attempted to estimate Ortoleva’s model, and, in particular, we have not considered a threshold below which there is no new selection of the prior. In an attempt to fit better the data, one could possibly estimate a richer model including a parameter for the threshold as in Ortoleva (2012)’s model. Nevertheless, our experiment gives support to his model of updating also in comparison to the BU and FBU models.

We have estimated the FBU model joint with maxmin preferences, the only one that, to the best of our knowledge, has been axiomatized (by Pires, 2002). It is sometimes claimed that maxmin preferences imply that agents are very pessimistic (since they consider the worst outcome), and one may think that they imply that subjects are too pessimistic in the context of our experiment. It should be noticed, however, that we did estimate the bounds $[\mu_*, \mu^*]$ and in this sense we did not constrain our subjects to be overly pessimistic (as it would have been the case had we imposed $\mu^* = 1$). Nonetheless, we also considered a more general criterion, proposed by Hurwicz (1951), in which an agent considers the best and worst outcomes of his decision and then makes his choice weighing the two extreme outcomes on the basis of his preferences. If he put all the weight, represented by a parameter λ , on the worst outcome, he would behave as in our FBU model; if he chose $\lambda = 0$, he would be extremely optimistic; intermediate values of λ indicate intermediate values of pessimism. Optimism in this model may help to explain our data. For instance, if $\mu_* = 0$ and $\mu^* = 1$ and $\lambda = 0$, an agent would choose 70 (the most extreme belief in the support) as a first action, and then 30 (again the most extreme belief) after receiving a contradicting signal, which is in line with the observed asymmetric updating. On the other hand, from a behavioral viewpoint, this is not the most appealing explanation: being optimistic means trusting the predecessor after observing him (“being optimistic that the predecessor is rational”), and, then distrusting him after receiving a contradicting signal (“being optimistic that the predecessor is a noise type”). Nevertheless, we estimated the model and obtained $\lambda = 0.17$, $\mu_* = 0.2$, $\mu^* = 0.68$, indicating some form of optimism. Using the same test for model selection explained above, we obtain a p -value of 0.6: that is, this model does not fit the data significantly better than the BU model.⁴⁰

A different approach to the problem would be to use the principle of indifference or insufficient reason. According to this principle, typically attributed to Jacob Bernoulli or Laplace, in the absence of a convincing reason, the sub-

⁴⁰ Another approach considered in the literature is the so-called minimax regret theory, first proposed by Savage (1954). An agent would compute, for each action, his maximum regret and then choose the action to minimize it. Intuitively, given that the action set is fixed, the predictions of this model would not be very different from the Hurwicz (1951)’s model for an intermediate value of λ (as the resulting behavior would be a good way to minimize the largest distance to the optimal action when varying the prior belief). It should be noticed that in the context of our experiment, regret modeled in such a way would represent a purely subjective construction in subjects’ mind. Subjects never had access to information about the predecessor’s type, actually not even to the signal the predecessor received. It is, therefore, not very compelling to assume that subjects could feel such mentally constructed regret.

ject would give the same probability to different events. In the context of our experiment, this would mean that a subject at time 2, not having any reason to attach a specific weight to the probability μ that the predecessor is noise, would simply use a uniform as a distribution of μ . In such a case, however, he would behave as in the Bayesian model. Clearly, this model cannot perform better than our BU model, in which we have estimated the parameter μ .

8 Conclusion

Our experiment is relevant for two different literatures: that on social learning and that on belief updating.

A long debate in the social learning literature has concerned how subjects treat their private information versus the information coming from the choices of others. This question is indeed at the core of this literature. A phenomenon frequently documented is that human subjects tend to rely more on their private information than on the public information, compared to the full rationality benchmark. Our experimental design let us study this issue in much more detail. We discovered that subjects tend to put more weight on their own information when it is in contrast with the public information (revealed by the choice of another subject), whereas they put approximately the correct weight when it agrees with it. This behavior could not be observed in previous experiments. Previous studies were mainly designed to study the occurrence of informational cascades. They found that when subjects are in a situation of potential herding (that is, they received a signal at odds with the history of predecessors' actions), they require a number of predecessors choosing the same action larger than the theoretical one in order to go against their signal. On the other hand, when subjects receive a signal in agreement with the previous history of actions, they typically follow it. The first type of decision is in line with our result (but gives coarser information on subjects' updating); the second is essentially uninformative on how subjects weigh the signal.

This result is incompatible with Bayesian updating of beliefs. It is instead explained by a form of updating of multiple priors known in the decision theory literature as Maximum Likelihood Updating. This updating rule consists in using new information for two purposes: first to select a prior in the set of multiple priors; second, to update that prior. There is an important issue in this updating rule. In our experiment, a subject has to update twice, first after observing a predecessor's action and, second, after observing a private signal too. In this multistage updating problem, from a theory viewpoint it is somehow unclear whether a different prior can be selected after new information arrives or whether once the prior is selected, the agent should stick to it (as if ambiguity were resolved for ever). Our experimental data are explained by a model in which the prior is selected after each new piece of information. In the decision theory literature, it is somehow claimed (Gilboa and Marinacci, 2013; Pires, 2002) that the MLU rule (with the property that a prior is picked once and for ever) is an extreme form of updating, since it only relies on one prior.

Our model, letting the agent change his prior after receiving new information, can be seen as an intermediate rule of updating between the standard MLU and the FBU in which all priors are updated. In our model, only one prior is selected, but after new information the selection can change. We hope that this and other future experiments will inform the debate in decision theory on how to update multiple priors.

