

De Filippis, Roberta; Guarino, Antonio; Jehiel, Philippe; Kitagawa, Toru

Working Paper

Non-Bayesian updating in a social learning experiment

cemmap working paper, No. CWP60/20

Provided in Cooperation with:

Institute for Fiscal Studies (IFS), London

Suggested Citation: De Filippis, Roberta; Guarino, Antonio; Jehiel, Philippe; Kitagawa, Toru (2020) : Non-Bayesian updating in a social learning experiment, cemmap working paper, No. CWP60/20, Centre for Microdata Methods and Practice (cemmap), London, <https://doi.org/10.47004/wp.cem.2020.6020>

This Version is available at:

<https://hdl.handle.net/10419/241935>

Standard-Nutzungsbedingungen:

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

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

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

Terms of use:

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

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

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

Non-Bayesian updating in a social learning experiment

Roberta De Filippis
Antonio Guarino
Philippe Jehiel
Toru Kitagawa

The Institute for Fiscal Studies
Department of Economics, UCL

cemmap working paper CWP60/20

Non-Bayesian Updating in a Social Learning Experiment

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

Abstract

In our laboratory experiment, subjects, in sequence, have to predict the value of a good. The second subject in the sequence makes his prediction twice: first (“first belief”), after he observes his predecessor’s prediction; second (“posterior belief”), after he observes his private signal. We find that the second subjects weigh their signal as a Bayesian agent would do when the signal confirms their first belief; they overweight the signal when it contradicts their first belief. This way of updating, incompatible with Bayesianism, can be explained by the Likelihood Ratio Test Updating (LRTU) model, a generalization of the Maximum Likelihood Updating rule. It is at odds with another family of updating, the Full Bayesian Updating. In another experiment, we directly test the LRTU model and find support for it.

1 Introduction

Suppose you are contemplating the possibility of investing in a new project. Since it seems, *a priori*, equally likely that it succeeds or fails, you ask for the opinion of an independent advisor. After collecting some information on the project, he evaluates the probability of success to be 70%. On the basis of this recommendation only, clearly your belief on the probability of succeeding depends on how much you trust your advisor’s ability. If you fully trust him, you may agree with him and evaluate the probability of success to be 70% as well. If you do not think he has done a good job, or you suspect he is not so talented as you originally thought, you may even completely discard his view and keep

*De Filippis: Department of Economics, University College London (UCL) (e-mail: roberta.filippis.10@ucl.ac.uk); Guarino: Department of Economics, UCL (e-mail: a.guarino@ucl.ac.uk); Jehiel: Paris School of Economics, and Department of Economics, UCL (e-mail: jehiel@enpc.fr); Kitagawa: Department of Economics, UCL (e-mail: t.kitagawa@ucl.ac.uk). We thank Marco Cipriani, Olivier Compte, Pietro Ortoleva and participants in conferences and seminars at many institutions for useful comments and suggestions. We thank Hugo Freeman, Andrea Giacometti, Huilei Kang, Justin Franco Lam, Jeff Rowley and Teresa Steininger for excellent research assistance. We gratefully acknowledge the financial support of the ERC and of the INET.

your prior belief of a 50% probability of success. For the sake of the example, let us assume you trust him, although not completely, and assess that the project will succeed with probability 65%. You now receive further information on the project, independently of your advisor's information. The information is of the same quality as that received by your advisor, but is negative, that is, if you had to base your evaluation on this information only, you would update your prior belief to a value lower than 50% (actually, 30%, since the information is of the same quality). How would you now use this information to make inferences on the quality of the project? Would you change your mind on your advisor's ability? And how would you revise your 65% belief?

One way of reasoning is that the negative private information you receive (contradicting the advisor's view) makes it more doubtful that he used his information correctly. You should, therefore, discount the advisor's evaluation even more than before and mainly rely on your information. Since you now trust him less, your belief based on his advice only would be less than 65%. Moreover, you have now negative information, which pushes your belief further down. For instance, if you completely lost trust in your advisor's ability, you would now totally discard his view and on the basis of his advice only you would evaluate the probability of success to be 50%; considering also the negative content of your information, you would evaluate the probability of success to be 30%.

While this reasoning seems intuitive, one has to be careful in applying it, at least if one wants to respect standard Bayesian updating.¹ Indeed, a Bayesian agent, once expressed his belief of 65%, would simply update it by considering the probability of the received information conditional on the project being a success or a failure. To the extent that the information received by the advisor and yourself are independent conditional on the profitability of the project, the 65% probability summarizes all the relevant information in order to update the belief. Given the values we have used in our example, a Bayesian agent would never give a valuation of 30%; rather, he would update the 65% belief to 44.3%.²

Decision problems like the one we have just, informally, described are very common in many applications. Traders in financial markets typically try to infer information, say about asset fundamental values, from the order flow and also use their own private information. Companies make investment decisions using consultants and their own internal profitability studies. Editors of scientific journals make their editorial decisions on the basis of referee reports and of their own reading of the papers.

In this work we aim to study human behavior in a controlled experiment

¹We presented this work in many seminars and conferences, and played this game with the audience. It occurred frequently that some professional economists gave an answer in agreement with this reasoning (and with our theory, as we shall see) and were not convinced they were violating Bayesianism until they saw the proof.

²This percentage is the solution to the equation $\frac{x}{100-x} = \frac{65}{35} \frac{30}{70}$ (where we have expressed the updating through the likelihood ratio), or, using the familiar Bayesian updating formula, to:

$$\frac{x}{100} = \frac{0.3(0.65)}{0.3(0.65)+0.7(0.35)}.$$

in which subjects have to make similar types of decisions. 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, if appropriate, to understand the determinants of deviations from it. To be specific, 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. 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 value of the good and ask him to make a new prediction.

The main result of our study is that 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 (i.e., *as if* they had observed more than one signal, or *as if* the signal had a higher precision than it actually has).

This asymmetric updating is incompatible with standard Bayesianism, even allowing for subjective precisions of the signals, in the vein of Savage (1954), as long as the two pieces of information are independent conditional on the value of the good. Given that the conditional independence follows from principles of logic, we conclude that there is a need to go beyond the Bayes/Savage paradigm. Moreover, as we will demonstrate, the asymmetric updating: *i*) cannot be explained by known psychological biases such as base-rate neglect or confirmation bias; *ii*) is not due to risk preferences; *iii*) is inherently related to social learning (indeed, it does not occur in control, individual decision making, treatments).

Another main contribution of the paper is to provide a simple explanation for such asymmetric updating, by building on the literature on ambiguity and multiple priors (and, in particular, on the branch devoted to belief updating). We propose that, like in models with multiple priors, 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 between 0 and 100 independently of whether he receives a good or bad 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

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.

using subjective representations of the precision of the signals). We label this mode of updating Likelihood Ratio Test Updating model, a generalization of the Maximum Likelihood Updating model.

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. It is important to remark that it is as if the subject put different weights on the signal depending on its realization, not that he consciously uses different weights.

If instead of considering the Likelihood Ratio Test Updating model one were to consider the Full Bayesian Updating model, a popular alternative model of updating with multiple priors, we establish that the agent would behave as if putting a *lower* (rather than higher) weight on his own signal when it is contradicting relative to when it is confirming. Thus, whereas the Likelihood Ratio Test Updating model can explain our experimental finding, the Full Bayesian Updating model cannot.

More formally, we offer statistical and experimental evidence in support of this theory. We perform a careful econometric analysis and find evidence in support of this theory against the Bayesian (Savage) approach and the Full Bayesian Updating model. As for experimental evidence, in a treatment we directly test how human subjects update on the rationality of the predecessor. Again, we find that the updating is at odds with the Savage approach and the Full Bayesian Updating model and in line with our model.

Our findings are of direct relevance for theory purposes. In particular, to the best of our knowledge, we are the first to use multiple priors to model beliefs about rationality of others and provide experimental evidence in its support. We show that ambiguity is relevant not just when it is about the composition of an urn (which has inspired an enormous literature, following Ellsberg’s famous experiments), but also in very common and economically relevant decision problems. In our analysis, we also contribute to the debate on how to update multiple priors (an area in which there is very little theoretical consensus) and

propose a novel identification strategy to distinguish between different theories of updating.

It should be clear that our results have implications for applied work. In applications, the Savage approach has been widely used. For instance, in behavioral finance, researchers model overconfidence by allowing subjective precisions of the signals (an agent considers his own signal more informative than the others'). While this is left for future research, our finding suggests the need to allow for multiple priors, which may unveil important phenomena, e.g., in the way economic agents react to good and bad news depending on the prevailing current belief, market conditions or state of the economy (recession or boom).

The paper is organized as follows. Section 2 describes the model, the Perfect Bayesian Equilibrium predictions, as well the predictions of the Likelihood Ratio Test Updating model and of the Full Bayesian Updating model. Section 3 presents the experiment. Section 4 contains the results of the main treatments. Section 5 illustrates the control treatments. Sections 6 and 7 present econometric and experimental evidence in support of the Likelihood Ratio Test Updating rule. Section 8 discusses the related literature and Section 9 concludes. The Appendix contains additional material and the instructions.

2 The Theoretical Framework

2.1 The Basic Model

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 assumed to be risk neutral 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 symmetric binary signal $s_t \in \{0, 1\}$ whose objective probability distribution is:

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

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.

At time 1, agent 1 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 $j = 1, 2$ is equal to $100 - 0.01(V - a_t^j)^2$.⁴

2.1.1 Theoretical Predictions of the Basic Model

Given the quadratic payoff function, both agents optimally state their belief (expressed as a percentage) that the value is 100. For given information I_t^j , the agent chooses a_t^j to maximize his expected payoff $E[100 - 0.01(V - a_t^j)^2 \mid I_t^j]$,

⁴The superscript j is immaterial when $t = 1$ since agent 1 acts only once.

that is, he chooses $a_t^{j*} = E(V|I_t^j)$. In the Perfect Bayesian Equilibrium (PBE), agent 1 chooses 70 upon observing $s_1 = 1$ and 30 upon observing $s_1 = 0$. After observing $a_1^* = 30$ or 70, agent 2 chooses as his first action $a_2^{1*} = a_1^*$. After observing the private signal s_2 , the agent updates his belief and chooses $a_2^{2*} = E(V|a_1 = 70, s_2 = 1) = 84.48$, $a_2^{2*} = E(V|a_1 = 30, s_2 = 0) = 15.52$ or $a_2^{2*} = E(V|a_1 = 30, s_2 = 1) = E(V|a_1 = 70, s_2 = 0) = 50$. In the first two cases, the signal points in the same direction as the predecessor’s action, therefore, we refer to them saying that the signal is “confirming;” in the last two cases, instead, the signal points in the opposite direction to the predecessor’s action, and we say that the signal is “contradicting.” Note that when the signal is contradicting, the agent’s posterior belief is equal to the prior belief, since the two pieces of information the agent receives “cancel out.”

2.2 The Extended Model

In the experiment, we can expect to see actions a_1 and a_2^j different from those computed in the PBE above for a variety of reasons. We summarize these reasons by letting agents attach subjective precisions to the signals, possibly different from the objective ones, in the spirit of Savage. We also wish to consider the possibility that agent 2 may “distrust” agent 1 by letting him believe that agent 1 did not make a correct inference from his signal.

Specifically, we let agents have subjective views about the signal precisions: $q_1^S \in (0.5, 1]$ and $q_2^{tS} \in (0.5, 1]$, where q_1^S stands for the subjective precision attached by agent 1 to the signal at time 1 and q_2^{tS} stands for the subjective precision attached by agent 2 to signals at times $t = 1, 2$.⁵ We also let agent 2 have subjective beliefs on the “rationality” of agent 1, that is, to think that agent 1 may be of two types, rational (τ_r) or noise (τ_n). The subjective probability that the subject is noise is denoted by $\Pr(\tau = \tau_n) \equiv \theta$ and the subjective probability that he is rational by $\Pr(\tau = \tau_r) = 1 - \theta$ (later on we will also use $\mu \equiv 1 - \theta$). The two types of agents behave as follows. A rational agent always chooses an action greater than 50 upon observing a good signal and an action lower than 50 upon observing a bad signal. Specifically, we let $f(a_1)$ denote the subjective distribution of actions on $(50, 100]$ as perceived by agent 2 when agent 1 is rational and receives a signal $s_1 = 1$, and by symmetry we let $f(100 - a_1)$ be the corresponding subjective distribution of actions on $a_1 \in [0, 50)$ when agent 1 receives a signal $s_1 = 0$.⁶ A noise agent, by contrast, chooses in a random way, according to a distribution $g(a_1)$ symmetric around 50 independently of the observed signal.⁷ We are defining an agent at time 1 as “rational” as long as he updates in the correct direction, since the only thing that agent 2 has to learn from agent 1 is, indeed, the signal realization (given that the objective

⁵The superscript S stands for “subjective.” Note that we let agent 2 assign a different precision to signal 1 (that he does not observe) and to signal 2. Also note that we maintain the assumption of symmetric signals.

⁶Symmetry is not required for our arguments; it is only to simplify the exposition.

⁷We could have instead assumed that the non-rational agent always chooses an action in the wrong direction. This would not alter our results.

precisions of the signals are known and do not have to be learned), and this is revealed under the minimal requirement that agent 1 updates in the right direction. We assume that, after observing a_1 , agent 2 thinks that agent 1 is rational with probability $\mu_1(a_1)$. This probability can be derived from the prior subjective probability θ and the subjective distributions of actions at time 1 by a rational or a noise agent 1.⁸ Formally, for any $a_1 \in (50, 100]$, we have that

$$\mu(a_1) = \frac{(1-\theta)f(a_1)}{(1-\theta)f(a_1) + 2g(a_1)\theta}.$$

A similar expression can be derived for the case of $a_1 \in [0, 50)$.

Given the subjective precisions and rationality beliefs and given his information I_t^j , the agent chooses a_t^j to maximize his expected payoff $E[100 - 0.01(V - a_t^j)^2 | I_t^j]$: his optimal action is $a_t^{j*} = E(V | I_t^j)$.

2.2.1 Theoretical Predictions of the Extended Model

The predictions of the extended model are qualitatively similar to the basic model's. Since $q_1^S > 0.5$, agent 1 updates in the "correct direction" although not necessarily as in the PBE. Given $\mu_1(a_1)$ and q_2^{1S} , agent 2 computes his expected value of the good, which he expresses by stating a_2^1 . In particular, a_2^{1*} is such that

$$\frac{a_2^{1*}}{100 - a_2^{1*}} = \frac{\Pr(V = 100 | a_1)}{\Pr(V = 0 | a_1)} = \begin{cases} \frac{\mu_1(a_1)q_2^{1S} + (1-\mu_1(a_1))g(a_1)}{\mu_1(a_1)(1-q_2^{1S}) + (1-\mu_1(a_1))g(a_1)} & \text{if } a_1 > 50, \\ \frac{\mu_1(a_1)(1-q_2^{1S}) + (1-\mu_1(a_1))g(a_1)}{\mu_1(a_1)q_2^{1S} + (1-\mu_1(a_1))g(a_1)} & \text{if } a_1 < 50, \end{cases} \quad (1)$$

where we are using the law of total probabilities and interpret $\Pr(\cdot)$ as agent's subjective probability.

Once he has stated this belief, he then updates it, using his subjective precision q_2^{2S} . A key aspect that we wish to highlight is how the second action a_2^2 of agent 2 relates to his first action a_2^1 and to the signal s_2 he observes. Simple application of Bayes's law yields:

$$\frac{\Pr(V = 100 | a_1, s_2)}{\Pr(V = 0 | a_1, s_2)} = \frac{\Pr(s_2 | V = 100, a_1)}{\Pr(s_2 | V = 0, a_1)} \frac{\Pr(V = 100 | a_1)}{\Pr(V = 0 | a_1)}. \quad (2)$$

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

$$\frac{\Pr(V = 100 | a_1, s_2)}{\Pr(V = 0 | a_1, s_2)} = \frac{\Pr(s_2 | V = 100)}{\Pr(s_2 | V = 0)} \frac{\Pr(V = 100 | a_1)}{\Pr(V = 0 | a_1)}, \quad (3)$$

that is, to

$$\frac{a_2^{2*}}{100 - a_2^{2*}} = \left(\frac{q_2^{2S}}{1 - q_2^{2S}} \right)^{2s_2 - 1} \frac{a_2^{1*}}{100 - a_2^{1*}}. \quad (4)$$

⁸The subjective distribution of actions by a rational agent 1 is derived from the distribution of q_2^{1S} .