References

- [1] Anderson, L., Holt, C., 1997. Information Cascades in the Laboratory. *American Economic Review*, 87(5), 847-862.
- [2] Banerjee, A., 1992. A Simple Model of Herd Behavior. *Quarterly Journal of Economics*, 107(3), 797-817.
- [3] Bikhchandani, S., Hirshleifer, D., Welch, I., 1992. A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades. *Journal of Political Economy*, 110(5), 992-1026.
- [4] Becker, G., DeGroot, M., Marschak, J., 1964. Measuring Utility by a Single-Response Sequential Method. *Behavioral Science*, 9(3), 226-232.
- [5] Berger, J., 1985. Statistical Decision Theory and Bayesian Analysis. *Springer Series in Statistics*.
- [6] Cipriani, M., Guarino, A., 2005. Herd Behavior in a Laboratory Financial Market. *American Economic Review*, 95(5), 1427-1443.
- [7] Cipriani, M., Guarino, A., 2009. Herd Behavior in Financial Markets: An Experiment with Financial Market Professionals. *Journal of the European Economic Association*, 7(1), 206-233.
- [8] Çelen, B., Kariv, S., 2004. Distinguishing Informational Cascades from Herd Behaviour in the Laboratory. *American Economic Review*, 94(3), 484-498.
- [9] Cohen, M., Gilboa, I., Jaffray, J., Schmeidler, D., 2000. An Experimental Study of Updating Ambiguous Beliefs. *Risk, Decision and Policy*, 5(02), 123-133.
- [10] Drehmann, M., Oechssler, J., Roider, A., 2005. Herding and Contrarian Behavior in Financial Markets: An Internet Experiment. *American Economic Review*, 95(5), 1403-1426.
- [11] Dominiak, A., Duersch, P., Lefort J., 2012. A Dynamic Ellsberg Urn Experiment. *Games and Economic Behavior*, 75(2), 625-638.
- [12] Epstein, L., G., Schneider, M., 2007. Learning Under Ambiguity. *The Review of Economic Studies*, 74(4), 1275-1303.

- [13] Fischbacher, U. , 2007. Z-Tree: Zurich Toolbox for Ready-Made Economic Experiments. *Experimental Economics*, 10(2), 171-178.
- [14] Gilboa, I., Marinacci, M., 2013. Ambiguity and the Bayesian Paradigm. In: Acemoglu, D., Arellano, M., Dekel, E., (eds.) *Advances in Economics and Econometrics. Cambridge University Press*, 1,179-242.
- [15] Gilboa, I., Schmeidler, D.,1993. Updating Ambiguous Beliefs. *The Journal of Economic theory*, 59(1), 33-49.
- [16] Goeree, J., Palfrey, T., Rogers, B., McKelvey, R., 2007. Self-Correcting Information Cascades. *Review of Economic Studies*, 74(3), 733-762.
- [17] Good, I., 1965. The Estimation of Probabilities: An Essay on Modern Bayesian Methods. *The MIT Press*.
- [18] Guarino, A., Jehiel, P., 2013. Social Learning with Coarse Inference. *American Economic Journal: Microeconomics*, 5(1), 147-74.
- [19] Hanany, E., Klibanoff, P., 2007. Updating Preferences with Multiple Priors. *Theoretical Economics*, 2(3), 261–298.
- [20] Hurwicz, L., 1951. The Generalised Bayes Minimax Principle: A Criterion for Decision Making Under Uncertainty. *Cowles Commission Discussion Paper Statistics*, 355.
- [21] Lee, I., 1993. On the Convergence of Informational Cascades. *Journal of Economic Theory*, 61(2), 395-411.
- [22] Newey, W., McFadden, D., 1994. Large Sample Estimation and Hypothesis Testing. In: Engle, R., McFadden, D., (eds.) *Handbook of Econometrics. Elsevier*, 4, 2111–2245.
- [23] Noth, M., Weber, M., 2003. Information Aggregation with Random Ordering: Cascades and Overconfidence. *The Economic Journal*, 113(84), 166-189.
- [24] Nyarko, Y., Schotter, A., 2002. An Experimental Study of Belief Learning Using Elicited Beliefs. *Econometrica*, 70(3), 971–1005.
- [25] Ortoleva, P.. Modelling the Change of Paradigme": Non-Bayesian Reactions to Unexpected News. *American Economic Review*, 102(6), 2410-2436.
- [26] Pires, C., 2002. A Rule for Updating Ambiguous Beliefs. *Theory and Decision*, 53(2), 137–152.
- [27] Savage, L., 1954. *The Foundations of Statistics. Wiley*, New York.
- [28] Schotter, A. 2003. Decision Making with Naive Advice. *American Economic Review PP*, 93(2), 196-201.
- [29] Schotter, A., Trevino, I., 2014. Belief Elicitation in the Lab. *Annual Review of Economics*, 6, 103-128.

APPENDIX (FOR ONLINE PUBLICATION ONLY)

9 Appendix A

9.1 More descriptive statistics

One could observe that if a subject chose, e.g., $a_{2i}^1 = 84$ and then, after receiving a bad signal, chose $a_{2i}^2 = 50$, the corresponding α_{2i}^2 would be 2, which is compatible with the overweight we documented. It must be noticed, though, that if we exclude the cases in which $a_{2i}^2 = 50$, nevertheless the asymmetry remains, as one can appreciate by looking at the following table.

	First Quartile	Median	Third Quartile
α_2^2	0.72	1.16	2.11
α_2^2 (upon observing confirming signal)	0.55	0.96	1.36
α_2^2 (upon observing contradicting signal)	1.30	2.07	2.98

Table 11: Distribution of weights on the own signal in the SL treatment.

The table shows the quartiles of the distribution of the weight on the own signal for the second action at time 2 in the SL treatment. The data refer to all cases in which the first action at time 2 was different from 50; moreover, cases in which the second action at time 2 was equal to 50 are excluded.

9.2 Social Learning Treatments

The social learning treatments SL1, SL2 and SL3 differ in some dimensions (length of the sequence, precision of the signal). Our results, however, are not significantly different across treatments. Specifically, we ran a Mann-Witney U test (Wilcoxon rank-sum test) on the medians of each session (the most conservative option to guarantee independence of observations) for time 1, as well as for the first decision at time 2 and the second decision at time 2 (for confirming and contradicting signals). The p-values are reported in Table A2.

	Time 1	Time 2.1	Time 2.2 - confirming	Time 2.2 - contradicting
SL1 versus SL2	0.50	0.09	0.84	0.10
SL1 versus SL3	0.08	0.008	0.30	0.14
SL2 versus SL3	0.22	0.69	0.29	1.00

Table 12: Tests for the SL treatments.

The table shows the results of Mann-Witney U test (Wilcoxon rank-sum test). The null hypothesis is that the medians come from the same distribution. In the table we report the p-values.

Ignoring the multiple hypothesis testing issue, we would reject the null hypothesis for one case (equivalence of SL1 versus SL3 at Time 2.1, i.e., for the first

action at time 2) at significance level 5%. The simple Bonferroni correction for multiple hypothesis tests controlling the family-wise error rate at 5%, however, lowers the critical p-value to 0.004, and we do not reject the joint null after this correction.