Our main focus is on the weight agent 2 attaches to his own signal and how this depends on the signal realization s_2 and on agent 1's action a_1 . To this end, let us define this weight as the value α_2^2 that satisfies the following equality:

$$\frac{a_2^{2*}}{100 - a_2^{2*}} = \left(\frac{q}{1 - q} \right)^{\alpha_2^2 \cdot (2s_2 - 1)} \frac{a_2^{1*}}{100 - a_2^{1*}}. \quad (5)$$

How α_2^2 compares to 1 indicates, for each (s_2, a_1) , whether agent 2 attaches more ($\alpha > 1$) or less ($\alpha < 1$) weight to his signal s_2 relative to the objective precision of s_2 ($q \equiv 0.7$). From (4) and (5), α_2^2 of the Bayesian agent is given by

$$\alpha_{2,B}^2(s_2, a_1; q_2^{2S}) \equiv \frac{\ln \left(\frac{q_2^{2S}}{1 - q_2^{2S}} \right)}{\ln \left(\frac{q}{1 - q} \right)}. \quad (6)$$

From this expression, it is clear that while $\alpha_{2,B}^2(s_2, a_1; q_2^{2S})$ may differ from 1 whenever $q_2^{2S} \neq q$, it is always the case that $\alpha_{2,B}^2(s_2, a_1; q_2^{2S})$ is the same, irrespective of a_1 and of s_2 . In particular, we have proven the following result:

Proposition 1 (Symmetric Updating)

In the PBE of the extended model, agent 2 attaches the same weight $\alpha_{2,B}^2(s_2, a_1; q_2^{2S}) = \ln \left(\frac{q_2^{2S}}{1 - q_2^{2S}} \right) / \ln \left(\frac{q}{1 - q} \right)$ to his private signal s_2 for either realization 0 or 1 and for any a_1 .

The signal realization s_2 affects how agent 2 updates his belief (this can be seen in the exponent $2s_2 - 1$ in (4)), but the weight attached to it as defined in (6) does not depend on the signal realization. In particular, it is the same whether agent 2 observes a signal that contradicts agent 1's action (e.g., $s_2 = 0$ after $a_1 > 50$) or a signal that confirms agent 1's action (e.g., $s_2 = 1$ after $a_1 > 50$). The fundamental reason for this "symmetric" updating is that, conditional on the value of the good, the signal at $t = 2$ is independent of the signal at $t = 1$ (by experimental design) as well as of the rationality of agent 1, which appears to derive from principles of logic.

So far we have considered risk neutral agents. The property of symmetric updating is, however, more general. Suppose agents have the following utility

$$u(V, a_t^j; \eta) = 100 \left(1 - \left(\frac{|V - a_t^j|}{100} \right)^{1 + \eta} \right), \quad (7)$$

where $\eta \in [0, \infty)$ is a risk preference parameter. For $\eta = 1$, $u(V, a_t; \eta)$ coincides with the monetary payoff, that is, this is the case of risk neutrality. When $\eta > 1$ the agent is risk averse, whereas a value $\eta < 1$ implies risk loving preferences.

By following the same steps as above, we obtain the relationship between the first and second optimal actions of agent 2, a generalization of (4):

$$\frac{a_2^{2*}}{1 - a_2^{2*}} = \left(\frac{q_2^{2S}}{1 - q_2^{2S}} \right)^{\eta^{-1}(2s_2 - 1)} \frac{a_2^{1*}}{1 - a_2^{1*}}. \quad (8)$$

In this extended model that allows for different risk preferences, by using (4) and (8), we obtain

$$\alpha_{2,B}^2(s_2, a_1; q_2^{2S}, \eta) = \eta^{-1} \frac{\ln\left(\frac{q_2^{2S}}{1-q_2^{2S}}\right)}{\ln\left(\frac{q}{1-q}\right)}. \quad (9)$$

We have therefore proven the following proposition:

Proposition 2 (Symmetric Updating for Different Risk Preferences)

Suppose agent 2 has preferences represented by (7). In the PBE of the extended model with risk preferences, he attaches the same weight $\alpha_{2,B}^2(s_2, a_1; q_2^{2S}, \eta) \equiv \eta^{-1} \ln\left(\frac{q_2^{2S}}{1-q_2^{2S}}\right) / \ln\left(\frac{q}{1-q}\right)$ to his private signal s_2 for either realization 0 or 1 and for any a_1 .

Intuitively, the effect of risk aversion is to lower the weight attached to the signal, as it makes actions closer to 50 more attractive (compared to the case of risk neutrality). This, however, does not affect the conclusion about the symmetry in updating, since it holds no matter what the signal realization is.

While we have established that the symmetric updating result holds for a broad class of extensions that adhere to the principle of Bayesian updating, it may be worth mentioning an incorrect intuition that would suggest otherwise. After observing a contradicting signal (e.g., $s_2 = 0$ after $a_1 > 50$), an agent would seem to be right in updating down his belief on the rationality of the predecessor (the belief previously used when stating a_2^{1*}) and, as a result, in inferring more from his own private signal, compared to the case of a confirming signal (e.g., $s_2 = 1$ after $a_1 > 50$). Yet, with Bayesian updating, while it is true that agent 2 does update on the rationality of the predecessor after observing his private signal, his posterior belief on the value of the good follows expression (4), as results from the statistical property that conditional on the value of the good, the signal s_2 is independent of the action a_1 . In other words, the intuition is incorrect because it implicitly amounts to incorrectly applying Bayes's law.⁹

In summary, we have shown that even enriching the model with risk aversion, subjective precisions and subjective expectations on the rationality of the

⁹To see this, note that a Bayesian agent updates on the rationality of the predecessor after observing the action a_1 and receiving the signal s_2 as follows:

$$\frac{\mu^2(a_1, s_2)}{1 - \mu^2(a_1, s_2)} = \frac{\sum_{k=0}^1 \Pr(\tau_r | a_1, s_2, V = k) \Pr(V = k | a_1, s_2)}{\sum_{k=0}^1 \Pr(\tau_n | a_1, s_2, V = k) \Pr(V = k | a_1, s_2)} = \frac{\sum_{k=0}^1 \Pr(\tau_r | a_1, V = k) \Pr(V = k | a_1, s_2)}{\sum_{k=0}^1 \Pr(\tau_n | a_1, V = k) \Pr(V = k | a_1, s_2)}. \quad (10)$$

If instead agent 2 after observing a_1 and s_2 first updated on the rationality of agent 1 and then plugged $\mu^2(a_1, s_2)$ into (1) to form a belief on V as in (4), he would violate Bayesian updating (as he would make the mistake of using the posterior belief —after observing a_1 and s_2 — about the rationality of agent 1 as if it were the belief after observing a_1 only).

predecessor, the Bayesian model dictates that the weight put on the signal does not depend on the signal realization. As we shall see, this basic implication of Bayesian updating is not consistent with our experimental data.

2.3 Updating Multiple Priors

In the Introduction we mentioned the (incorrect) 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 own signal. While this intuition is incorrect in the realm of Bayesianism, it 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, update on the first belief. We first review this method of updating referred to as the Likelihood Ratio Test Updating rule, and then discuss an alternative popular method of updating referred to as Full Bayesian Updating, which will be shown to imply markedly different predictions.

2.3.1 Likelihood Ratio Test Updating

Compared to the previous Extended Model, we now introduce one modification: we allow agent 2 to have multiple priors about $\theta = \Pr(\tau = \tau_n)$. Whereas a Bayesian agent has a unique prior θ , a subject at time 2 has now ambiguous beliefs on θ , that is, multiple priors belonging to the set $[\theta_*, \theta^*] \subseteq [0, 1]$, $\theta_* < \theta^*$.

To update his belief upon observing an event E , first the agent 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 θ_* , in the complementary case, he will pick up θ^* . In other words,

$$\begin{aligned} \text{if } \frac{\Pr(E|\tau = \tau_r)}{\Pr(E|\tau = \tau_n)} &\geq c, \text{ then } \theta = \theta_*, \text{ and} & (11) \\ \text{if } \frac{\Pr(E|\tau = \tau_r)}{\Pr(E|\tau = \tau_n)} &< c, \text{ then } \theta = \theta^*, \end{aligned}$$

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

We refer to this model of updating based on the likelihood ratio $\frac{\Pr(E|\tau=\tau_r)}{\Pr(E|\tau=\tau_n)}$ as Likelihood Ratio Test Updating (LRTU) rule. It can be seen as a simple generalization of the Maximum Likelihood Updating (MLU) model, in which the time 2 subject estimates θ to be the value in $[\theta_*, \theta^*]$ that maximizes the likelihood of observing the event E . Indeed, since

$$\Pr(E) = \Pr(E|\tau = \tau_r)\Pr(\tau = \tau_r) + \Pr(E|\tau = \tau_n)\Pr(\tau = \tau_n),$$

that is,

$$\Pr(E) = \Pr(E|\tau = \tau_r)(1 - \theta) + \Pr(E|\tau = \tau_n)\theta,$$

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

$$\begin{aligned} \text{if } \frac{\Pr(E|\tau = \tau_r)}{\Pr(E|\tau = \tau_n)} &\geq 1, \text{ then } \theta = \theta_*, \text{ and} & (12) \\ \text{if } \frac{\Pr(E|\tau = \tau_r)}{\Pr(E|\tau = \tau_n)} &< 1, \text{ then } \theta = \theta^*. \end{aligned}$$

The LRTU model generalizes the MLU model to take into account that agents may need stronger or weaker evidence in favor of one type in order to select a specific prior.¹⁰

Let us now describe the LRTU in the context of our model. In our model subject 2 makes a decision twice, first after observing the event $E \equiv \{a_1\}$ and then after observing the event $E \equiv \{a_1, s_2\}$.¹¹ Thus, there are two events in which agent 2 updates his belief about the rationality of agent 1. Note also that after observing $\{a_1, s_2\}$ the agent also uses the signal realization s_2 to update on the first belief.

From now on we consider the case in which $a_1 \in (50, 100]$; the other case is analogous. In this case, for $E \equiv \{a_1\}$, $\frac{\Pr(E|\tau = \tau_r)}{\Pr(E|\tau = \tau_n)} = \frac{f(a_1)}{2g(a_1)}$; hence,

$$\theta^1(a_1) = \begin{cases} \theta_* & \text{if } \frac{f(a_1)}{2g(a_1)} \geq c, \\ \theta^* & \text{if } \frac{f(a_1)}{2g(a_1)} < c. \end{cases}$$

Suppose now that agent 2 receives a confirming signal, $s_2 = 1$, then $E = \{a_1, s_2 = 1\}$ and

$$\frac{\Pr(E|\tau = \tau_r)}{\Pr(E|\tau = \tau_n)} = \frac{[q_2^{1S} q_2^{2S} + (1 - q_2^{1S})(1 - q_2^{2S})]f(a_1)}{g(a_1)}.$$

¹⁰This is equivalent to assuming that the subject acts as if he received another signal φ about the predecessor's type (and uncorrelated with the event). In this case, he would choose the prior to maximize the following probability:

$$\Pr(E, \varphi) = \Pr(E, \varphi|\tau = \tau_r) \Pr(\tau = \tau_r) + \Pr(E, \varphi|\tau = \tau_n) \Pr(\tau = \tau_n).$$

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

$$\frac{\Pr(E, \varphi|\tau = \tau_r)}{\Pr(E, \varphi|\tau = \tau_n)} \geq 1,$$

that is,

$$\frac{\Pr(E|\tau = \tau_r) \Pr(\varphi|\tau = \tau_r)}{\Pr(E|\tau = \tau_n) \Pr(\varphi|\tau = \tau_n)} \geq 1,$$

or

$$\frac{\Pr(E|\tau = \tau_r)}{\Pr(E|\tau = \tau_n)} \geq \frac{\Pr(\varphi|\tau = \tau_n)}{\Pr(\varphi|\tau = \tau_r)}.$$

By setting $\frac{\Pr(\varphi|\tau = \tau_n)}{\Pr(\varphi|\tau = \tau_r)} \equiv c$, one obtains the LRTU model.

¹¹Since a_1 is a continuous variable, $\Pr(\{a_1\}|T = t_r)$ should be read as a conditional density function.

If, instead, agent 2 receives a contradicting signal, $s_2 = 0$, then $E = \{a_1, s_2 = 0\}$ and

$$\frac{\Pr(E|\tau = \tau_r)}{\Pr(E|\tau = \tau_n)} = \frac{[q_2^{1S}(1 - q_2^{2S}) + (1 - q_2^{1S})q_2^{2S}]f(a_1)}{g(a_1)}.$$

Note that $[q_2^{1S}q_2^{2S} + (1 - q_2^{1S})(1 - q_2^{2S})] > \frac{1}{2} > [q_2^{1S}(1 - q_2^{2S}) + (1 - q_2^{1S})q_2^{2S}]$: a confirming signal makes it more likely that agent 1 is rational; conversely, a contradicting signal makes it less likely.¹²

Therefore, we have the following proposition:

Proposition 3 (Belief on the Predecessor's Rationality with LRTU Rule)

Suppose $\frac{f(a_1)}{2g(a_1)} \geq c$ so that $\theta^1(a_1) = \theta_*$. Then

$$\theta^2(a_1, s_2 = 1) = \theta_*$$

and

$$\theta^2(a_1, s_2 = 0) = \begin{cases} \theta_* & \text{if } [q_2^{1S}(1 - q_2^{2S}) + (1 - q_2^{1S})q_2^{2S}] \frac{f(a_1)}{g(a_1)} \geq c, \\ \theta^* & \text{if } [q_2^{1S}(1 - q_2^{2S}) + (1 - q_2^{1S})q_2^{2S}] \frac{f(a_1)}{g(a_1)} < c. \end{cases}$$

Suppose $\frac{f(a_1)}{2g(a_1)} < c$ so that $\theta^1(a_1) = \theta^*$. Then

$$\theta^2(a_1, s_2 = 0) = \theta^*$$

and

$$\theta^2(a_1, s_2 = 1) = \begin{cases} \theta_* & \text{if } [q_2^{1S}q_2^{2S} + (1 - q_2^{1S})(1 - q_2^{2S})] \frac{f(a_1)}{g(a_1)} \geq c, \\ \theta^* & \text{if } [q_2^{1S}q_2^{2S} + (1 - q_2^{1S})(1 - q_2^{2S})] \frac{f(a_1)}{g(a_1)} < c. \end{cases}$$

Suppose, as it will prove more relevant to interpret our data, that agent 2 initially trusts that agent 1 is more likely to be rational and thus chooses θ_* . Then Proposition 3 says that agent 2 must *a fortiori* do so when receiving a confirming signal, since the likelihood that agent 1 is rational increases. This is not necessarily true after a contradicting signal, since the likelihood that agent 1 is rational decreases. Thus, when agent 2 observes a confirming signal, he keeps the same belief θ_* about agent 1 being a noise agent, and, therefore, his two choices (one before and one after his own signal) are made as if the agent 2 held a unique prior θ_* , that is, as if he were Bayesian. In contrast, when a contradicting signal is observed and agent 2 switches from θ_* to θ^* , his two

¹²A simple proof of the inequality proceeds as follows. Observe that $q_1^S \rightarrow q_1^S q_2^S + (1 - q_1^S)(1 - q_2^S)$ is a strictly increasing function for all $q_2^S \in (0.5, 1]$ and thus in the considered range, $q_1^S q_2^S + (1 - q_1^S)(1 - q_2^S)$ is strictly larger than $q_2^S \frac{1}{2} + (1 - \frac{1}{2})(1 - q_2^S) = \frac{1}{2}$.