9.3 Estimation and test

Let us illustrate the details of the GMM estimation and of the model specification test.

9.3.1 GMM estimation

Estimating the LRTU model

Let us consider first the estimation of the LRTU model. The parameters to be estimated are $\theta \equiv (\mu_*, \mu^*, c)$, $0 \leq \mu_* \leq \mu^* \leq 1$, and $c \geq 0$. To make the dependence on the parameters explicit, we express the LRTU model actions obtained in the main text as $a_{2i}^1(\mu_{2i}^1, q_{1i}^S; \theta)$, $a_{2i}^2(\mu_{2i,confirm}^2, q_{1i}^S, q_{2i}^S; \theta)$, and $a_{2i}^2(\mu_{2i,contrdict}^2, q_{1i}^S, q_{2i}^S; \theta)$. For given θ , a_{1i} , and $s_{2i} = 1$, the heterogeneity in subjective signal precisions generates the joint distribution of the time 2 actions $(a_{2i}^1(\mu_{2i}^1, q_{1i}^S; \theta), a_{2i}^2(\mu_{2i,confirm}^2, q_{1i}^S, q_{2i}^S; \theta))$. If the LRTU model were the true data generating process, then, at the true value of θ , the conditional distribution of $(a_{2i}^1(\mu_{2i}^1, q_{1i}^S; \theta), a_{2i}^2(\mu_{2i,confirm}^2, q_{1i}^S, q_{2i}^S; \theta))$ given $(a_{1i}, s_{2i} = 1)$ generated from heterogeneous (q_{1i}^S, q_{2i}^S) would coincide with the actual conditional distribution of (a_{2i}^1, a_{2i}^2) . This implies that, for any integrable function $h(a_{2i}^1, a_{2i}^2)$,

$$E[h(a_{2i}^1, a_{2i}^2) - E_Q[h(a_{2i}^1(\mu_{2i}^1, q_{1i}^S; \theta), a_{2i}^2(\mu_{2i,confirm}^2, q_{1i}^S, q_{2i}^S; \theta)) | a_{1i}, s_{2i} = 1]] = 0$$

holds at the true θ for every a_{1i} , where the inner expectation $E_Q[\cdot]$ is the expectation with respect to the joint distribution Q of (q_1, q_2) , which we assume be independent of (a_{1i}, s_{2i}) , and the outer expectation is with respect to the actual sampling distribution of (a_{2i}^1, a_{2i}^2) conditional on a_{1i} and $s_{2i} = 1$. Specifically, as we said, for Q we use the empirical distribution of precisions. Hence,

$$\begin{aligned} & E_Q[h(a_{2i}^1(\mu_{2i}^1, q_{1i}^S; \theta), a_{2i}^2(\mu_{2i,confirm}^2, q_{1i}^S, q_{2i}^S; \theta))] \\ & \approx \frac{1}{J} \sum_j h(a_{2i}^1(\mu_{2i}^1, q_{1j}^S; \theta), a_{2i}^2(\mu_{2i,confirm}^2, q_{1j}^S, q_{2j}^S; \theta)), \end{aligned}$$

where the index j indicates an observation of (q_{1j}^S, q_{2j}^S) and J is the number of observations of (q_{1j}^S, q_{2j}^S) available in our dataset. Specifically, when $h(\cdot, \cdot)$ involves only a_{2i}^1 , the marginal distribution of q_1^S suffices to compute $E_Q(h(a_{2i}^1))$. Therefore, we construct the empirical distribution of q_1^S by pooling the rational actions at time 1 ($a_{1i} \geq 50$) in the SL and IDM treatments

($J = 1331$). When $h(\cdot, \cdot)$ involves both a_{2i}^1 and a_{2i}^2 , we construct the empirical distribution of (q_1^S, q_2^S) using the observations (a_{1i}, a_{2i}) in the IDM treatment only, restricted to $50 \leq a_{1i} < 100$ and $a_{2i} \geq 50$.⁴¹ The total number of observations used to construct the empirical distribution of (q_1^S, q_2^S) amounts to $J = 440$.

Similarly, for the contradicting signal case we have that

$$E [h(a_{2i}^1, a_{2i}^2) - E_Q [h(a_{2i}^1(\mu_{2i}^1, q_1^S; \theta), a_{2i}^2(\mu_{2i}^2, \text{contradict}, q_1^S, q_2^S; \theta))] | a_{1i}, s_{2i} = 0] = 0$$

holds for any a_{1i} .

These moment conditions imply the following unconditional moment conditions:

$$E [s_{2i} \cdot (h(a_{2i}^1, a_{2i}^2) - E_Q [h(a_{2i}^1(\mu_{2i}^1, q_1^S; \theta), a_{2i}^2(\mu_{2i}^2, \text{confirm}, q_1^S, q_2^S; \theta))])] = 0, \quad (30)$$

$$E [(1 - s_{2i}) \cdot (h(a_{2i}^1, a_{2i}^2) - E_Q [h(a_{2i}^1(\mu_{2i}^1, q_1^S; \theta), a_{2i}^2(\mu_{2i}^2, \text{contradict}, q_1^S, q_2^S; \theta))])] = 0. \quad (31)$$

When $h(a_{2i}^1, a_{2i}^2)$ only depends on a_{2i}^1 , s_{2i} plays no role and the moment conditions (30) and (31) reduce (with a slight abuse of notation) to

$$E [h(a_{2i}^1) - E_Q [h(a_{2i}^1(\mu_{2i}^1, q_1^S; \theta))]] = 0. \quad (32)$$

Given a specification for $h(\cdot)$, we estimate θ by applying GMM to the unconditional moment conditions (30) - (32).