Similarly, $q_1^S \rightarrow q_1^S (1 - q_2^S) + (1 - q_1^S)q_2^S$ is strictly decreasing for all $q_2^S \in (0.5, 1]$ and thus $q_1^S (1 - q_2^S) + (1 - q_1^S)q_2^S < \frac{1}{2}$.

choices are made as follows: first, his belief on the predecessor's rationality is computed as

$$\mu(a_1) = \frac{(1 - \theta_*)f(a_1)}{(1 - \theta_*)f(a_1) + 2\theta_*g(a_1)},$$

and he makes the first action $a_{2,L}^1$ such that

$$\frac{a_{2,L}^1}{100 - a_{2,L}^1} = \frac{\mu(a_1)q_2^{1S} + (1 - \mu(a_1))g(a_1)}{\mu(a_1)(1 - q_2^{1S}) + (1 - \mu(a_1))g(a_1)};$$

then, after observing the signal $s_2 = 0$, his belief on the predecessor's rationality is

$$\mu^2(a_1, s_2 = 0) = \frac{(1 - \theta^*)f(a_1)}{(1 - \theta^*)f(a_1) + 2\theta^*g(a_1)}.$$

and he makes the second action $a_{2,L}^{2*}$ such that

$$\frac{a_{2,L}^{2*}}{100 - a_{2,L}^{2*}} = \left(\frac{q_2^{2S}}{1 - q_2^{2S}} \right)^{2s_2 - 1} \frac{a_2^{1**}}{100 - a_2^{1**}}, \quad (13)$$

where

$$\frac{a_2^{1**}}{100 - a_2^{1**}} = \frac{\mu^2(a_1, s_2 = 0)q_2^{1S} + (1 - \mu^2(a_1, s_2 = 0))g(a_1)}{\mu^2(a_1, s_2 = 0)(1 - q_2^{1S}) + (1 - \mu^2(a_1, s_2 = 0))g(a_1)} < \frac{a_{2,L}^1}{100 - a_{2,L}^1}.$$

Similarly to the Bayesian model, let $\alpha_{2,L}^2(s_2, a_1; q_2^{1S}, q_2^{2S})$ be the weight attached to s_2 as defined by

$$\frac{a_{2,L}^2}{100 - a_{2,L}^2} = \left(\frac{q}{1 - q} \right)^{\alpha_2^2 \cdot (2s_2 - 1)} \frac{a_{2,L}^1}{100 - a_{2,L}^1},$$

where agent 2's action is now obtained using the LRTU rule. When $\theta(a_1) = \theta_*$ and the signal is confirming ($s_2 = 1$), it is immediate that $\alpha_{2,L}^2(s_2 = 1, a_1; q_2^{1S}, q_2^{2S}) = \ln\left(\frac{q_2^{2S}}{1 - q_2^{2S}}\right) / \ln\left(\frac{q}{1 - q}\right)$, like in the Bayesian case (since in this case there is no change of prior belief θ_*). When, instead, in response to a contradicting signal, agent 2's belief on the predecessor's rationality switches from θ_* to θ^* we obtain

$$\begin{aligned} \alpha_{2,L}^2(s_2 = 0, a_1; q_2^{1S}, q_2^{2S}) &= \left[\frac{\ln\left(\frac{q_2^{2S}}{1 - q_2^{2S}}\right)}{\ln\left(\frac{q}{1 - q}\right)} \right] \cdot \left(1 + \frac{\ln\left(\frac{a_{2,L}^{1*}}{100 - a_{2,L}^{1*}}\right) - \ln\left(\frac{a_2^{1**}}{100 - a_2^{1**}}\right)}{\ln\left(\frac{q_2^{2S}}{1 - q_2^{2S}}\right)} \right) \\ &> \ln\left(\frac{q_2^{2S}}{1 - q_2^{2S}}\right) / \ln\left(\frac{q}{1 - q}\right) = \alpha_{2,L}^2(s_2 = 1, a_1; q_2^{1S}, q_2^{2S}). \end{aligned}$$

The greater distrust of agent 1's rationality after a contradicting signal is equivalent to a larger weight attached by agent 2 to his own signal. Agent 2's weights

$\alpha_{2,L}^2(s_2, a_1; q_2^{1S}, q_2^{2S})$ are therefore no longer necessarily equal for the confirming and contradicting signal cases.

A similar analysis holds for the case in which $\frac{f(a_1)}{2g(a_1)} < c$, so that $\theta^1(a_1) = \theta^*$. Agent 2 can revise up the trust about the predecessor's rationality in response to a confirming signal $s_2 = 1$, while remaining distrustful if the signal is contradicting. We summarize these results in the next proposition:

Proposition 4 (Asymmetric Updating with LRTU Rule)

Under the LRTU updating rule, the weight that agent 2 attaches to his signal satisfies the following conditions:

(i) if $\frac{f(a_1)}{2g(a_1)} \geq c$ so that $\theta^1(a_1) = \theta_*$,

$$\alpha_{2,L}^2(s_2 = 0, a_1; q_2^{1S}, q_2^{2S}) \geq \alpha_{2,L}^2(s_2 = 1, a_1; q_2^{1S}, q_2^{2S}) = \ln \left(\frac{q_2^{2S}}{1 - q_2^{2S}} \right) / \ln \left(\frac{q}{1 - q} \right),$$

where the inequality is strict if and only if $[q_2^{1S}(1 - q_2^{2S}) + (1 - q_2^{1S})q_2^{2S}] \frac{f(a_1)}{g(a_1)} < c$;

(ii) if $\frac{f(a_1)}{2g(a_1)} < c$ so that $\theta^1(a_1) = \theta^*$,

$$\alpha_{2,L}^2(s_2 = 1, a_1; q_2^{1S}, q_2^{2S}) \geq \alpha_{2,L}^2(s_2 = 0, a_1; q_2^{1S}, q_2^{2S}) = \ln \left(\frac{q_2^{2S}}{1 - q_2^{2S}} \right) / \ln \left(\frac{q}{1 - q} \right),$$

where the inequality is strict if and only if $[q_2^{1S}q_2^{2S} + (1 - q_2^{1S})(1 - q_2^{2S})] \frac{f(a_1)}{g(a_1)} \geq c$.

We now illustrate our results on the LRTU rule with a simple example in which agent 2 first trusts the predecessor's rationality (i.e., $\frac{f(a_1)}{2g(a_1)} \geq c$).

An Example Suppose that agent 2 has multiple priors $[\theta_*, \theta^*] = [0, 1]$ on the predecessor's type. He observes $a_1 = 70$ and then the signal $s_2 = 0$. Let us assume the threshold is $c = 1$, so that the model is equivalent to the MLU model. Suppose he has expectations on the rational and noise types' actions at time 1 such that $\frac{\Pr(a_1=70|\tau=\tau_r)}{\Pr(a_1=70|\tau=\tau_n)} \geq 1$. In this case, the agent selects the prior $\theta_* = 0$. The agent is confident about the predecessor's rationality, and, therefore, chooses $a_{2,L}^1 = 70$. After receiving the signal $s_2 = 0$, the agent 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, $\frac{\Pr(a_1=70, s_2=0|\tau=\tau_r)}{\Pr(a_1=70, s_2=0|\tau=\tau_n)} < 1$, then the agent chooses $\theta^* = 1$. Being now confident that the predecessor was a noise type, the agent 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 belief of 0.5, the subject obtains a posterior belief of 0.3 on the value being 100 and, as a result, chooses $a_{2,L}^2 = 30$. In terms of our previous analysis, this is equivalent to a subject overweighting the signal, with $\alpha_2^2 = 2$, since $30 = 100 \frac{(1-0.7)^2 \frac{70}{100}}{(1-0.7)^2 \frac{70}{100} + 0.7^2 (1 - \frac{70}{100})}$. A similar analysis applies to the case in which the subject observes a signal $s_2 = 1$. It is easy to see that if

$\frac{\Pr(a_1=70|\tau=\tau_r)}{\Pr(a_1=70|\tau=\tau_n)} \geq 1$, then *a fortiori* $\frac{\Pr(a_1=70, s_2=1|\tau=\tau_r)}{\Pr(a_1=70, s_2=1|\tau=\tau_n)} \geq 1$ (as per Proposition 3). Therefore, in this case the subject sticks to the prior $\theta_* = 0$. Since the agent is still confident that the predecessor was rational, he does not change his first belief on $V = 100$, which remains 0.7. Since the agent 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 $a_{2,L}^2 = 84$. This is equivalent to an agent weighing the signal as a Bayesian agent would do, with $\alpha_2^2 = 1$.

2.3.2 Full Bayesian Updating

The LRTU (MLU) is only one family of updating rules for multiple priors. Another, well known, family goes under the name of Full Bayesian Updating (FBU) (for a discussion, see Gilboa and Marinacci, 2013). We now show that this updating rule does not produce the same asymmetry we illustrated for LRTU (MLU), and is not compatible with the intuition mentioned in the Introduction.

Here we consider the FBU model with maxmin preferences, axiomatized by Pires (2002). From now on, our statements about the FBU model should be read as statements about the FBU rule with maxmin preferences. In the Appendix, we will consider FBU with Hurwitz preferences as an alternative to the maxmin preferences (and we will show that this model is not compatible with the asymmetry explained before either).

Let us consider, as usual, the case in which $a_1 \in (50, 100]$. Suppose, similarly to the previous section, that an agent at time 2 starts with a set of priors $[\theta_*, \theta^*]$ on the proportion of noise-type agents at time 1. Upon observing the predecessor's action a_1 , the agent applies Bayes's rule to each prior θ in $[\theta_*, \theta^*]$ and obtains the following likelihood ratio:

$$\frac{p_2^1(\theta, q_2^{1S})}{1 - p_2^1(\theta, q_2^{1S})} \equiv \frac{\Pr(V = 100|a_1; \theta, q_2^{1S})}{\Pr(V = 0|a_1; \theta, q_2^{1S})} = \frac{(1 - \theta)q_2^{1S}f(a_1) + g(a_1)\theta}{(1 - \theta)(1 - q_2^{1S})f(a_1) + g(a_1)\theta}. \quad (14)$$

According to the FBU, the range of θ spans a range of beliefs on the value of the good being 100: $[p_2^1(\theta^*, q_2^{1S}), p_2^1(\theta_*, q_2^{1S})]$.

For any given prior θ in $[\theta_*, \theta^*]$, after observing a signal s_2 the agent updates his likelihood ratio as follows:

$$\frac{p_2^2(\theta, q_2^{1S}, q_2^{2S})}{1 - p_2^2(\theta, q_2^{1S}, q_2^{2S})} = \frac{(1 - \theta)q_2^{1S}f(a_1) + g(a_1)\theta}{(1 - \theta)(1 - q_2^{1S})f(a_1) + g(a_1)\theta} \left(\frac{q_2^{2S}}{1 - q_2^{2S}} \right)^{2s_2 - 1}.$$

Again, the range of θ spans a range of beliefs on the value of the good being 100: $[p_2^2(\theta^*, q_2^{1S}, q_2^{2S}), p_2^2(\theta_*, q_2^{1S}, q_2^{2S})]$.

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

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

that is,

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

Therefore, for the case $a_1 \geq 50$, since the lower bound of the beliefs satisfies $\underline{p} = p_2^1(\theta^*, q_2^{1S}) \geq \frac{1}{2}$, the agent's first action is based on the most pessimistic prior, $\theta = \theta^*$:

$$a_{2,F}^1 = 100p_2^1(\theta^*, q_2^{1S}).$$

This action equal to that of the Bayesian model with $\theta = \theta^*$. The agent acts as if he distrusted the predecessor's rationality (a pessimistic agent maximizes the minimum payoff he can obtain, and the minimum payoff coincides with the predecessor's action being the least informative).

Similarly, the second action is

$$a_{2,F}^2 = \begin{cases} 100p_2^2(\theta^*, q_2^{1S}, q_2^{2S}), & \text{if } p_2^2(\theta^*, q_2^{1S}, q_2^{2S}) \geq \frac{1}{2}, \\ 50, & \text{if } p_2^2(\theta^*, q_2^{1S}, q_2^{2S}) < \frac{1}{2}, \text{ and } p_2^2(\theta_*, q_2^{1S}, q_2^{2S}) > \frac{1}{2}, \\ 100p_2^2(\theta_*, q_2^{1S}, q_2^{2S}), & \text{if } p_2^2(\theta_*, q_2^{1S}, q_2^{2S}) \leq \frac{1}{2}. \end{cases}$$

After receiving a confirming signal, the agent keeps choosing according to the Bayesian model with the same prior $\theta = \theta^*$. Instead, after receiving a contradicting signal he chooses as in the Bayesian model but with $\theta \leq \theta^*$, since now the worst-case scenario is that in which the predecessor's action was the most informative (if $p_2^2(\theta_*, q_2^{1S}, q_2^{2S}) \leq \frac{1}{2}$ he chooses as in the Bayesian model with $\theta = \theta_*$).

It is useful to rewrite these conclusions in terms of weight on the private signal, and contrast them to those we obtained for the LRTU model (as shown in Proposition 4).

Let $\alpha_{2,F}^2(s_2, a_1; q_2^{1S}, q_2^{2S})$ be the weight attached to s_2 as defined by

$$\frac{a_{2,F}^2}{100 - a_{2,F}^2} = \left(\frac{q}{1-q} \right)^{\alpha_{2,F}^2 \cdot (2s_2 - 1)} \frac{a_{2,F}^1}{100 - a_{2,F}^1},$$

where agent 2's action is now obtained using the FBU rule. When agent 2 receives a confirming signal, $\alpha_{2,F}^2(s_2 = 1, a_1; q_2^{1S}, q_2^{2S}) = \ln \left(\frac{q_2^{2S}}{1 - q_2^{2S}} \right) / \ln \left(\frac{q}{1-q} \right)$ since the two actions by agent 2 (before and after receiving the signal) are based on the same prior belief $\theta = \theta^*$ (note though that this is not necessarily the same prior belief as in the LRTU model). When, instead, agent 2 receives a contradicting signal $s_2 = 0$, $\alpha_{2,F}^2(s_2 = 0, a_1; q_2^{1S}, q_2^{2S})$ is *smaller* than in the Bayesian case:

$$\begin{aligned} \alpha_{2,F}^2(s_2 = 0, a_1; q_2^{1S}, q_2^{2S}) &= \left[\frac{\ln \left(\frac{q_2^{2S}}{1 - q_2^{2S}} \right)}{\ln \left(\frac{q}{1-q} \right)} \right] \cdot \left(1 - \frac{\ln \left(\frac{a_{2,F}^2}{100 - a_{2,F}^2} \right) - \ln \left(\frac{a_{2,F}^1}{100 - a_{2,F}^1} \right)}{\ln \left(\frac{q_2^{2S}}{1 - q_2^{2S}} \right)} \right) \\ &\leq \ln \left(\frac{q_2^{2S}}{1 - q_2^{2S}} \right) / \ln \left(\frac{q}{1-q} \right) = \alpha_{2,F}^2(s_2 = 1, a_1; q_2^{1S}, q_2^{2S}), \end{aligned}$$

where a_2^{2*} is as shown in the expression (13) above. Here, the inequality follows from $a_{2,F}^2 \geq a_2^{2*}$ and is strict when $p_2^2(\theta^*, q_2^{1S}, q_2^{2S}) < \frac{1}{2}$. Hence, the FBU model yields an asymmetric way of updating, as summarized in the next proposition.

Proposition 5 (Asymmetric Updating with FBU Rule)

Under the FBU updating rule, the weight agent 2 attaches to his private signal depends on its realization and satisfies

$$\ln\left(\frac{q_2^{2S}}{1 - q_2^{2S}}\right) / \ln\left(\frac{q}{1 - q}\right) = \alpha_{2,F}^2(s_2 = 1, a_1; q_2^{1S}, q_2^{2S}) \geq \alpha_{2,F}^2(s_2 = 0, a_1; q_2^{1S}, q_2^{2S}),$$

where the inequality is strict if and only if $p_2^2(\theta^*, q_2^{1S}, q_2^{2S}) < \frac{1}{2}$.

Essentially, in the FBU model the agent updates like a Bayesian with a prior belief θ^* when receiving a confirming signal, while he puts a *lower* weight on his own signal when the signal is contradicting. This is in stark contrast with the LRTU model (case *i* of Proposition 4), in which the agent updates like a Bayesian with prior belief θ_* when receiving a confirming signal, while he puts a *higher* weight on his own signal when it is contradicting.

The essential observation from this analysis, which we will use for the interpretation of our experimental data, is that neither the Bayesian model nor the FBU model can explain that the weight on the own signal is greater upon receiving a contradicting signal than upon receiving a confirming signal, whereas the LRTU model can. For theoretical purposes, one may further observe that there is a similarity between the weights in the FBU model and in the LRTU model when the agent originally distrusts the predecessor's rationality (case *ii* of Proposition 4). Indeed, for both models the weight on the contradicting signal is lower than on the confirming signal. But note the difference in this case too. Take the case in which $q_2^{2S} = q$. In the FBU model, the agent updates like a Bayesian (he puts a weight 1 on the own signal) for the confirming case, while he puts a weight lower than 1 for the contradicting case. In the case *ii* of Proposition 4 for the LRTU rule, instead, the agent, after distrusting the predecessor, puts a weight higher than 1 on a confirming signal and equal to 1 on a contradicting signal.

We now illustrate the working of FBU in our model with a simple example.

An Example Suppose that agent 2 has multiple priors $[\theta_*, \theta^*] = [0, 1]$ on the predecessor's type. He observes $a_1 = 70$. The agent updates his belief on the value of the good using each prior $\theta \in [0, 1]$. This means that his posterior beliefs on $V = 100$ lie in $[0.5, 0.7]$. Therefore, he chooses $a_{2,F}^1 = 50$, the action that maximizes the minimum payoff he can obtain. After receiving the contradicting signal $s_2 = 0$, the subject updates his set of beliefs to $[0.3, 0.5]$. This implies that again he chooses $a_{2,F}^2 = 50$, which is equivalent to $\alpha_2^2 = 0$. After receiving the signal $s_2 = 1$, instead, the agent updates his set of beliefs to $[0.7, 0.84]$. He will then maximize his utility by choosing $a_{2,F}^2 = 70$, which is equivalent to $\alpha_2^2 = 1$. Note that in this example this updating rule implies no updating at all

(rather than overweighting the signal) after receiving a contradicting signal, and updating as a Bayesian after observing a confirming signal. This asymmetric way of updating is sharply different from that in the previous example (about the LRTU rule).

3 The Experiment and the Experimental Design

3.1 The Experiment

We ran the experiment in the Experimental Laboratory for Finance and Economics (ELFE) at the Department of Economics at University College London (UCL) in 2009, 2010, 2011, 2014, 2018 and 2019.¹³ 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 287 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:

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. 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 the good signal, $s_t = 1$, and the red signal to the bad signal, $s_t = 0$.
3. 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 places. The other subjects observed the decision made by the first subject on their screen. The identity of the subject was not revealed.
4. In a second period, another subject was randomly chosen and asked to choose a number between 0 and 100, having observed the first subject’s choice only.

¹³ELFE is the new name of the laboratory, formerly known as ELSE Laboratory. The SL treatments (described below) were conducted with a built-on-purpose program. The IDM and GC treatments were programmed and conducted with the software *z-Tree* (Fischbacher, 2007) in 2014 and 2019.

5. After he had made that choice, he received a signal (another “ball draw” from the same urn with replacement) 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.
6. The experiment then continued with a third, fourth, ..., tenth period, until all 10 subjects had acted.¹⁴
7. 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 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.¹⁵

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 rounds 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 (SL). 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 $[0.7, 0.71]$ (instead of having a precision always exactly equal to 0.7). Each subject observed not only the signal (the “ball color”) but also the exact precision of his own signal.¹⁶ A third treatment (SL3) was identical to SL2, with the

¹⁴The experiment was designed to address many research questions. The data collected on periods beyond 2 are not relevant for this paper’s research question.

¹⁵This 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 Schotter and Trevino (2014).

¹⁶Drawing 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 (since it becomes relevant only for periods beyond the second).

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.

Individual Decision Making (IDM). In the social learning treatments subjects make decisions after observing private signals and the actions of others. 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 two treatments in which subjects observed more than one signal and made more than one decision. In a first treatment (IDM1), a subject received a signal and had to make a choice in the interval $[0, 100]$ (as subject 1 in the SL treatments). Then, with a 50% probability, he received another signal and had to make a second decision (like the second decision of subject 2 in the SL treatments).¹⁷ In a second treatment (IDM2) we wanted to take care of another concern. In social learning, subject 2 observes a signal (subject 1's action) about another signal (subject 1's ball color): could it be that the behavior we observe is due to this sort of compound lottery? In the IDM2 treatment, the computer drew a first signal. The subject did not observe it; rather, he observed a computer's "announcement" about this signal. Such an announcement coincides with the true signal with probability 0.75 and is equal to the other signal with probability 0.25 (it is important to remark that these probabilities are known to the subject). After observing this announcement, the subject had to make a choice in the interval $[0, 100]$. Then, he received another signal and had to make a second decision (similarly to the SL treatments). The announcement is essentially a signal with precision 0.6 ($= 0.7 * 0.75 + 0.3 * 0.25$). Following the same logic of our basic model, in the PBE the first action is 0.6 or 0.4 (for a good or bad signal, respectively). The second action is 0.22 after a bad announcement and bad signal; 0.39 after a good announcement and a bad signal; 0.61 after a bad announcement and a good signal; 0.78 after a good announcement and good signal. Allowing for risk aversion and subjective precisions can change the specific predictions but does not alter the result for the symmetric updating, as we proved in Proposition 2.

Guess the Color (GC). In the social learning treatments, subject 2 observes subject 1's action, can infer his signal, and, presumably, on the basis of this, state the value of the good. We also ran a treatment in which we ask subject 2 to report his belief on the signal observed by his predecessor. Specifically, subject 1 had exactly the same task as in all other treatments, that is, to state his belief in the interval $[0, 100]$ after observing a signal. Subject 2 observed this choice and was asked to choose a number in $[0, 100]$, expressing the probability that the first subject had observed a good signal (green ball). Then he received

¹⁷We did not ask subjects to make a second decision in all rounds so that the first decision was made without knowing whether there would be a second decision; with this method, the subject was in a condition, we believe, very similar to that of subject 1 in the SL treatments.

a signal and had to make the same choice again. The elicitation method was identical to the other treatments. We ran 5 sessions for this treatment, using the same protocol as in Treatment SL and with a sequence of only two subjects. In each session there were 10 or 12 subjects and in each round they were randomly matched in pairs: one would be the first in the sequence and one the second.¹⁸

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
IDM1	0.7	1 or 2	-	-	36	30
IDM2	0.7	2	-	-	32	15
GC	0.7	2	10	5	56	15

Table 3.1: Treatments' main features.

SL: Social Learning; IDM: Individual Decision Making; GC: Guess the Color. Note that in SL2 there are 49 subjects since one session was run with 9 participants rather than 10 due to a last minute unavailability of one subject.

Finally, to study the robustness of our results to different specifications, we also ran a treatment in which 44 subjects received signals of stochastic precision in $(0.5, 1]$. The precision was only known to the subject receiving the signal. We used sequences of two subjects (like in the GC treatment). In the interest of space, we discuss this treatment in the Appendix.

4 Results of the SL treatments

Our main interest is to understand how human subjects weigh their private signal relative to the public information contained in the observation of the predecessor's action. To this aim, we focus on subjects' second decisions at time 2, when both private and public information are available. Before doing so, however, let us very briefly discuss the decisions of subjects at time 1 (who have only observed a private signal) and the first decisions of subjects at time 2 (based on the observation of their predecessor's choice only). The interested reader can find a detailed analysis in the Appendix.

4.1 How do subjects weigh their signal relative to their predecessor's action?

Let us now study how subjects update their beliefs upon receiving the signal at time 2, that is, let us focus on the second action at time 2. We will refer

¹⁸The total number of participants was 56 since in two sessions we recruited 10 subjects and in three sessions we recruited 12 subjects (which made the probability of a subject being matched with the same other participant even lower). The exchange rate for this treatment was 100 liras for 6 GBP.




Figure 4.1: Distribution of actions at time 1. The top (bottom) panel refers to actions upon receiving $s_1 = 1$ ($s_1 = 0$).




Figure 4.2: 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 4.3: Distribution of a_2^2 for confirming (top panel) or contradicting (bottom panel) s_2 .

to the first action that subjects take at time 2 as their “first belief” and to the second as their “posterior belief.” Figure 4.3 shows the frequency of posterior beliefs conditional on whether the subject received a signal confirming his first belief (i.e., $s_2 = 1$ after an action $a_2^1 > 50$ or $s_2 = 0$ after an action $a_2^1 < 50$) or contradicting it (i.e., $s_2 = 1$ after an action $a_2^1 < 50$ or $s_2 = 0$ after an action $a_2^1 > 50$).¹⁹ The figure is obtained (and all the following analysis is conducted) after transforming an action a_2^1 into $100 - a_2^1$ and the corresponding signal s_1 into $1 - s_1$ whenever $a_1 < 50$.

If subjects acted as in the PBE of the basic model (with objective precision, risk neutrality and no doubt about the rationality of their predecessor), in the case of confirming signal we would observe the entire distribution concentrated on 84. The empirical distribution shows much more heterogeneity. Nevertheless,

¹⁹In this analysis we exclude the cases in which the action at time 1 was $a_1 = 50$, since observing a 50 is uninformative. We do study the case in which a subject at time 2 observed an informative action at time 1 and chose $a_2^1 = 50$; in this case we distinguish whether the action observed at time 1 confirmed or contradicted the realization of the signal s_2 . 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 $a_1 < 50$).

the median action as well as the mode are indeed close to this prediction. 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 quantify these observations, we compute the weight that the subject puts on his signal by using the notation of Section 2 (to alleviate notation, we omit the arguments of α_2^2):

$$\frac{a_2^2}{100} = s_1 \frac{0.7\alpha_2^2 \frac{a_1^1}{100}}{0.7\alpha_2^2 \frac{a_1^1}{100} + (1-0.7)\alpha_2^2 \left(1 - \frac{a_1^1}{100}\right)} + (1-s_1) \frac{(1-0.7)\alpha_2^2 \frac{a_1^1}{100}}{(1-0.7)\alpha_2^2 \frac{a_1^1}{100} + 0.7\alpha_2^2 \left(1 - \frac{a_1^1}{100}\right)}. \quad (15)$$

In this formulation α_2^2 can be thought of as expressing subjective precisions (like in expression (6)) or as reflecting both subjective precisions and risk preferences (like in expression (9)). We prefer to think of α_2^2 as expressing subjective beliefs only, but this is a matter of interpretation.²⁰

In this interpretation, for $\alpha_2^2 = 1$, expression (15) gives the Bayesian updating formula, and so $\alpha_2^2 = 1$ 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 $\alpha_2^2 = 2$, 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 ($\alpha_2^2 = 0$) does not update at all upon observing it ($a_2^2 = 50$), whereas a subject who puts an infinite weight on it chooses an extreme action ($a_2^2 = 0$ or $a_2^2 = 100$), as if he were convinced that the signal fully reveals the value of the good. Finally, a negative value of α_2^2 indicates that the subject misreads the signal, e.g., interpreting a good signal as a bad one.

Table 4.1 reports the results.²¹ While in the case of a confirming signal the 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, 1.70.²² 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.²³

²⁰In the laboratory, subjects may choose different actions for many different reasons, including errors in updating beliefs or implementing the actions, risk preferences, subjective beliefs. As we said, we find subjective precisions a simple approach to capture all these departures, in the spirit of Savage.

²¹The value of α_2^2 is undetermined when $a_2^1 = 100$, therefore we exclude these cases. When $a_2^2 = 100$ we use the same approximation as previously discussed.

²²We ran a Mann-Whitney 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 = 0.000002).

²³As 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_1^1 , since the difference is for one observation only. Moreover, we cannot reject the hypothesis that the results, both

	1st Quartile	Median	3rd Quartile
α_2^2	0.69	1.16	2.03
α_2^2 (upon observing confirming signal)	0.55	0.96	1.34
α_2^2 (upon observing contradicting signal)	1.00	1.70	2.72

Table 4.1: 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.

To sharpen our understanding of the behavior in the SL treatment, we now look at how the weight on the signal changes with the first belief. Table 4.1 reports the quartiles for α_2^2 for three different classes of a_2^1 . As one can immediately observe, the asymmetry occurs for the last two classes, but not for the first.

	1st Quartile	Median	3rd Quartile
α_2^2 (upon observing confirming signal)			
Conditional on $50 < a_2^1 \leq 66.7$	0.66	0.97	1.16
Conditional on $66.7 < a_2^1 \leq 83.4$	0.18	0.91	1.55
Conditional on $a_2^1 > 83.4$	0.54	2.10	3.74
α_2^2 (upon observing contradicting signal)			
Conditional on $50 < a_2^1 \leq 66.7$	0.55	0.96	1.60
Conditional on $66.7 < a_2^1 \leq 83.4$	1.03	1.68	2.11
Conditional on $a_2^1 > 83.4$	2.53	3.34	4.23

Table 4.2: 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.

As we mentioned above, the median subject chose an action $a_2^1 > 66.7$ 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.

This result is in line with the LRTU rule, as shown in case i of Proposition 4. Essentially, subjects at time 2 who trust the predecessor’s rationality, stop trusting it after a contradicting signal. This asymmetric updating is not in agreement with the Bayesian model, even considering the extended one, and cannot be explained by the FBU rule. In Section 6, we will perform an econometric comparison of the different models of updating, taking into account the heterogeneity of actions, and we will show that indeed the LRTU model is the

for confirming and contradicting signals, are the same for the three treatments SL1, SL2 and SL3. (see the Appendix for details).

best fit of the data. Before we do that, let us make some important observations and (in the next section) discuss the results of our control treatments.

First, 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 $s_2 = 1$ 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 equal to 0, whereas it is slightly greater than 1). 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 would 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 to subjects in the experiment (i.e., first the private signal and then the predecessor’s action) we would not have been able to rule out this possibility.²⁴

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.²⁵

5 Controls: the IDM treatments

One may wonder whether the asymmetric updating 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 (possibly corresponding to a psychological bias other than base rate neglect or confirmation bias as mentioned above). To tackle this issue, as we explained, we ran two, individual decision making, control treatments. Table 5.1 shows the results of the IDM1 treatment, in which subjects receive two consecutive signals of the same precision 0.7. As one can see, the asymmetry and the overweight of the contradicting signal are

²⁴It is also worth mentioning that whereas in the social learning literature, as in much psychological literature, researchers have talked about “overconfidence” (see Section 7), in other experiments 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.

²⁵One could observe that if a subject chose, e.g., $a_2^1 = 84$ and then, after receiving a bad signal, $a_2^2 = 50$, the corresponding α_2^2 would be 2, which is compatible with the overweight we documented. It must be noticed, though, that if we exclude the cases in which $a_2^2 = 50$, the asymmetry remains and is actually even stronger (see Table B.3 in the Appendix).

not present in this treatment: 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). Table 5.2 shows the results of the IDM2 treatment. Remember that in such a treatment we want to mimic the fact that in social learning subject 2 only observes “a signal (subject 1’s action) about a signal (subject 1’s ball color).” In the IDM2 treatment a subject first observes a computer announcement about the signal (color of the ball) and then directly observes the color of a second ball drawn from the same urn. Again, there is no asymmetry in updating. Subjects slightly overweigh the signal they observe directly, but this slight overweight does not depend on whether the signal confirms or contradicts the announcement (we cannot reject the hypothesis that the two median weights are the same; Mann-Whitney U test, p -value = 0.27).

From these two control treatments, we can conclude that the asymmetric updating we observe in the SL treatment is not just a deviation from Bayesian updating after receiving two pieces of information; it is, instead, inherent to the social learning aspect of our study.