Specifically, our approach is to match the cumulative distribution functions (cdfs) of α predicted by the models with the empirical distributions. Recall that $(\alpha_{2i}^1, \alpha_{2i}^2)$ can be written in terms of (a_{2i}^1, a_{2i}^2) as

$$\begin{aligned} \text{time 2.1:} \quad \alpha_{2i}^1 &= \frac{l(a_{2i}^1/100)}{l(0.7)}, \\ \text{time 2.2-confirming:} \quad \alpha_{2i}^2 &= \frac{l(a_{2i}^2/100) - l(a_{2i}^1/100)}{l(0.7)}, \\ \text{time 2.2-contradicting:} \quad \alpha_{2i}^2 &= \frac{l(a_{2i}^2/100) - l(a_{2i}^1/100)}{l(0.3)}. \end{aligned}$$

To match the cdfs of α 's evaluated at $t \in [0, \infty)$, we specify $h(\cdot, \cdot)$ as

$$h(a_{2i}^1) = 1 \left\{ \frac{l(a_{2i}^1/100)}{l(0.7)} \leq t \right\},$$

when we match the cdf of α_{2i}^1 , and specify $h(\cdot, \cdot)$ as

$$\begin{aligned} h(a_{2i}^1, a_{2i}^2) &= 1 \left\{ \frac{l(a_{2i}^2/100) - l(a_{2i}^1/100)}{l(0.7)} \leq t \right\} \text{ and} \\ h(a_{2i}^1, a_{2i}^2) &= 1 \left\{ \frac{l(a_{2i}^2/100) - l(a_{2i}^1/100)}{l(0.3)} \leq t \right\}, \end{aligned}$$

⁴¹We drop observations $a_{1i} = 100$ since we cannot impute a unique value of q_{2i}^S on the basis of the observed a_{2i} .

when we match the cdf of α_{2i}^2 for the confirming and contradicting signal case, respectively.

Since we discretise the action space of a_{1i} into three intervals (“bins”) $B_1 = (50, 66.7]$, $B_2 = (66.7, 83.4]$ and $B_3 = (83.4, 1]$ and the theoretical predictive distribution of α vary over a_{1i} only across these three bins, we focus on the distributions of α_{2i}^1 and α_{2i}^2 conditional on a_{1i} being in each of these three bins. We compute the distributions of α for time 2.1 as well as for time 2.2, distinguishing between the confirming and the contradicting signal case. Overall, we obtain nine empirical distributions of α (three for each bin) to be matched with the corresponding distributions of α 's predicted by the theoretical model.

We match the cdfs of α at the three points of the support corresponding to the empirical quartiles of α conditional on $a_{1i} \in B$, with $B \in \{B_1, B_2, B_3\}$. For $p \in \{0.25, 0.5, 0.75\}$ and $B \in \{B_1, B_2, B_3\}$, we denote the p -th quartile of α_{2i}^1 conditional on action $a_{1i} \in B$ by $t_{2,p,B}^1$, the p -th quartile of α_{2i}^2 conditional on action $a_{1i} \in B$ and $s_{2i} = 1$ by $t_{2,conf,p,B}^2$, and the p -th quartile of α_{2i}^2 conditional on action $a_{1i} \in B$ and $s_{2i} = 0$ by $t_{2,cont,p,B}^2$.

Given the underlying parameter vector θ and the signal precisions (q_1^S, q_2^S) , the theoretical α 's can be written as

$$\text{time 2.1: } \alpha_{2i}^1(\theta, q_1^S) = \frac{l(a_{2i}^1(\mu_{2i}^1, q_1^S; \theta) / 100)}{l(0.7)},$$

time 2.2-confirming :

$$\begin{aligned} \alpha_{2i,conf}^2(\theta, q_1^S, q_2^S) &= \frac{l(a_{2i}^2(\mu_{2i,conf}^2, q_1^S, q_2^S; \theta) / 100) - l(a_{2i}^1(\mu_{2i}^1, q_1^S; \theta) / 100)}{l(0.7)} \\ &= \frac{l(q_2^S) + l(a_{2i}^1(\mu_{2i,conf}^2, q_1^S; \theta) / 100) - l(a_{2i}^1(\mu_{2i}^1, q_1^S; \theta) / 100)}{l(0.7)}, \end{aligned}$$

time 2.2-contradicting :

$$\begin{aligned} \alpha_{2i,cont}^2(\theta, q_1^S, q_2^S) &= \frac{l(a_{2i}^2(\mu_{2i,contradict}^2, q_1^S, q_2^S; \theta) / 100) - l(a_{2i}^1(\mu_{2i}^1, q_1^S; \theta) / 100)}{l(0.3)} \\ &= \frac{l(1 - q_2^S) + l(a_{2i}^1(\mu_{2i,contradict}^2, q_1^S; \theta) / 100) - l(a_{2i}^1(\mu_{2i}^1, q_1^S; \theta) / 100)}{l(0.3)}. \end{aligned}$$

The predicted distributions of α given $a_{1i} \in B$ (and s_{2i} for the second action at time 2) is obtained by viewing $\alpha_{2i}^1(\theta, q_1^S)$ and $\alpha_{2i}^2(\theta, q_1^S, q_2^S)$ as random variables with their probability distributions generated from the empirical distribution of the heterogeneous signal precisions $(q_1^S, q_2^S) \sim Q$.

Since we match the 9 distributions of α at three points of the support, we have in total the following 27 moment conditions:

$$\underbrace{\mathbf{m}_i^L(\theta)}_{27 \times 1} = \begin{pmatrix} \mathbf{m}_{1i}^L(\theta) \\ \mathbf{m}_{2i,conf}^L(\theta) \\ \mathbf{m}_{2i,cont}^L(\theta) \end{pmatrix},$$

where $\mathbf{m}_{1i}^{LRT}(\theta)$ is a 9×1 vector of moment conditions concerning the cdfs of α_{2i}^1 :

$$\underbrace{\mathbf{m}_{1i}^L(\theta)}_{9 \times 1} = \begin{pmatrix} 1\{a_{1i} \in B_1\} \cdot \begin{pmatrix} 1\{\alpha_{2i}^1 \leq t_{2,0.25,B_1}^1\} - E_Q(1\{\alpha_{2i}^1(\theta, q_1^S) \leq t_{2,0.25,B_1}^1\}) \\ 1\{\alpha_{2i}^1 \leq t_{2,0.5,B_1}^1\} - E_Q(1\{\alpha_{2i}^1(\theta, q_1^S) \leq t_{2,0.5,B_1}^1\}) \\ 1\{\alpha_{2i}^1 \leq t_{2,0.75,B_1}^1\} - E_Q(1\{\alpha_{2i}^1(\theta, q_1^S) \leq t_{2,0.75,B_1}^1\}) \end{pmatrix} \\ \vdots \\ 1\{a_{1i} \in B_3\} \cdot \begin{pmatrix} 1\{\alpha_{2i}^1 \leq t_{2,0.25,B_3}^1\} - E_Q(1\{\alpha_{2i}^1(\theta, q_1^S) \leq t_{2,0.25,B_3}^1\}) \\ 1\{\alpha_{2i}^1 \leq t_{2,0.5,B_3}^1\} - E_Q(1\{\alpha_{2i}^1(\theta, q_1^S) \leq t_{2,0.5,B_3}^1\}) \\ 1\{\alpha_{2i}^1 \leq t_{2,0.75,B_3}^1\} - E_Q(1\{\alpha_{2i}^1(\theta, q_1^S) \leq t_{2,0.75,B_3}^1\}) \end{pmatrix} \end{pmatrix}, \quad (33)$$

and $\mathbf{m}_{2i,conf}^L(\theta)$ and $\mathbf{m}_{2i,cont}^L(\theta)$ are 9×1 vectors of moment conditions concerning the cdfs of α_{2i}^2 for confirming and contradicting signal cases, respectively:

$$\begin{aligned} & \underbrace{\mathbf{m}_{2i,conf}^{LRTU}(\theta)}_{9 \times 1} \\ &= s_{2i} \cdot \begin{pmatrix} 1\{a_{1i} \in B_1\} \cdot \begin{pmatrix} 1\{\alpha_{2i}^2 \leq t_{2,conf,0.25,B_1}^2\} - E_Q(1\{\alpha_{2i,conf}^2(\theta, q_1^S, q_2^S) \leq t_{2,conf,0.25,B_1}^2\}) \\ 1\{\alpha_{2i}^2 \leq t_{2,conf,0.5,B_1}^2\} - E_Q(1\{\alpha_{2i,conf}^2(\theta, q_1^S, q_2^S) \leq t_{2,conf,0.5,B_1}^2\}) \\ 1\{\alpha_{2i}^2 \leq t_{2,conf,0.75,B_1}^2\} - E_Q(1\{\alpha_{2i,conf}^2(\theta, q_1^S, q_2^S) \leq t_{2,conf,0.75,B_1}^2\}) \end{pmatrix} \\ \vdots \\ 1\{a_{1i} \in B_3\} \cdot \begin{pmatrix} 1\{\alpha_{2i}^2 \leq t_{2,conf,0.25,B_3}^2\} - E_Q(1\{\alpha_{2i,conf}^2(\theta, q_1^S, q_2^S) \leq t_{2,conf,0.25,B_3}^2\}) \\ 1\{\alpha_{2i}^2 \leq t_{2,conf,0.5,B_3}^2\} - E_Q(1\{\alpha_{2i,conf}^2(\theta, q_1^S, q_2^S) \leq t_{2,conf,0.5,B_3}^2\}) \\ 1\{\alpha_{2i}^2 \leq t_{2,conf,0.75,B_3}^2\} - E_Q(1\{\alpha_{2i,conf}^2(\theta, q_1^S, q_2^S) \leq t_{2,conf,0.75,B_3}^2\}) \end{pmatrix} \end{pmatrix} \\ & \underbrace{\mathbf{m}_{2i,cont}^{LRTU}(\theta)}_{9 \times 1} \\ &= (1 - s_{2i}) \cdot \begin{pmatrix} 1\{a_{1i} \in B_1\} \cdot \begin{pmatrix} 1\{\alpha_{2i}^2 \leq t_{2,cont,0.25,B_1}^2\} - E_Q(1\{\alpha_{2i,cont}^2(\theta, q_1^S, q_2^S) \leq t_{2,cont,0.25,B_1}^2\}) \\ 1\{\alpha_{2i}^2 \leq t_{2,cont,0.5,B_1}^2\} - E_Q(1\{\alpha_{2i,cont}^2(\theta, q_1^S, q_2^S) \leq t_{2,cont,0.5,B_1}^2\}) \\ 1\{\alpha_{2i}^2 \leq t_{2,cont,0.75,B_1}^2\} - E_Q(1\{\alpha_{2i,cont}^2(\theta, q_1^S, q_2^S) \leq t_{2,cont,0.75,B_1}^2\}) \end{pmatrix} \\ \vdots \\ 1\{a_{1i} \in B_3\} \cdot \begin{pmatrix} 1\{\alpha_{2i}^2 \leq t_{2,cont,0.25,B_3}^2\} - E_Q(1\{\alpha_{2i,cont}^2(\theta, q_1^S, q_2^S) \leq t_{2,cont,0.25,B_3}^2\}) \\ 1\{\alpha_{2i}^2 \leq t_{2,cont,0.5,B_3}^2\} - E_Q(1\{\alpha_{2i,cont}^2(\theta, q_1^S, q_2^S) \leq t_{2,cont,0.5,B_3}^2\}) \\ 1\{\alpha_{2i}^2 \leq t_{2,cont,0.75,B_3}^2\} - E_Q(1\{\alpha_{2i,cont}^2(\theta, q_1^S, q_2^S) \leq t_{2,cont,0.75,B_3}^2\}) \end{pmatrix} \end{pmatrix} \end{aligned}$$

Since the number of moment conditions is greater than the number of unknown parameters, we obtain a point estimator of θ by minimizing the overidentified GMM objective function in two steps. In the first step, we solve

$$\hat{\theta} = \arg \min_{\theta} \left(\sum_{i=1}^n \mathbf{m}_i^L(\theta) \right)' \left(\sum_{i=1}^n \mathbf{m}_i^L(\theta) \right),$$

and, in the second step, we solve

$$\hat{\theta}_{GMM} = \arg \min_{\theta} \left(\frac{1}{n} \sum_{i=1}^n \mathbf{m}_i^L(\theta) \right)' \hat{W}^{-1} \left(\frac{1}{n} \sum_{i=1}^n \mathbf{m}_i^L(\theta) \right),$$

where

$$\hat{W} = \frac{1}{n} \sum_{i=1}^n \mathbf{m}_i^L(\hat{\theta}) \mathbf{m}_i^L(\hat{\theta})'.$$

The optimization for $\hat{\theta}$ and $\hat{\theta}_{GMM}$ is carried out by grid search with grid size 0.01.