	1st Quartile	Median	3rd Quartile
α_2^2	0.64	1.08	2.06
α_2^2 (upon observing confirming signal)	0.64	1.30	2.48
α_2^2 (upon observing contradicting signal)	0.94	1.00	1.74

Table 5.1: Distribution of weights on the own signal in the IDM1 treatment.

The table shows the quartiles of the distribution of the weight on the own signal for the action at time 2 in the IDM1 treatment. The data refer to all cases in which the action at time 1 was different from 50.

	1st Quartile	Median	3rd Quartile
α_2^2	0.76	1.24	2.03
α_2^2 (upon observing confirming signal)	0.71	1.30	2.11
α_2^2 (upon observing contradicting signal)	0.81	1.24	1.73

Table 5.2: Distribution of weights on the own signal in the IDM2 treatment.

The table shows the quartiles of the distribution of the weight on the own signal for the action at time 2 in the IDM2 treatment. The data refer to all cases in which the action at time 1 was different from 50.

6 Econometric analysis

We now estimate the three models of updating presented in Section 2, the Bayesian Updating (BU) model, the LRTU model and the FBU model, with the data of the SL treatments. We will then perform a formal statistical comparison among them to quantify the evidence in favor of one model against the others.

As we know from the descriptive statistics, there is a lot of heterogeneity in the actions. We also know from the description of the three models that subjective precisions play a big role. To take this into account we use the data from the IDM1 treatment, in which we can infer subjective signal precisions directly from the chosen actions. Inferring some features of a model directly from the data of an auxiliary treatment in lieu of imposing extra parametric assumptions is, to the best of our knowledge, another novel contribution of our work to the experimental literature.

For the BU model, we estimate the parameter θ ; for the LRTU model we estimate the parameters (θ_*, θ^*, c) ; finally, for the FBU model we estimate the parameters (θ_*, θ^*) . We obtain the distribution of actions f and g as described in Section 2 directly from the subjects' actions at time 1.

We proceed to the estimation of the parameters by the generalized method of moments (GMM). In each of our models, the heterogeneity in the subjective precision of signals (q_2^{1S}, q_2^{2S}) induces a distribution of weights of updating of agent 2 for any fixed value of the parameters. The estimation strategy consists in finding the parameter values such that the distribution of weights predicted by a model is closest to the actual distribution of weights. With maximum likelihood, we would need to specify a parametric distribution for (q_2^{1S}, q_2^{2S}) . In our IDM experiments, however, we do observe the empirical distribution of (q_2^{1S}, q_2^{2S}) . 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.

As we did in our descriptive analysis, we focus on the three quartiles of the empirical distribution of the updating weights 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. In each of the three cases, we consider the distribution of weights 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 $B_3 = (83.4, 100]$. We then match such distributions of weights predicted by the model with the empirical ones at each of the three quartiles reported in Table 4.1. The estimate will, therefore, result from 27 moment conditions.²⁶

In the Appendix we describe all the other details of the econometric models and the estimation procedure. Here we only discuss the estimation results, reported in Table 6.1 (non-parametric bootstrapped standard errors in parenthesis).²⁷

It is useful to start our description with the estimates for the LRTU model. The lower bound estimate of c shows that, to “trust” a predecessor’s action, a subject needs the likelihood ratio to be greater than a threshold equal to 1.65,

²⁶For the BU model, as observed above, given agent 2’s first action, his second action only depends on q_2^{2S} . For this reason, the identifying information for θ is only based on the first action at time 2 (i.e., 9 moment conditions).

²⁷The GMM objective function does not have a unique minimizer for the parameter c : $c \in [1.65, 1.71]$. Table 6.1 hence reports the set of minimizers for c . Nevertheless, the other parameters have the same estimate for any $c \in [1.65, 1.71]$.

Model	θ	θ_*	θ^*	c
BU	0.30 (0.052)			
LRTU		0 (0.016)	0.30 (0.045)	[1.65, 1.71] (0.071)
FBU (and maxmin)		0.30 (0.070)	0.30 (0.067)	

Table 6.1: 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 the lower bound of the set of minimizers.

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 the proportion of a noise type is $\theta_* = 0$). When, instead, the threshold is not reached, he attaches a probability of 0.30 to the action coming from a noise type.

It is interesting to note one implication of these parameter estimates for subjects' behavior. After receiving a confirming signal, subjects keep trusting an action $a_1 \in (83.4, 100]$ (which they trust even before receiving the signal). After receiving a contradicting signal, in 68.9% of the cases they stop trusting such an action.

The BU model is a special case of the LRTU model in which the support of prior beliefs collapses to a singleton, $\theta_* = \theta^* = \theta$. We estimated θ for the BU model and obtained an estimate of 0.30 (standard error: 0.052), which coincides with the estimate of θ^* in the LRTU model. The crucial question is whether the LRTU model provides a better explanation for the observed behavior than the BU model. 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 degrees of freedom (has more parameters) than the other. We develop a model comparison hypothesis test between BU and LRTU models (see Appendix). The p -value of the test is 0.005, that is, we can reject the null hypothesis that the BU model with parameter value $\theta = 0.3$ is the true data generating process 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.

As for FBU model, the estimates for θ_* and θ^* are both identical to 0.30. That is, under the FBU specification, the estimates indicate that subject 2 does not have multiple priors for the predecessor's rationality. Furthermore, if we compare this result to that for the BU model, we note that the FBU model's estimates coincide with the estimated BU, despite the fact that FBU model has an extra degree of freedom. In other words, adding multiple priors and assuming FBU does not improve the fitness of the more restrictive BU model. Combined with the model comparison result of LRTU versus BU model, we conclude that

the fitness of the LRTU model outperforms the FBU model. In the Appendix, we also estimate the FBU model with the Hurwitz criterion instead of the maxmin criterion, and again find that the FBU model with Hurwitz criterion fits the data significantly worse than the LRTU model.

We would like to emphasize that our approach does not impose any parametric restriction on the subjective precisions: the evidence in favor of the LRTU model does not depend on any distributional specification for heterogeneous beliefs.

7 Results of the GC Treatment

We also tested the hypothesis that subjects change their mind about the predecessors' rationality in a Bayesian fashion (or not) directly, through the GC treatment. Recall that in such a treatment, subject 2 is asked to state his belief that the predecessor observed the good signal (the green ball) first after observing the predecessor's action only and, then, after receiving the private signal too. The stated belief is equivalent to the stated belief that the first decision maker is rational, according to the definition used in Section 2.

Our main interest is in the beliefs after the second subject has also received the private signal. Figure 7.1 reports the distributions of these beliefs, distinguishing the cases of contradicting and confirming signals.²⁸ We denote these beliefs about the predecessor's signal by x_t^j , to avoid confusion with the beliefs about the value, previously denoted by a_t^j . Analogously to the analysis of Section 4, the figure is obtained after transforming an action $x_2^1 < 50$ into $100 - x_2^1$ and the corresponding signal s_1 into $1 - s_1$. For each case, we plot two frequencies: the left bar in each bin shows the empirical frequency of x_2^2 ; the right bar, instead, shows the frequency of the theoretical beliefs (denoted by \hat{x}_2^2) that subjects would have reported, had they updated in a Bayesian way (with signal precisions equal to the objective precisions).²⁹

The empirical and theoretical frequencies are remarkably different in the contradicting signal case, whereas they are remarkably similar in the confirming signal case. For the contradicting case, the theoretical belief distribution first-order stochastically dominates the empirical belief distribution: subjects' beliefs about their predecessors' signals (equivalent to their predecessors' rationality, according to the definitions of Section 2) shift more drastically than what Bayesian updating predicts. Note, in particular, that the mode of the empirical distribution is 50, that is, after a contradicting signal, subjects give equal chance to the predecessor having received a green or red ball.

For confirming signals, the theoretical, Bayesian, belief distribution matches the empirical belief distribution remarkably well, indicating that subjects up-

²⁸The signals are defined as confirming or contradicting with respect to a_1 . There are no relevant differences when the definition of confirming or contradicting is with respect to x_2^1 .

²⁹We also simulated the updated beliefs using subjective precisions, elicited from the IDM treatment (as in the econometrics analysis). The results are almost identical and presented in the Appendix (Figure D.3).




Figure 7.1: Distribution of second decisions (beliefs on the predecessor’s signal) at time 2. The top panel refers to confirming signals; the bottom to contradicting signals. In each panel, we report the frequency of the actual decisions (left bars) and of the simulated (using objective precisions), theoretical ones (right bars).

date their beliefs in line with Bayesian updating. We perform two-sample Kolmogorov-Smirnov tests and reject the null hypothesis of equivalence of the distributions (p -value < 0.01) for the contradicting signal case, whereas we cannot reject the null for the confirming signal case (p -value = 0.13).

In the Appendix we quantify the departure from Bayesian updating in the GC treatment, by computing a subject’s weight on the signal with a similar method to that of Section 4. In the case of a confirming signal the median weight coincides with the Bayesian one, whereas in the case of a contradicting signal it is considerably higher, 2.86.

We also simulated the actions (beliefs on the value) implied by the beliefs about the predecessor’s signal and obtained distributions remarkably similar to the empirical distributions reported in Figure 4.3 of Section 4 (see Figure D.5 in Appendix).

Overall, the results of this treatment show how subjects change their view about the predecessors’ rationality in a direct way. The changes are not in line with Bayesian updating, but in line with the LRTU, the theoretical model we proposed to explain the asymmetric updating observed in the SL treatments.

8 Discussion and related literature

To conclude, we want to explain how our study relates to different strands of literature.

First, the idea that one uses the data to select the prior from a set of priors is well and long established in statistics. It 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). Selecting the prior is at the root of the debate in Bayesian statistics and, in particular, in empirical Bayes procedures (see, also Cox and Hinkley, 1979; Good, 1983; Berger and Berliner, 1986). In this methodology, the set of priors from which one prior is estimated is invariant to the arrival of new information, similarly to our LRTU rule.

Second, in economic decision theory, MLU was first introduced in the literature on ambiguity aversion. It is, indeed, one of the two main families of updating with ambiguous beliefs, the other being FBU (see Gilboa and Marinacci, 2013 for a survey). The literature has mostly considered one-shot decision problems, noting important conceptual difficulties when dealing with the arrival of new information and belief updating.³⁰ 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. Our LRTU model generalizes the MLU rule since it lets the agent require particularly strong evidence to choose a prior. More importantly, the set of priors from which the agent chooses does not shrink to a singleton as, implicitly, in Gilboa and Schmeidler (1993).³¹

Third, more recently, a model of updating based on hypothesis testing has been repropounded in the economics literature in a contribution by Ortoleva (2012) who axiomatizes this model to deal with the important question of belief updat-

³⁰There is also an important experimental literature devoted to estimating ambiguity aversion in one-shot decision problems, see, e.g., Ahn *et al.* (2014) and the papers discussed therein.

³¹Gilboa 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. Perhaps, our model can be seen somehow in between these two extremes. In the LRTU model, the subject does pick one prior, but this does not eliminate the other priors forever, since the subject can pick another prior after new information arrives. A model in which the agent picks different priors every time new information arrives exhibits a form of time inconsistency. 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). For a theoretical investigation of dynamically consistent updating of ambiguous beliefs see Hanany *et al.* (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). As we emphasized already before, we view this as being very much in the tradition of the statistics literature dating back to the *Type-II* maximum likelihood of Good (1965).

ing after observing unexpected events.³² In his Hypothesis Testing 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 lower than a specific threshold, 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. Ortoleva (2012) highlights the case in which paradigm shift would occur essentially after unforeseen events (which amounts to setting the threshold probability to 0), but the general version of his theory is very similar in spirit to our approach, here applied to the modelling of an agent’s belief about the rationality of another agent. Our parameter c can be viewed as playing a role similar to that of 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 — a bias we actually estimate). Our evidence can also be read as supportive of Ortoleva (2012)’s approach.

Finally, our experiment is related to the experimental social learning literature. Compared to the canonical experiments (following Anderson and Holt, 1997), we made two important procedural changes. First, in previous experiments subjects made a decision in a binary action space, whereas we ask subjects to choose 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 the another.³³ Second, in previous experiments subjects made one decision only, after observing both the predecessors and the private signal. In our experiment, instead, subject 2 states his belief twice (before and after the private signal), which allows us to identify the weight on the signal in a neat way. Many previous papers (e.g., Nöth and Weber, 2003; Çelen and Kariv, 2004; Goeree *et al.*, 2007) have reported that subjects overweigh their own signal. Our results are clearly in line with this finding. In those experiments, though, the behavior we document and explain could not be observed, exactly because each subject only made one decision and in a binary action space. When subjects had a signal in agreement with the previous actions, they typically followed it and chose the same action. This decision is essentially uninformative for the experimenter on how subjects update on their private information. In fact, on the basis of previous 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.

³²In another recent contribution, Weinstein (2017) also proposes that agents may be non Bayesian when something unexpected happens and the agent decides for a “paradigm shift.”

³³Within the discrete action space experiments, exceptions to the binary action space are the financial market experiments of Cipriani and Guarino (2005, 2009) where subjects can choose to buy, to sell or not to trade. In the interesting experimental design of Çelen 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.

9 Conclusion

We studied the behavior of human subjects in an experiment in which they report their belief about the value of a good, first after observing a predecessor's evaluation and then upon receiving a private signal. We found an important deviation from Bayesian updating: news contradicting the previous belief are overweighted compared to news that confirm it. This asymmetric updating is inherent to how people learn from others (it is not a psychological bias in updating) and cannot be explained by subjective beliefs in the vein of Savage (despite the flexibility of this approach). We explain this asymmetry by using a model of updating multiple priors on the rationality of others. We see our work as a step in understanding how to update multiple priors (something very little explored) and how to model beliefs about others' rationality. In future work, this approach could be combined with models of bounded rationality, such as level-k, QRE, ABEE. We suspect our findings are relevant in the field, given that rarely there are data on the probability distribution of others' rationality. The way subjects update, as we have observed it in our study, may also have many interesting implications for social and economic applications.³⁴ For instance, in a financial market, after a period of boom (recession), the arrival of new, negative (positive) information may trigger a stronger reaction by traders and so higher price volatility. More empirical work to explore such applications would be, we believe, very valuable.

References

- [1] Ahn, D., Choi, S., Gale, D., Kariv, S., 2014. Estimating Ambiguity Aversion in a Portfolio Choice Experiment. *Quantitative Economics*, 5(2), 195-223.
- [2] Anderson, L., Holt, C., 1997. Information Cascades in the Laboratory. *American Economic Review*, 87(5), 847-862.
- [3] Becker, G., DeGroot, M., Marschak, J., 1964. Measuring Utility by a Single-Response Sequential Method. *Behavioral Science*, 9(3), 226-232.
- [4] Berger, J., 1985. *Statistical Decision Theory and Bayesian Analysis*. Springer.
- [5] Berger, J., Berliner, L.M., 1986. Robust Bayes and Empirical Bayes Analysis with ε -contaminated Priors. *The Annals of Statistics*, 461-486.
- [6] Çelen, B., Kariv, S., 2004. Distinguishing Informational Cascades from Herd Behavior in the Laboratory. *American Economic Review*, 94(3), 484-498.

³⁴In a recent study, Giustinelli and Pavoni (2017) use survey measures of ambiguous beliefs in the context of high school track choice.

- [7] Cipriani, M., Guarino, A., 2005. Herd Behavior in a Laboratory Financial Market. *American Economic Review*, 95(5), 1427-1443.
- [8] 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.
- [9] Cox, D. Hinkley, D., 1979. Theoretical Statistics. CRC Press.
- [10] Epstein, L. G., Schneider, M., 2007. Learning Under Ambiguity. *Review of Economic Studies*, 74(4), 1275–1303.
- [11] Fischbacher, U. , 2007. Z-Tree: Zurich Toolbox for Ready-Made Economic Experiments. *Experimental Economics*, 10(2), 171-178.
- [12] Gale, D., 1996, What Have We Learned from Social Learning?, *European Economic Review*, 40(3-5), 617-628.
- [13] 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.
- [14] Gilboa, I., Schmeidler, D.,1993. Updating Ambiguous Beliefs. *Journal of Economic theory*, 59(1), 33-49.
- [15] Giustinelli, P., Pavoni, N., 2017. The Evolution of Awareness and Belief Ambiguity in the Process of High School Track Choice. *Review of Economic Dynamics*, 25, 93-120.
- [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] Good, I., 1983. Good thinking: The foundations of probability and its applications. University of Minnesota Press.
- [19] Guarino, A., Jehiel, P., 2013. Social Learning with Coarse Inference. *American Economic Journal: Microeconomics*, 5(1), 147-74.
- [20] Hanany, E., Klibanoff, P., 2007. Updating Preferences with Multiple Priors. *Theoretical Economics*, 2(3), 261–298.
- [21] Hurwicz, L., 1951. The Generalised Bayes Minimax Principle: A Criterion for Decision Making Under Uncertainty. *Cowles Commission Discussion Paper Statistics*, 355.
- [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] Nöth, M., Weber, M., 2003. Information Aggregation with Random Ordering: Cascades and Overconfidence. *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., 2012. Modelling the Change of Paradigm: 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.
- [30] Weinstein, J., 2017. Bayesian Inference Tempered by Classical Hypothesis Testing, mimeo.

APPENDIX (FOR ONLINE PUBLICATION ONLY)

A SL treatment, time 1 and time 2 first action

We illustrate the decisions of subjects at time 1 (when they have only observed a private signal) and the first decisions of subjects at time 2, based on the observation of their predecessor’s choice only.

A.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 IDM1 treatment; for this reason we pool all data together (for a total of 1380 observations).³⁵

Figure A.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 (39%) are in line with Bayesian updating, deviating from it by less than 5 units; 22% of actions are smaller than the Bayesian one and 39% of actions are larger. Note, in particular, that in 9% 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 (3%) 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 13% in 50 and of 12% in 0, and with other actions distributed in a similar manner to what explained above.

Similarly to α_2^2 defined in equation (6) of the main text, we define α_1 the weight that the agent 1 attaches to his own signal as

$$a_1 = 100 \left(s_1 \frac{q^{\alpha_1}}{q^{\alpha_1} + (1-q)^{\alpha_1}} + (1-s_1) \frac{(1-q)^{\alpha_1}}{q^{\alpha_1} + (1-q)^{\alpha_1}} \right), \quad (16)$$

³⁵We ran a Mann-Whitney 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 = 0.86). 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 = 0.22), SL1 with SL3 (p-value = 0.46), or SL2 with SL3 (p-value = 0.07).




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

where α_1 satisfies³⁶

$$\alpha_1 = \frac{\ln\left(\frac{a_1/100}{1-a_1/100}\right)}{\ln\left(\frac{q}{1-q}\right)}.$$

Note that for $\alpha_1 = 1$ expression (16) gives the Bayesian updating formula, and so $\alpha_1 = 1$ 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 $\alpha_1 = 2$, 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 ($\alpha_1 = 0$) does not update at all upon observing it ($a_1 = 50$), whereas a subject who puts an infinite weight on it chooses an extreme action ($a_1 = 0$ or $a_1 = 100$), as if he were convinced that the signal fully reveals the value of the good. Finally, a negative value of α_1 indicates that the subject misreads the signal, e.g., interpreting a good signal as a bad one.

³⁶Recall that a subject made many choices in the same experiment, since he participated in several rounds. 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 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.

Table A.1 reports the quartiles of the distribution of the computed α_1 .³⁷ Note that the median α_1 is 1, indicating that the median subject is actually Bayesian.³⁸

	1st Quartile	Median	3rd Quartile
α_1	0.73	1.00	2.05

Table A.1: 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.

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 these subjective beliefs, there is no systematic bias to overweight or underweight the signal: the median belief corresponds to Bayesian updating.

A.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 only. For that purpose, 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. Here we only consider the data from the SL treatment.

³⁷When $a_1 = 0$ or 100, we compute α_1 by approximating $a_1 = 0$ with ε and $a_1 = 100$ with $100 - \varepsilon$ (with $\varepsilon = 0.01$). We prefer to report the quantiles rather than the mean or other statistics whose computations are affected by the approximation of α_1 .

³⁸In this analysis, we have allowed for heterogeneous weights on the signal and assumed that subjects did state their beliefs correctly. Another approach would be to take into account that subjects could have made mistakes while reporting their beliefs, as in the following model:

$$a_1 = 100 \left(s_1 \frac{q^{\alpha_1}}{q^{\alpha_1} + (1-q)^{\alpha_1}} + (1-s_1) \frac{(1-q)^{\alpha_1}}{q^{\alpha_1} + (1-q)^{\alpha_1}} \right) + \varepsilon_1, \quad (17)$$

where the weight on the signal is the same for all subjects but each subject makes a random mistake ε_1 . It is easy to show that, as long as the error term has zero median, constant α_1 implied by the median of a_1 in this model coincides with the median of heterogeneous α_1 considered in the text.

Other interpretations are possible. One may, for instance, argue that the fact that a subject chooses 70, while compatible with Bayesian updating, does not necessarily indicate 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. See also our discussion of risk preferences in the main text.




Figure A.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 $a_1 < 50$).

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 3.6% 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 A.2). Compared to Figure A.1, Figure A.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.

We replicated the analysis for agent 1, by replacing the case in which the subject observed a good (bad) signal with the case in which the subject observed $a_1 > 50$ ($a_1 < 50$, respectively), and so chose a_2^1 such that

$$a_2^1 = 100 \left(1 \{a_1 > 50\} \cdot \frac{q^{\alpha_2^1}}{q^{\alpha_2^1} + (1-q)^{\alpha_2^1}} + 1 \{a_1 < 50\} \cdot \frac{(1-q)^{\alpha_2^1}}{q^{\alpha_2^1} + (1-q)^{\alpha_2^1}} \right). \quad (18)$$

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 0.7. By applying this model, we obtain the results reported in Table A.2. The median weight is (slightly) lower than 1 and the first and third quartiles are 0.13 and 1.4 (versus 0.73 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.^{39 40}

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 heterogenous 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 A.2.⁴¹

	1st Quartile	Median	3rd Quartile
α_2^1	0.14	0.94	1.39
α_2^1 (upon observing $50 < a_1 \leq 66.7$)	0	0.48	0.87
α_2^1 (upon observing $66.7 < a_1 \leq 83.4$)	0	0.89	1.33
α_2^1 (upon observing $a_1 > 83.4$)	0.93	1.31	2.73

Table A.2: 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 ($50 < a_1 \leq 66.7$) and, although less, those in a neighborhood of the Bayesian one ($66.7 < a_1 \leq 83.4$). 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 increasing in the distance from 50. A well known model in which errors are

³⁹We 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 = 0.015). We repeated the same test considering only the IDM treatment for a_1 , and obtain p-value = 0.149.

⁴⁰Discounting 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.”

⁴¹We 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.

inversely related to their costs is the Quantal Response Equilibrium (which also assumes expectations are rational). Our results are, however, not compatible with such 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 3.6% 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 which, as illustrated in Section 2, agent 2 entertains a subjective belief about the rationality of agent 1.

B Additional descriptive statistics

B.1 Behavior after choosing 50 at time 2

	1st Quartile	Median	3rd Quartile
α_2^2	0.00	1.02	2.35
α_2^2 (upon observing confirming signal)	0.36	1.06	2.37
α_2^2 (upon observing contradicting signal)	0.00	0.98	1.86

Table B.1: 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	1.15
α_2^2 (upon observing confirming signal)	0.00	1.00	1.74
α_2^2 (upon observing contradicting signal)	0.00	0.00	0.00

Table B.2: 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.

B.2 More descriptive statistics in the second action at time 2

One could observe that if a subject chose, e.g., $a_2^1 = 84$ and then, after receiving a bad signal, chose $a_2^2 = 50$, the corresponding α_2^2 would be 2, which is compatible with the overweight we documented. It must be noticed, though, that if we

exclude the cases in which $a_2^2 = 50$, nevertheless the asymmetry remains, as one can appreciate by looking at Table B.3.

	1st Quartile	Median	3rd Quartile
α_2^2	0.71	1.16	2.11
α_2^2 (upon observing confirming signal)	0.55	0.96	1.35
α_2^2 (upon observing contradicting signal)	1.34	2.07	2.96

Table B.3: 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.

B.3 Social Learning treatments: tests

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-Whitney 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 B.4.

	Time 1	Time 2.1	Time 2.2 - confirming	Time 2.2 - contradicting
SL1 versus SL2	0.22	1.00	0.75	0.55
SL1 versus SL3	0.46	0.13	0.29	0.75
SL2 versus SL3	0.07	0.06	0.35	0.17

Table B.4: 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.

C Econometric analysis

In this appendix, we present the econometric analysis yielding the results summarized in Section 6 of the main text. The first part illustrates the estimation

and model comparison hypothesis testing of the LRTU and BU models. The second part presents the estimation and model comparison hypothesis testing for the FBU model.

Through the entire section, we index an observation of $(s_1, a_1, a_2^1, s_2, a_2^2)$ by i (as customary in econometric expositions).

C.1 LRTU and BU models

As we have discussed in the main text, the reported beliefs at time 1 and at time 2 are quite heterogeneous, with non-regular features (e.g., multi-modality, asymmetry). To take this into account in our econometric models, rather than imposing parametric specifications, we use the IDM treatment observations and obtain a nonparametric estimator for the distribution of the unobservable heterogeneity.

Our purpose is to understand which model explains the behavior of subjects at time 2 best. The LRTU and BU models will have two common ingredients:

- a) subjective beliefs on the informativeness (precision) of the private signal;
- b) 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.

Let us start discussing point a) 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, we let the subjective precisions $q_{2i}^{1S} = \Pr(s_{1i} = 1|V = 100) = \Pr(s_{1i} = 0|V = 0)$ and $q_{2i}^{2S} = \Pr(s_{2i} = 1|V = 100) = \Pr(s_{2i} = 0|V = 0)$ vary for each observation i . 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_{2i}^{1S} , and q_{2i}^{2S} , 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 $q_{2i}^{1S} = 0.73$; in the IDM treatment, $a_{2i} = 80$ after $a_{1i} = 73$ and $s_{2i} = 1$ is equivalent to $q_{2i}^{2S} = 0.60$). We will use the so-recovered empirical distribution of q_{2i}^{1S} , 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 $(q_{2i}^{1S}, q_{2i}^{2S})$ can be recovered in the IDM treatment.⁴²

Let us move to point b). As we said in the main text, we assume that a subject at time 2 believes that the predecessor is of two types: either “rational” (τ_r) or “noise” (τ_n), with $\Pr(\tau_n) \equiv \theta$. 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,

⁴²In 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.

instead, chooses any action between 0 and 100 independently of the signal.⁴³ The BU model assumes a unique θ ; in the LRTU model, instead, the beliefs on the predecessor's rationality consist in a set of priors $[\theta_*, \theta^*]$. We will estimate the unique θ or the lower and upper bounds θ_* and θ^* by fitting the models to the data. For the LRTU model, we will also estimate the threshold parameter c .

As we know from Section 4, the empirical distribution of actions at time 1 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 Table A.2, 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. 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,

$$g_\sigma(a_1) = \frac{1}{\Phi_\sigma([0, 100])} \sum_{l=1}^3 \frac{\Phi_\sigma(B_l)}{|B_l|} \cdot 1\{a_1 \in B_l\}, \text{ for } a_1 > 50, \quad (19)$$

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. 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.⁴⁴ 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 $\Pr(a_1 > 50 | s_1 = 0) = \Pr(a_1 < 50 | s_1 = 1) = \frac{\theta}{2}$.

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

⁴³As 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 the signal realization, and this is revealed under the minimal requirement that the subject updates in the right direction.

⁴⁴Given the transformation of data, all incorrect actions are below 50.




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

the empirical distribution of time 1 subject's actions. The histogram density for the actions greater than 50 is

$$h(a_1) = \sum_{l=1}^3 \hat{b}_l 1_{\{a_1 \in B_l\}} \quad \text{for } a_1 > 50, \quad (20)$$

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.⁴⁵ Note, however, that not 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 be⁴⁶

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

Figure C.1 shows the estimated histograms.

Given these histograms, a (rational) subject i at time 2, observing an action

⁴⁵Note that we exclude $a_{1i} = 50$. This action is uninformative and, therefore, has a different status from any other action.

⁴⁶Recall 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.

$a_{1i} > 50$, has the following conditional beliefs (density functions):

$$\begin{aligned} \Pr(a_{1i}|V = 100, \tau_r) &= \tag{21} \\ \Pr(a_{1i}|s_{1i} = 1, V = 100, \tau_r)q_{2i}^{1S} + \Pr(a_{1i}|s_{1i} = 0, V = 100, \tau_r)(1 - q_{2i}^{1S}) &= q_{2i}^{1S} f(a_{1i}), \\ \Pr(a_{1i}|V = 0, \tau_r) &= (1 - q_{2i}^{1S})f(a_{1i}), \\ \Pr(a_{1i}|V = 100, \tau_n) &= \Pr(a_{1i}|V = 0, \tau_n) = g(a_{1i}). \end{aligned}$$

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 (θ) and of the rational type ($1 - \theta$).

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

The BU model

According to the BU model, given a prior belief θ on the proportion of noise type subjects at time 1, a subject applies Bayes's rule to determine his first action. Being explicit about its dependence on θ and q_{2i}^{1S} , we denote it by

$$\begin{aligned} a_{2i,B}^1 &\equiv a_2^{1*}(\theta, q_{2i}^{1S}) \equiv 100 \Pr(V = 100|a_{1i}) = 100 \frac{(1 - \theta)q_{2i}^{1S} f(a_{1i}) + g(a_{1i})\theta}{(1 - \theta)f(a_{1i}) + 2g(a_{1i})\theta} \\ &= 100 \frac{(1 - \theta)q_{2i}^{1S} \frac{f(a_{1i})}{g(a_{1i})} + \theta}{(1 - \theta)\frac{f(a_{1i})}{g(a_{1i})} + 2\theta}. \end{aligned}$$

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 $a_{2i,B}^2 \equiv a_{2,B}^2(\theta, q_{2i}^{1S}, q_{2i}^{2S})$ such that the following equality holds:

$$l\left(\frac{a_{2,B}^2(\theta, q_{2i}^{1S}, q_{2i}^{2S})}{100}\right) = l\left(\frac{a_2^{1*}(\theta, q_{2i}^{1S})}{100}\right) + l(q_{2i}^{2S}); \tag{23}$$

similarly, after a contradicting signal, action $a_{2i,B}^2 = a_{2,B}^2(\theta, q_{2i}^{1S}, q_{2i}^{2S})$ will satisfy

$$l\left(\frac{a_{2,B}^2(\theta, q_{2i}^{1S}, q_{2i}^{2S})}{100}\right) = l\left(\frac{a_2^{1*}(\theta, q_{2i}^{1S})}{100}\right) + l(1 - q_{2i}^{2S}). \tag{24}$$

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

The LRTU model

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

$$\frac{\Pr(a_{1i}|\tau = \tau_r)}{\Pr(a_{1i}|\tau = \tau_n)} = \frac{\frac{1}{2}q_{2i}^{1S} f(a_{1i}) + \frac{1}{2}(1 - q_{2i}^{1S})f(a_{1i})}{g(a_{1i})} = \frac{f(a_{1i})}{2g(a_{1i})}. \tag{25}$$

In particular, he selects $\theta_i^1 = \theta^1(a_{1i})$ as follows:

$$\theta_i^1 = \begin{cases} \theta_* & \text{if } \frac{f(a_{1i})}{g(a_{1i})} \geq 2c, \\ \theta^* & \text{if } \frac{f(a_{1i})}{g(a_{1i})} < 2c. \end{cases} \quad (26)$$

He then applies Bayes's rule to determine his first action, $a_{2i,L}^1 \equiv a_2^{1*}(\theta_i^1, q_{2i}^{1S})$, which is identical to expression (22), after substituting θ_i^1 to θ . Note that $a_{2i,L}^1$ varies from $100q_{2i}^{1S}$ to 50 as θ_i^1 varies from 0 to 1. Moreover, note that although the same q_{2i}^{1S} was used both in (25) and in (22), (25) does not depend on q_{2i}^{1S} .