Estimating the BU model

The BU model is a special case of the LRTU model in which $\mu_* = \mu^* = \mu$. In this case c becomes an irrelevant parameter, and the only parameter to estimate is $\theta = \mu \in [0, 1]$. Furthermore, note that the theoretical α_{2i}^2 is given by $s_{2i} l(q_2^S) + (1 - s_{2i}) l(1 - q_2^S)$ (which is independent of the parameters) when $\mu_* = \mu^* = \mu$. Hence, the identifying information for μ only comes from the cdf of α_{2i}^1 . Nevertheless, in the two-step GMM procedure, we make use of the full set of moment conditions (27×1), since the first-stage estimate does not necessarily equal to the second-stage estimate due to the non-block-diagonal weighting matrix. The set of moment conditions is given by

$$\underbrace{\mathbf{m}_i^B(\theta)}_{27 \times 1} = \begin{pmatrix} \mathbf{m}_{1i}^B(\theta) \\ \mathbf{m}_{2i,conf}^B \\ \mathbf{m}_{2i,cont}^B \end{pmatrix},$$

where these moment conditions are the moment conditions of the LRTU model constrained to $\mu_* = \mu^* = \mu$. Since only the first set of moment conditions $\mathbf{m}_{1i}^B(\theta)$ depends on μ , an initial GMM estimator minimizes

$$\hat{\mu} = \arg \min_{\mu} \left(\sum_{i=1}^n \mathbf{m}_{1i}^B(\theta) \right)' \left(\sum_{i=1}^n \mathbf{m}_i^B(\theta) \right). \quad (34)$$

The optimal 2-step GMM estimator then minimizes the variance weighted GMM objective functions with the *full* set of moment conditions,

$$\begin{aligned} \hat{\mu}_{GMM} &= \arg \min_{\mu} \left(\frac{1}{n} \sum_{i=1}^n \mathbf{m}_i^B(\theta) \right)' \hat{W}^{-1} \left(\frac{1}{n} \sum_{i=1}^n \mathbf{m}_i^B(\theta) \right), \\ \hat{W} &= \frac{1}{n} \sum_{i=1}^n \mathbf{m}_i^B(\hat{\theta}) \mathbf{m}_i^B(\hat{\theta})' \text{ with } \hat{\theta} = \hat{\mu}. \end{aligned} \quad (35)$$

Again, a grid search with grid size 0.01 is used to find $\hat{\mu}$ and $\hat{\mu}_{GMM}$.

Estimating the FBU model

In the FBU model, the unknown parameters are $\theta = (\mu_*, \mu^*)$, $0 \leq \mu_* \leq \mu^* \leq 1$. Since we only consider the realization of q_{1i}^S greater than 0.5, the range of beliefs for the first action at time 2 is a subset of $[\frac{1}{2}, 1]$ (see expression (??)), and the maximin action $a_{2i, \max \min}^1$ is the Bayes's action with the implied prior μ^* . Hence, the moment conditions for the FBU model concerning the cdf of α_2^1 are obtained by replacing $\alpha_2^1(\theta, q_1^S)$ in (33) with

$$\alpha_{2i, F}^1(\theta, q_1^S) = \frac{l(a_{2i}^1(\mu^*, q_1^S)/100)}{l(0.7)}.$$

We then denote the resulting 9 moment conditions by $\mathbf{m}_{1i}^F(\theta)$.

As for the moment conditions for the cdfs of α_2^2 , we cannot fix the implied prior as it depends on the individual's (q_{1i}, q_{2i}) . Nevertheless, given $(\theta, q_{1i}^S, q_{2i}^S)$, the maxmin action can be pinned down according to the formula $a_{2i, F}^2(\theta, q_{1i}^S, q_{2i}^S)$ given in Section 6. Accordingly, we can obtain the moment conditions concerning the cdfs of α_2^2 by

$$\begin{aligned} & \underbrace{\mathbf{m}_{2, \text{conf}, i}^F(\theta)}_{9 \times 1} \\ &= s_{2i} \cdot \begin{pmatrix} 1\{a_{1i} \in B_1\} \cdot \begin{pmatrix} 1\{\alpha_{2i}^2 \leq t_{2, \text{conf}, 0.25, B_1}^2\} - E_Q \left(1\{\alpha_{2i, \text{conf}}^{2, \max \min}(\theta, q_1^S, q_2^S) \leq t_{2, \text{conf}, 0.25, B_1}^2\} \right) \\ 1\{\alpha_{2i}^2 \leq t_{2, \text{conf}, 0.5, B_1}^2\} - E_Q \left(1\{\alpha_{2i, \text{conf}}^{2, \max \min}(\theta, q_1^S, q_2^S) \leq t_{2, \text{conf}, 0.5, B_1}^2\} \right) \\ 1\{\alpha_{2i}^2 \leq t_{2, \text{conf}, 0.75, B_1}^2\} - E_Q \left(1\{\alpha_{2i, \text{conf}}^{2, \max \min}(\theta, q_1^S, q_2^S) \leq t_{2, \text{conf}, 0.75, B_1}^2\} \right) \end{pmatrix} \\ \vdots \\ 1\{a_{1i} \in B_3\} \cdot \begin{pmatrix} 1\{\alpha_{2i}^2 \leq t_{2, \text{conf}, 0.25, B_3}^2\} - E_Q \left(1\{\alpha_{2i, \text{conf}}^{2, \max \min}(\theta, q_1^S, q_2^S) \leq t_{2, \text{conf}, 0.25, B_3}^2\} \right) \\ 1\{\alpha_{2i}^2 \leq t_{2, \text{conf}, 0.5, B_3}^2\} - E_Q \left(1\{\alpha_{2i, \text{conf}}^{2, \max \min}(\theta, q_1^S, q_2^S) \leq t_{2, \text{conf}, 0.5, B_3}^2\} \right) \\ 1\{\alpha_{2i}^2 \leq t_{2, \text{conf}, 0.75, B_3}^2\} - E_Q \left(1\{\alpha_{2i, \text{conf}}^{2, \max \min}(\theta, q_1^S, q_2^S) \leq t_{2, \text{conf}, 0.75, B_3}^2\} \right) \end{pmatrix} \end{pmatrix} \\ & \underbrace{\mathbf{m}_{2i, \text{cont}}^{\max \min}(\theta)}_{9 \times 1} \\ &= (1 - s_{2i}) \cdot \begin{pmatrix} 1\{a_{1i} \in B_1\} \cdot \begin{pmatrix} 1\{\alpha_{2i}^2 \leq t_{2, \text{cont}, 0.25, B_1}^2\} - E_Q \left(1\{\alpha_{2i, \text{cont}}^{2, \max \min}(\theta, q_1, q_2) \leq t_{2, \text{cont}, 0.25, B_1}^2\} \right) \\ 1\{\alpha_{2i}^2 \leq t_{2, \text{cont}, 0.5, B_1}^2\} - E_Q \left(1\{\alpha_{2i, \text{cont}}^{2, \max \min}(\theta, q_1, q_2) \leq t_{2, \text{cont}, 0.5, B_1}^2\} \right) \\ 1\{\alpha_{2i}^2 \leq t_{2, \text{cont}, 0.75, B_1}^2\} - E_Q \left(1\{\alpha_{2i, \text{cont}}^{2, \max \min}(\theta, q_1, q_2) \leq t_{2, \text{cont}, 0.75, B_1}^2\} \right) \end{pmatrix} \\ \vdots \\ 1\{a_{1i} \in B_3\} \cdot \begin{pmatrix} 1\{\alpha_{2i}^2 \leq t_{2, \text{cont}, 0.25, B_3}^2\} - E_Q \left(1\{\alpha_{2i, \text{cont}}^{2, \max \min}(\theta, q_1, q_2) \leq t_{2, \text{cont}, 0.25, B_3}^2\} \right) \\ 1\{\alpha_{2i}^2 \leq t_{2, \text{cont}, 0.5, B_3}^2\} - E_Q \left(1\{\alpha_{2i, \text{cont}}^{2, \max \min}(\theta, q_1, q_2) \leq t_{2, \text{cont}, 0.5, B_3}^2\} \right) \\ 1\{\alpha_{2i}^2 \leq t_{2, \text{cont}, 0.75, B_3}^2\} - E_Q \left(1\{\alpha_{2i, \text{cont}}^{2, \max \min}(\theta, q_1, q_2) \leq t_{2, \text{cont}, 0.75, B_3}^2\} \right) \end{pmatrix} \end{pmatrix} \end{aligned}$$