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

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

Therefore,

$$\theta_{2i,conf}^2 \equiv \theta^2(a_{1i}, s_2 = 1; q_{2i}^{1S}, q_{2i}^{2S}) = \begin{cases} \theta_* & \text{if } \frac{f(a_{1i})}{g(a_{1i})} \geq \frac{c}{q_{2i}^{1S} q_{2i}^{2S} + (1 - q_{2i}^{1S})(1 - q_{2i}^{2S})}, \\ \theta^* & \text{if } \frac{f(a_{1i})}{g(a_{1i})} < \frac{c}{q_{2i}^{1S} q_{2i}^{2S} + (1 - q_{2i}^{1S})(1 - q_{2i}^{2S})}. \end{cases} \quad (27)$$

Given $\theta_{2i,conf}^2$ and q_{2i}^{2S} , $a_{2i,L,conf}^2 \equiv a_{2,L}^2(\theta_{2i,conf}^2, q_{2i}^{1S}, q_{2i}^{2S})$ satisfies

$$l\left(\frac{a_{2,L}^2(\theta_{2i,conf}^2, q_{2i}^{1S}, q_{2i}^{2S})}{100}\right) \equiv l\left(\frac{a_2^{1*}(\theta_{2i,conf}^2, q_{2i}^{1S})}{100}\right) + l(q_{2i}^{2S}), \quad (28)$$

where $a_2^{1*}(\theta_{2i,conf}^2, q_{2i}^{1S})$ is equal to (22) with $\theta_{2i,conf}^2$ in place of θ .

Note that the threshold in (27) is lower than that in (26).

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

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

and, therefore,

$$\theta_{2i,cont}^2 \equiv \theta^2(a_{1i}, s_2 = 0; q_{2i}^{1S}, q_{2i}^{2S}) = \begin{cases} \theta_* & \text{if } \frac{f(a_{1i})}{g(a_{1i})} \geq \frac{c}{q_{2i}^{1S}(1 - q_{2i}^{2S}) + (1 - q_{2i}^{1S})q_{2i}^{2S}}, \\ \theta^* & \text{if } \frac{f(a_{1i})}{g(a_{1i})} < \frac{c}{q_{2i}^{1S}(1 - q_{2i}^{2S}) + (1 - q_{2i}^{1S})q_{2i}^{2S}}. \end{cases} \quad (29)$$

Given $\theta_{2i,cont}^2$ and q_{2i}^{2S} , $a_{2i,L,cont}^2 = a_{2,L}^2(\theta_{2i,cont}^2, q_{2i}^{1S}, q_{2i}^{2S})$ satisfies

$$l\left(\frac{a_{2,L}^2(\theta_{2i,cont}^2, q_{2i}^{1S}, q_{2i}^{2S})}{100}\right) \equiv l\left(\frac{a_2^{1*}(\theta_{2i,cont}^2, q_{2i}^{1S})}{100}\right) + l(1 - q_{2i}^{2S}). \quad (30)$$

Note that the threshold in (29) is higher than that in (26): a confirming signal lowers the threshold to trust the predecessor's rationality, whereas a contradicting signal raises it.

C.2 Estimation methodology and results

We estimate the two 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 for 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 $(q_{2i}^{1S}, q_{2i}^{2S})$. In our experiment, however, we do observe the empirical distribution of $(q_{2i}^{1S}, q_{2i}^{2S})$. 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.

In the descriptive analysis in the main text, 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).⁴⁷

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_{2i}^{1S} and q_{2i}^{2S} to those that are weakly greater than 0.5.

C.2.1 Estimating the LRTU model

Consider first the estimation of the LRTU model. The parameters to be estimated are $\Phi \equiv (\theta_*, \theta^*, c)$, $0 \leq \theta_* \leq \theta^* \leq 1$, and $c \geq 0$. To make the dependence on the parameters explicit, we express the LRTU model actions as $a_{2i,L}^1(\Phi) \equiv a_2^{1*}(\theta_i^1, q_{2i}^{1S}; \Phi)$, $a_{2i,L,conf}^2(\Phi) \equiv a_{2,L}^2(\theta_{2i,conf}^2, q_{2i}^{1S}, q_{2i}^{2S}; \Phi)$ for the confirming signal case, where $\theta_{2i,conf}^2$ is defined in (27), and $a_{2i,L,cont}^2(\Phi) \equiv a_{2,L}^2(\theta_{2i,cont}^2, q_{2i}^{1S}, q_{2i}^{2S}; \Phi)$ for the contradicting signal case, where $\theta_{2i,cont}^2$ is defined in (29). For given Φ , a_{1i} , and $s_{2i} = 1$, the heterogeneity in subjective signal precisions generates the joint distribution of the time 2 actions $(a_{2i,L}^1(\Phi), a_{2i,L,conf}^2(\Phi))$.

⁴⁷For the BU model, as observed above, given $a_{2i,B}^1(\theta, q_{2i}^{1S})$, action $a_{2i,B}^2(\theta, q_{2i}^{1S}, q_{2i}^{2S})$ only depends on q_{2i} . For this reason, the estimate of θ is only based on the first action at time 2 (i.e., on 9 moment conditions).

If the LRTU model were the true data generating process, then, at the true value of Φ , the conditional distribution of $(a_{2i,L}^1(\Phi), a_{2i,L,conf}^2(\Phi))$ given $(a_{1i}, s_{2i} = 1)$ generated from heterogeneous $(q_{2i}^{1S}, q_{2i}^{2S})$ would coincide with the actual conditional distribution of (a_{2i}^1, a_{2i}^2) . This implies that, for any integrable function $h(a_{2i}^1, a_{2i}^2)$,

$$E [h(a_{2i}^1, a_{2i}^2) - E_Q [h(a_{2i,L}^1(\Phi), a_{2i,L,conf}^2(\Phi)) | a_{1i}, s_{2i} = 1]] = 0$$

holds at the true Φ for every a_{1i} , where the inner expectation $E_Q[\cdot]$ is the expectation with respect to the joint distribution Q of (q_2^{1S}, q_2^{2S}) , which we assume to 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} . Specifically, as we said, for Q we use the empirical distribution of precisions. Hence,

$$\begin{aligned} & E_Q [h(a_{2i,L}^1(\Phi), a_{2i,L,conf}^2(\Phi))] \\ & \approx \frac{1}{J} \sum_j h(a_2^{1*}(\theta_i^1, q_{2j}^{1S}; \Phi), a_{2,L}^2(\theta^2(a_{1i}, s_2 = 1; q_{2j}^{1S}, q_{2j}^{2S}), q_{2j}^{1S}, q_{2j}^{2S}; \Phi)), \end{aligned}$$

where the index j indicates an observation of $(q_{2j}^{1S}, q_{2j}^{2S})$ and J is the number of observations of $(q_{2j}^{1S}, q_{2j}^{2S})$ available in our dataset. Specifically, when $h(\cdot, \cdot)$ involves only $a_{2i,L}^1(\Phi)$, the marginal distribution of q_2^{1S} suffices to compute $E_Q(h(a_{2i,L}^1(\Phi)))$. Therefore, we construct the empirical distribution of q_2^{1S} 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 $a_{2i,L}^1(\Phi)$ and $a_{2i,L,conf}^2(\Phi)$, we construct the empirical distribution of (q_2^{1S}, q_2^{2S}) using the observations (a_{1i}, a_{2i}) in the IDM treatment only, restricted to $50 \leq a_{1i} < 100$ and $a_{2i} \geq 50$.⁴⁸ The total number of observations used to construct the empirical distribution of (q_2^{1S}, q_2^{2S}) amounts to $J = 440$.

Similarly, for the contradicting signal case we have that

$$E [h(a_{2i}^1, a_{2i}^2) - E_Q [h(a_{2i,L}^1(\Phi), a_{2i,L,cont}^2(\Phi)) | a_{1i}, s_{2i} = 0]] = 0$$

holds for any a_{1i} .

These moment conditions imply the following unconditional moment conditions:

$$E [s_{2i} \cdot (h(a_{2i}^1, a_{2i}^2) - E_Q [h(a_{2i,L}^1(\Phi), a_{2i,L,conf}^2(\Phi))])] = 0, \quad (31)$$

$$E [(1 - s_{2i}) \cdot (h(a_{2i}^1, a_{2i}^2) - E_Q [h(a_{2i,L}^1(\Phi), a_{2i,L,cont}^2(\Phi))])] = 0. \quad (32)$$

When $h(a_{2i}^1, a_{2i}^2)$ only depends on a_{2i}^1 , s_{2i} plays no role and the moment conditions (31) and (32) reduce (with a slight abuse of notation) to

$$E [h(a_{2i}^1) - E_Q [h(a_2^{1*}(\theta_i^1, q_2^{1S}; \Phi))]] = 0. \quad (33)$$

⁴⁸We drop observations $a_{1i} = 100$ since we cannot impute a unique value of q_{2i}^{2S} on the basis of the observed a_{2i} .

Given a specification for $h(\cdot)$, we estimate θ by applying GMM to the unconditional moment conditions (31) - (33).

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 $h(\cdot)$ 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

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

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

Since we discretize the action space of a_{1i} into three intervals ("bins") $B_1 = (50, 66.7]$, $B_2 = (66.7, 83.4]$ and $B_3 = (83.4, 100]$ and the theoretical distribution of α varies 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 Φ and the signal precisions (q_2^{1S}, q_2^{2S}) , the theoretical α 's can be written as

$$\text{time 2.1:} \quad \alpha_{2i}^1(\Phi, q_2^{1S}) = \frac{l(a_{2i}^{1*}(\theta_i^1, q_2^{1S}; \Phi)/100)}{l(0.7)},$$

$$\begin{aligned}
\text{time 2.2-confirming} & : \\
\alpha_{2i,conf}^2(\Phi, q_2^{1S}, q_2^{2S}) & = \frac{l(a_{2,L}^2(\theta_{2i,conf}^2, q_2^{1S}, q_2^{2S}; \Phi)/100) - l(a_2^{1*}(\theta_i^1, q_2^{1S}; \Phi)/100)}{l(0.7)} \\
& = \frac{l(q_2^S) + l(a_2^{1*}(\theta_{2i,conf}^2, q_2^{1S}; \theta)/100) - l(a_2^{1*}(\theta_i^1, q_2^{1S}; \Phi)/100)}{l(0.7)},
\end{aligned}$$

$$\begin{aligned}
\text{time 2.2-contradicting} & : \\
\alpha_{2i,cont}^2(\Phi, q_2^{1S}, q_2^{2S}) & = \frac{l(a_{2,L}^2(\theta_{2i,cont}^2, q_2^{1S}, q_2^{2S}; \Phi)/100) - l(a_2^{1*}(\theta_i^1, q_2^{1S}; \Phi)/100)}{l(0.3)} \\
& = \frac{l(1 - q_2^S) + l(a_2^{1*}(\theta_{2i,cont}^2, q_2^{1S}; \Phi)/100) - l(a_2^{1*}(\theta_i^1, q_2^{1S}; \Phi)/100)}{l(0.3)}.
\end{aligned}$$

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

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

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

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

$$\underbrace{\mathbf{m}_{1i}^L(\Phi)}_{9 \times 1} = \begin{pmatrix} 1\{a_{1i} \in B_1\} \cdot \begin{pmatrix} 1\{\alpha_{2i}^1 \leq t_{2,0.25,B_1}^1\} - E_Q(1\{\alpha_{2i}^1(\Phi, q_2^{1S}) \leq t_{2,0.25,B_1}^1\}) \\ 1\{\alpha_{2i}^1 \leq t_{2,0.5,B_1}^1\} - E_Q(1\{\alpha_{2i}^1(\Phi, q_2^{1S}) \leq t_{2,0.5,B_1}^1\}) \\ 1\{\alpha_{2i}^1 \leq t_{2,0.75,B_1}^1\} - E_Q(1\{\alpha_{2i}^1(\Phi, q_2^{1S}) \leq t_{2,0.75,B_1}^1\}) \end{pmatrix} \\ \vdots \\ 1\{a_{1i} \in B_3\} \cdot \begin{pmatrix} 1\{\alpha_{2i}^1 \leq t_{2,0.25,B_3}^1\} - E_Q(1\{\alpha_{2i}^1(\Phi, q_2^{1S}) \leq t_{2,0.25,B_3}^1\}) \\ 1\{\alpha_{2i}^1 \leq t_{2,0.5,B_3}^1\} - E_Q(1\{\alpha_{2i}^1(\Phi, q_2^{1S}) \leq t_{2,0.5,B_3}^1\}) \\ 1\{\alpha_{2i}^1 \leq t_{2,0.75,B_3}^1\} - E_Q(1\{\alpha_{2i}^1(\Phi, q_2^{1S}) \leq t_{2,0.75,B_3}^1\}) \end{pmatrix} \end{pmatrix}, \quad (34)$$

and $\mathbf{m}_{2i,conf}^L(\Phi)$ and $\mathbf{m}_{2i,cont}^L(\Phi)$ are 9×1 vectors of moment conditions concern-

ing the cdfs of α_{2i}^2 for confirming and contradicting signal cases, respectively:

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

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

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

and, in the second step, we solve

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

where

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

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

C.2.2 Estimating the BU model

The BU model is a special case of the LRTU model in which $\theta_* = \theta^* = \theta$. In this case c becomes an irrelevant parameter, and the only parameter to estimate is $\Phi = \theta \in [0, 1]$. Furthermore, note that the theoretical α_{2i}^2 is given by $s_{2i}l(q_2^{2S}) + (1 - s_{2i})l(1 - q_2^{2S})$ (which is independent of the parameters) when $\theta_* = \theta^* = \theta$. Hence, the identifying information for Φ 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

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

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

$$\hat{\Phi} = \arg \min_{\Phi} \left(\sum_{i=1}^n \mathbf{m}_{1i}^B(\Phi) \right)' \left(\sum_{i=1}^n \mathbf{m}_i^B(\Phi) \right). \quad (35)$$

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

$$\begin{aligned} \hat{\Phi}_{GMM} &= \arg \min_{\Phi} \left(\frac{1}{n} \sum_{i=1}^n \mathbf{m}_{1i}^B(\Phi) \right)' \hat{W}^{-1} \left(\frac{1}{n} \sum_{i=1}^n \mathbf{m}_i^B(\Phi) \right), \quad (36) \\ \hat{W} &= \frac{1}{n} \sum_{i=1}^n \mathbf{m}_i^B(\hat{\Phi}) \mathbf{m}_i^B(\hat{\Phi})' \text{ with } \hat{\Phi} = \hat{\theta}. \end{aligned}$$

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

C.2.3 Estimation Results

For convenience, Table C.1 reports the results of the second stage GMM estimation (non-parametric bootstrapped standard errors in parenthesis).

The estimated proportion of noise type subjects in the BU model is $\theta = 0.3$. This 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.

It should be mentioned that $\theta = 0.3$ 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

Model	θ	θ_*	θ^*	c
BU	0.30 (0.052)			
LRTU		0 (0.016)	0.30 (0.045)	[1.65, 1.71] (0.071)

Table C.1: 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 the lower bound of the set of maximum likelihood estimates.

subjects at time 2 did not have rational expectations on the proportion of noise and rational predecessors.