where

$$\begin{aligned} \text{for time 2.2-confirming} & : \\ \alpha_{2i,conf}^{2,\max\min}(\theta, q_1^S, q_2^S) & = \frac{l(a_{2i,F}^2(\theta, q_1^S, q_2^S)/100) - l(a_{2i}^1(\mu^*, q_1^S)/100)}{l(0.7)}, \end{aligned}$$

$$\begin{aligned} \text{for time 2.2-contradicing} & : \\ \alpha_{2i,cont}^{2,\max\min}(\theta, q_1^S, q_2^S) & = \frac{l(a_{2i,F}^2(\theta, q_1^S, q_2^S)/100) - l(a_{2i}^1(\mu^*, q_1^S)/100)}{l(0.3)}. \end{aligned}$$

The estimation of $\theta = (\mu_*, \mu^*)$ then proceeds by forming the moment vector

$$\underbrace{\mathbf{m}_i^F(\theta)}_{27 \times 1} = \begin{pmatrix} \mathbf{m}_{1i}^F(\theta) \\ \mathbf{m}_{2i,conf}^F(\theta) \\ \mathbf{m}_{2i,cont}^F(\theta) \end{pmatrix}$$

and running the same estimation procedure as in the LRTU model.

9.3.2 Resampling-based model comparison

We now turn to presenting the details of the implementation of the model comparison procedure shown in Section 6.

We consider as the null model the BU model with parameter value $\hat{\mu}_{GMM}$ (as reported in Table 10). As usual, we sample (q_{1i}^S, q_{2i}^S) randomly and with replacement from the empirical distribution. We then plug them into the formulae of the theoretical α 's, with (a_{1i}, s_{2i}) set at the values observed in the actual dataset. Having a random draw of (q_{1i}^S, q_{2i}^S) for each observation and computing the α_{2i}^1 and α_{2i}^2 for each i , we obtain a simulated sample from the null BU model with the same size as the actual data. We generate 1000 such samples and index them by $j = 1, 2, \dots, 1000$.

For each simulated dataset, we minimize the GMM objective functions in the BU model and the LRTU model. The minimized values of the objective functions are denoted by \hat{L}_B^j and \hat{L}_L^j , $j = 1, \dots, 1000$, respectively. To keep the weights on the moment conditions identical in the estimation of the BU and the LRTU models, we construct the GMM objective functions by choosing the weighting matrix used to obtain $\hat{\mu}_{GMM}$ for the actual data. We keep this weighting matrix fixed across samples.

We then approximate the null distribution of the difference of the GMM objective functions by the empirical distribution of $\Delta \hat{L}^j = \hat{L}_L^j - \hat{L}_B^j$, for $j = 1, \dots, 1000$. To obtain the p -value for the null model (the BU model) against the LRTU model, we compute $\Delta \hat{L}$, the difference of the GMM objective functions for our actual data. Of course, we use the same weighting matrix as the one used to compute $\Delta \hat{L}^j$, $j = 1, \dots, 1000$. The p -value is then obtained by the proportion of $\Delta \hat{L}^j$'s that are greater than $\Delta \hat{L}$. A small p -value (e.g., less than 5%) indicates that the LRTU model fits the actual data significantly better than the BU model, even taking into account the fitness gain only due to the over-parametrization of the LRTU model.

10 Appendix B: Instructions

Welcome to our experiment! We hope you will enjoy it.

You are about to take part in a study on decision making with 9 other participants. Everyone in the experiment has the same instructions. If something in the instructions is not clear and you have questions, please, do not hesitate to ask for clarification. We will be happy to answer your questions privately.

Depending on your choices, the other participants' choices and some luck you will earn some money. You will receive the money immediately after the experiment.

10.1 The Experiment

The experiment consists of 15 rounds of decision making. In each round you will make two consecutive decisions. All of you will participate in each round.

What you have to do

In each round, you have simply to choose a number between 0 and 100. You will make this choice twice, before and after receiving some information. The reason for these choices is the following. There is a good whose value can be either 0 or 100 units of a fictitious currency called "lira." You will not be told whether the good is worth 0 or 100 liras, but will receive some information about which value is more likely to have been chosen by a computer. We will ask you to predict the value of the good, that is, to indicate the chance that the value is 100 liras.

The value of the good

Whether the good will be worth 0 or 100 liras will be determined randomly at the beginning of each round: there will be a probability of 50% that the value is 0 and a probability of 50% that it is 100 liras, like in the toss of a coin. The computer chooses the value of the good in each round afresh. The value of the good in one round never depends on the value of the good in one of the previous rounds.

What you will know about the value

Although you will not be told the value of the good, you will, however, receive some information about which value is more likely to have been chosen. For each of you, the computer will use two "virtual urns" both containing green and red balls. The proportion of the two types of balls in each urn, however, is different. One urn contains more red than green balls, whereas the other urn contains more green than red balls. If the value of the good is 0, you will observe a ball drawn from an urn containing more red balls. If the value is 100, instead, you will observe a ball drawn from an urn containing more green balls. To recap:

- If the value is 100, then there are more GREEN balls in the urn.

- If the value is 0, then there are more RED balls in the urn.

Therefore, the ball color will give you some information about the value of the good. Below we will tell you more about how many balls there are in the urns. First, though, let us see more precisely what will happen in each round.

10.2 Procedures for each round

In each of the 15 rounds you will make decisions in sequence, one after the other. There will be 10 periods. Each of you will make her/his two choices only in one period, randomly chosen. Since there are 10 participants, this means that all of you will participate in each round.

The precise sequence of events is the following:

First: the computer program will decide randomly if the good for that round is worth 0 or 100 liras. You will not be told this value. On your screen you will read “Round 1 of 15. The computer is deciding the value of the good by flipping a coin.”

Second: the computer program will randomly select who is the first person who has to make a choice. Each of you has the same ($1/10th$) chance of being selected.

Third: the computer will draw a ball from the “virtual urn” and inform the first person (only the first person) of the drawn ball color. The first person will see this information on the screen. No one else will see it. The other participants will be waiting.

Fourth: after the person sees this information, (s)he has to choose a number between 0 and 100. This can be done by moving a slider on the screen (to select a precise number, please, use the arrows on your keyboard). The decision made will be visible to all participants.

Fifth: the computer will now randomly choose another person. Again, all the remaining 9 people have the same ($1/9th$) chance of being chosen.

Sixth: this second person, having observed the first person’s prediction, will be asked to make her/his prediction, choosing a number between 0 and 100. This decision will not be visible to other participants.

Seventh: after the decision, the computer will draw a ball from the “virtual urn” and inform (only) the second person of its color.

Eighth: the second person, after observing the ball color, will now make a new prediction, choosing again a number between 0 and 100. This choice is visible to all participants.

Ninth: the computer will choose a third person. This person will have to make two predictions, before and after receiving information, exactly as the second person. The first decision is after having observed the first two persons’ predictions. The second prediction is after having observed the ball color too. The decision made after seeing the ball color will be visible to everyone. Then the computer will choose the fourth person and so on, until all ten people have had the opportunity to participate.

Tenth: the computer will reveal the value of the good for the round and the payoff you earned in the round.

Observation 1: All 10 participants have to make the same type of decision, predicting the value of the good. However, the first person in the sequence is asked to make only one prediction, while the others will make two. The reason is simple. Since the first person knows nothing, the only sensible prediction is 50, given that there is a 50 – 50 chance that the value is 0 or 100 liras. Therefore, if you are the first, we do not ask you to make the prediction before seeing the ball color. Instead, if you are a subsequent person, we will ask you to make a prediction even before seeing the ball color, just after observing the predecessors' predictions. **Always recall that the predecessors' predictions that you will observe are the second predictions that they made, that is, the predictions they made after receiving information about the ball color.**

Observation 2: As we said, when it is your turn, the computer will draw a ball from one of two virtual urns: the urn containing more red than green balls if the value is zero; and the urn containing more green than red balls if the value is 100. But, exactly how many red and green balls are there in the urns? If the value is 0, then there are 70 red balls and 30 green balls. If the value is 100, then there are 70 green balls and 30 red balls.

10.3 Your per-round payoff

Your earnings depend on how well you predict the value of the good. If you are the first person in the sequence, your payoff will depend on the only prediction that you are asked to make. If you are a subsequent decision maker, your payoff will depend on the first or the second prediction you make, with the same chance (like in the toss of a coin).

If you predict the value exactly, you will earn 100 liras. If your prediction differs from the true value by an amount x , you will earn $100 - 0.01x^2$. This means that the further your prediction is from the true value, the less you will earn. Moreover, if your mistake is small, you will be penalized only a small amount; if your mistake is big, you will be penalized more than proportionally.

To make this rule clear, let us see some examples.

Example 1: Suppose the true value is 100. Suppose you predict 80. In this case you made a mistake of 20. We will give you $100 - 0.01 * 20^2 = 96.0$ liras.

Example 2: Suppose the true value is 0. Suppose you predict 10. In this case you made a mistake of 10. We will give you $100 - 0.01 * 10^2 = 99$ liras.

Example 3: Suppose the true value is 100. Suppose you predict 25. In this case you made a mistake of 75. We will give you $100 - 0.01 * 75^2 = 43.75$ liras.

Example 4: Suppose the true value is 0. Suppose you predict 50. In this case you made a mistake of 50. We will give you $100 - 0.01 * 50^2 = 75$ liras.

Note that the worst you can do under this payoff scheme is to state that you believe that there is a 100% chance that the value is 100 when in fact it is 0, or you believe that there is a 100% chance that the value is 0 when in fact it is 100. Here your payoff from prediction would be 0. Similarly, the best you can do is to guess correctly and assign 100% to the value which turns out to be the actual value of the good. Here your payoff will be 100 liras.

Note that with this payoff scheme, the best thing you can do to maximize the expected size of your payoff is simply to state your true belief about what you think the true value of the good is. Any other prediction will decrease the amount you can expect to earn. For instance, suppose you think there is a 90% chance that the value of the good is 100 and, hence, a 10% chance that value is 0. If this is your belief about the likely value of the good, to maximize your expected payoff, choose 90 as your prediction. Similarly, if you think the value is 100 with chance 33% and 0 with chance 67%, then select 33.

10.4 The other rounds

We will repeat the procedures described in the first round for 14 more rounds. As we said, at the beginning of each new round, the value of the good is again randomly chosen by the computer. Therefore, the value of the good in round 2 is independent of the value in round 1 and so on.

10.5 The final payment

To compute your payment, we will randomly choose (with equal chance) one round among the first five, one among the rounds 6 – 10 and one among the last five rounds. For each of these round we will then choose either prediction 1 or prediction 2 (with equal chance), unless you turn was 1, in which case there is nothing to choose since you only made one prediction. We will sum the payoffs that you have obtained for those predictions and rounds. We will then convert your payoff into pounds at the exchange rate of 100 liras = £7. That is, for every 100 liras you earn, you will get 7 pounds. Moreover, you will receive a participation fee of £5 just for showing up on time. You will be paid in cash, in private, at the end of the experiment.