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 : $c \in [1.65, 1.71]$. Nevertheless, the other parameters have the same estimate for any $c \in [1.65, 1.71]$. This parameter c co-determines the thresholds to trust or not the predecessor. It is clear that the inequalities in (26), (28), (30) may be satisfied for a set of parameter values. The estimates show 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 $\theta_* = 0$). When, instead, the threshold is not reached, he updates as if the probability of a noise predecessor were $\theta^* = 0.3$. 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’ 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 $\theta^* = 0.3$); they do trust an action $a_{1i} \in (83.4, 100]$. Let us consider now a_{2i}^2 . In this case, the decision to trust or not the predecessor depends on the subjective precisions of signals, as one can notice from (27) and (29). 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 $a_{1i} \in (66.7, 83.4]$.⁴⁹ They keep trusting a more extreme action $a_{1i} \in (83.4, 100]$. After receiving a contradicting signal, they keep not trusting an action $a_{1i} \in (50, 66.7]$ or $a_{1i} \in (66.7, 83.4]$, and in 68.9% of the cases they stop trusting an action $a_{1i} \in (83.4, 100]$.

⁴⁹This 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.

C.2.4 Resampling-based model comparison

We now test whether the LRTU model provides a better explanation for the observed behavior than the BU model. 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 degrees of freedom (has more parameters) than the other. There is a large literature on model specification tests 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) for one parameter of the LRTU we have multiple maxima; and (iii) the LRTU nests the BU model at the boundary of the parameter space (e.g., $\theta_* = \theta^*$). 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 the evidence against a null model without relying on an asymptotic theory (at the cost of being computationally intensive).

We set the null model to the BU model with parameter value $\theta = 0.3$ (as reported in Table C.2.3). As usual, we sample $(q_{2i}^{1S}, q_{2i}^{2S})$ 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 $(q_{2i}^{1S}, q_{2i}^{2S})$ 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 $\hat{\Phi}_{GMM}$ for the actual data. We keep this weighting matrix fixed across the 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 $j = 1, \dots, 1000$. Note that $\Delta\hat{L}^j$ 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 $\Delta\hat{L} = \hat{L}_B - \hat{L}_L$ as the difference between the minimized GMM objective functions of the BU and LRTU models for our dataset. We use the same weighting matrix as the one used to compute $\Delta\hat{L}^j$, $j = 1, \dots, 1000$. 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 so-computed p -value is 0.005, 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 gains 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).

C.3 The Full Bayesian Updating (FBU) model and its estimation

Finally, we illustrate the estimation of the FBU models. We first consider the FBU model with maxmin preferences, axiomatized by Pires (2002).

Before we go into the formal analysis, let us consider an illustrative example. Suppose a subject at time 2 has a prior θ ranging from $\theta_* = 0$ to $\theta^* = 1$. Suppose the predecessor chooses $a_1 = 70$. The subject updates his belief on the value of the good using each prior $\theta \in [0, 1]$. 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 $[0.3, 0.5]$. 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 $\alpha_2^2 = 1$. Note that in this example this updating rule implies 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.

Suppose a subject at time 2 starts with a set of priors $[\theta_*, \theta^*]$ on the proportion of noise type subjects at time 1. The subject applies Bayes's rule for each prior θ_{2i}^1 in $[\theta_*, \theta^*]$ and obtains a belief

$$p_{2i}^1(\theta, q_{2i}^{1S}) \equiv \Pr(V = 1 | a_{1i}; \theta, q_{2i}^{1S}) = \frac{(1 - \theta)q_{2i}^{1S}f(a_{1i}) + g(a_{1i})\theta}{(1 - \theta)f(a_{1i}) + 2g(a_{1i})\theta}, \quad (37)$$

where $f(\cdot)$ and $g(\cdot)$ are the distributions of actions of rational and noise-type agents, respectively, as specified in the main text. In FBU, the range of θ spans a range of beliefs on the value of the good being 100: $[p_{2i}^1(\theta^*, q_{2i}^{1S}), p_{2i}^1(\theta_*, q_{2i}^{1S})]$.

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

$$[l(p_{2i}^2(\theta^*, q_{2i}^{1S}, q_{2i}^{2S})), l(p_{2i}^2(\theta_*, q_{2i}^{1S}, q_{2i}^{2S}))] = [l(p_{2i}^1(\theta^*, q_{2i}^{1S})), l(p_{2i}^1(\theta_*, q_{2i}^{1S}))] + l(q_{2i}^{2S}), \quad (38)$$

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

$$[l(p_{2i}^2(\theta^*, q_{2i}^{1S}, q_{2i}^{2S}), l(p_{2i}^2(\theta_*, q_{2i}^{1S}, q_{2i}^{2S})))] = [l(p_{2i}^1(\theta^*, q_{2i}^{1S}), l(p_{2i}^1(\theta_*, q_{2i}^{1S})))] + l(1 - q_{2i}^{2S}). \quad (39)$$

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

$$a_{2i, F}^1 = a_2^{1*}(\theta^*, q_{2i}^{1S}) = 100p_{2i}^1(\theta^*, q_{2i}^{1S}).$$

Similarly, the second action is

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

C.3.1 Estimating the FBU model

Adopting the GMM approach used to estimate LRTU and BU models in the main text, we can estimate the unknown parameters $\Phi = (\theta_*, \theta^*)$, $0 \leq \theta_* \leq \theta^* \leq 1$, in the FBU model. Since we only consider the realization of q_{2i}^{1S} 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 (37)), and the maximin action $a_{2i, \max \min}^1$ is the Bayes's action with the implied prior θ^* . Hence, the moment conditions for the FBU model concerning the cdf of α_2^1 are obtained by replacing $\alpha_{2i}^1(\Phi, q_{2i}^{1S})$ in (34) with

$$\alpha_{2i, F}^1(\Phi, q_{2i}^{1S}) = \frac{l(a_2^{1*}(\theta^*, q_{2i}^{1S})/100)}{l(0.7)}.$$

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

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 $(q_{2i}^{1S}, q_{2i}^{2S})$. Nevertheless, given $(\Phi, q_{2i}^{1S}, q_{2i}^{2S})$, the maxmin action can be pinned down according to the formula of $a_{2i, F}^2$ given above. Let us denote by $a_{2i, F, conf}^2(\Phi, q_{2i}^{1S}, q_{2i}^{2S})$ the second action of ambiguity averse subject 2 predicted by FBU, given precisions $(q_{2i}^{1S}, q_{2i}^{2S})$ and in response to the confirming signal. Similarly, let $a_{2i, F, cont}^2(\Phi, q_{2i}^{1S}, q_{2i}^{2S})$ be his

action given precisions (q_2^{1S}, q_2^{2S}) and in response to the contradicting signal. We accordingly obtain the weights for the ambiguity averse agents as follows:

$$\begin{aligned} \text{time 2.2-confirming} & : \\ \alpha_{2i,F,conf}^2(\Phi, q_2^{1S}, q_2^{2S}) & \equiv \frac{l(a_{2i,F,conf}^2(\Phi, q_2^{1S}, q_2^{2S})/100) - l(a_2^{1*}(\theta^*, q_2^{1S}; \Phi)/100)}{l(0.7)}, \end{aligned}$$

$$\begin{aligned} \text{time 2.2-contradicting} & : \\ \alpha_{2i,F,cont}^2(\Phi, q_2^{1S}, q_2^{2S}) & = \frac{l(a_{2i,F,cont}^2(\Phi, q_2^{1S}, q_2^{2S})/100) - l(a_2^{1*}(\theta^*, q_2^{1S}; \Phi)/100)}{l(0.3)}. \end{aligned}$$

We then obtain the moment conditions concerning the cdfs of α_{2i}^2 by

$$\begin{aligned} & \underbrace{\mathbf{m}_{2,conf,i}^F(\Phi)}_{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\left(1\{\alpha_{2i,F,conf}^2(\Phi, q_2^{1S}, q_2^{2S}) \leq t_{2,conf,0.25,B_1}^2\}\right) \\ 1\{\alpha_{2i}^2 \leq t_{2,conf,0.5,B_1}^2\} - E_Q\left(1\{\alpha_{2i,F,conf}^2(\Phi, q_2^{1S}, q_2^{2S}) \leq t_{2,conf,0.5,B_1}^2\}\right) \\ 1\{\alpha_{2i}^2 \leq t_{2,conf,0.75,B_1}^2\} - E_Q\left(1\{\alpha_{2i,F,conf}^2(\Phi, q_2^{1S}, q_2^{2S}) \leq t_{2,conf,0.75,B_1}^2\}\right) \\ \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\left(1\{\alpha_{2i,F,conf}^2(\Phi, q_2^{1S}, q_2^{2S}) \leq t_{2,conf,0.25,B_3}^2\}\right) \\ 1\{\alpha_{2i}^2 \leq t_{2,conf,0.5,B_3}^2\} - E_Q\left(1\{\alpha_{2i,F,conf}^2(\Phi, q_2^{1S}, q_2^{2S}) \leq t_{2,conf,0.5,B_3}^2\}\right) \\ 1\{\alpha_{2i}^2 \leq t_{2,conf,0.75,B_3}^2\} - E_Q\left(1\{\alpha_{2i,F,conf}^2(\Phi, q_2^{1S}, q_2^{2S}) \leq t_{2,conf,0.75,B_3}^2\}\right) \end{pmatrix} \end{pmatrix} \\ & \underbrace{\mathbf{m}_{2i,cont}^F(\Phi)}_{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\left(1\{\alpha_{2i,F,cont}^2(\Phi, q_2^{1S}, q_2^{2S}) \leq t_{2,cont,0.25,B_1}^2\}\right) \\ 1\{\alpha_{2i}^2 \leq t_{2,cont,0.5,B_1}^2\} - E_Q\left(1\{\alpha_{2i,F,cont}^2(\Phi, q_2^{1S}, q_2^{2S}) \leq t_{2,cont,0.5,B_1}^2\}\right) \\ 1\{\alpha_{2i}^2 \leq t_{2,cont,0.75,B_1}^2\} - E_Q\left(1\{\alpha_{2i,F,cont}^2(\Phi, q_2^{1S}, q_2^{2S}) \leq t_{2,cont,0.75,B_1}^2\}\right) \\ \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\left(1\{\alpha_{2i,F,cont}^2(\Phi, q_2^{1S}, q_2^{2S}) \leq t_{2,cont,0.25,B_3}^2\}\right) \\ 1\{\alpha_{2i}^2 \leq t_{2,cont,0.5,B_3}^2\} - E_Q\left(1\{\alpha_{2i,F,cont}^2(\Phi, q_2^{1S}, q_2^{2S}) \leq t_{2,cont,0.5,B_3}^2\}\right) \\ 1\{\alpha_{2i}^2 \leq t_{2,cont,0.75,B_3}^2\} - E_Q\left(1\{\alpha_{2i,F,cont}^2(\Phi, q_2^{1S}, q_2^{2S}) \leq t_{2,cont,0.75,B_3}^2\}\right) \end{pmatrix} \end{pmatrix} \end{aligned}$$

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

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

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

C.3.2 Estimation Results

The optimal two-step GMM estimates show that $\hat{\theta}_* = \hat{\theta}^* = 0.30$ with their (non-parametric bootstrap) standard errors estimated as 0.070 and 0.067, respectively. That is, under FBU specification, the estimates indicate that subject 2 does not have multiple priors for the predecessor’s rationality. Furthermore, if we compare this result to that for the BU model, we note that the FBU model’s estimates coincide with the estimated BU, despite the fact that FBU model has an extra degree of freedom. In other words, adding multiple priors and assuming FBU does not improve the fitness of more restrictive BU model.

So far we have estimated the FBU model joint with maxmin preferences. 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 $[\theta_*, \theta^*]$ and in this sense we did not constrain our subjects to be overly pessimistic (as it would have been the case had we imposed $\theta^* = 1$). 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 $\theta_* = 0$, $\theta^* = 1$ 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 $\lambda = 0.45$, $\theta_* = 0.08$ and $\theta^* = 0.52$, indicating some form of optimism. Using the same test for model selection explained above, we obtain a p -value of 0.59: that is, this model does not fit the data significantly better than the BU model.⁵⁰

A different approach to the problem would be to use the principle of indifference or insufficient reason. According to this principle, typically attributed

⁵⁰Another 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.

to Jacob Bernoulli or Laplace, in the absence of a convincing reason, the subject 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 θ that the predecessor is noise, would simply use a uniform as a distribution of θ . 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 θ .

D The GC Treatment

Recall that in the GC Treatment we elicit subjects' beliefs about the predecessor's signal, rather than about the value of the good. Formally, agent 1 still observes the signal s_1 and takes an action a_1 to maximize $100 - 0.01(V - a_1)^2$. Agent 2, instead, after observing a_1 chooses a first action x_2^1 to maximize the quadratic payoff $100 - 0.01(100s_1 - x_2^1)^2$. He then observes the private signal s_2 and takes a second action x_2^2 to maximize a payoff equal to $100 - 0.01(100s_1 - x_2^2)^2$. That is, agent 2 is asked to report what he thinks about the binary signal received by agent 1, first based on agent 1's action a_1 only, then after additionally receiving s_2 .

Hence, in this model specification, agent 2 chooses his optimal action $x_2^{j*}/100 = \Pr(s_1 = 1|I_2^j)$, whereas nothing changes for agent 1 compared to the discussion in the main text. In the PBE, agent 2 chooses $x_2^1 = 100$ after observing $a_1 = 70$ and $x_2^1 = 0$ after observing $a_1 = 30$, since action $a_1 = 70$ (respectively, 30) perfectly reveals that the signal observed by agent 1 is $s_1 = 1$ (respectively, $s_1 = 0$). After observing the private signal s_2 , the agent does not update his degenerate beliefs on the signal observed by agent 1 and so, after observing $a_1 = 70$, $x_2^2/100 = \Pr(s_1|a_1 = 70, s_2 = 1) = \Pr(s_1|a_1 = 70, s_2 = 0) = 1$ and after observing $a_1 = 30$, $x_2^2 = \Pr(s_1|a_1 = 30, s_2 = 1) = \Pr(s_1|a_1 = 30, s_2 = 0) = 0$.⁵¹

The same analysis carries through if, as in the main text, we extend the model to allow for subjective precisions.

After observing an action greater than 50, agent 2 chooses his first action as follows:

$$\frac{x_2^1}{100} =: \Pr(s_1 = 1|a_1) = \tag{40}$$

$$\Pr(s_1 = 1|a_1, \tau_r) \Pr(\tau_r|a_1) + \Pr(s_1 = 1|a_1, \tau_i) \Pr(\tau_i|a_1) = \Pr(\tau_r|a_1) = \mu_1(a_1).$$

Similarly, in the case in which $a_1 < 50$,

$$\frac{x_2^1}{100} = \Pr(s_1 = 1|a_1) = 1 - \mu_1(a_1).$$

Therefore, his first action simply reflects his belief $\mu_1(a_1)$.

⁵¹Actions a_1 other than 30 and 70 cannot arise, and one need not specify the actions of agent 2 after such off-the-equilibrium-path actions.

Regarding agent 2's second action, let us stick to the case in which $a_1 > 50$ and discuss how the agent updates his beliefs and chooses x_2^2 after observing a “contradicting” signal $s_2 = 0$ or a “confirming” signal $s_2 = 1$.⁵² In the case of $s_2 = 0$, with a slight abuse of notation, we have

$$\begin{aligned} \frac{x_2^2(a_1, s_2 = 0)}{100} &=: \Pr(s_1 = 1 | a_1, s_2 = 0) = \\ &\Pr(s_1 = 1 | a_1, s_2 = 0, \tau_r) \Pr(\tau_r | a_1, s_2 = 0) + \\ &\Pr(s_1 = 1 | a_1, s_2 = 0, \tau_i) \Pr(\tau_i | a_1, s_2 = 0) = \\ &\Pr(\tau_r | a_1, s_2 = 0) := \mu_2(a_1, s_2 = 0), \end{aligned}$$

and, after some algebra, we obtain

$$\begin{aligned} \frac{x_2^2(a_1, s_2 = 0)}{100} &= \Pr(\tau_r | a_1, s_2 = 0) = & (41) \\ &\frac{[q_2^{1S}(1 - q_2^{2S}) + q_2^{2S}(1 - q_2^{1S})] \mu_1(a_1)}{[q_2^{1S}(1 - q_2^{2S}) + q_2^{2S}(1 - q_2^{1S})] \mu_1(a_1) + [(q_2^{1S} q_2^{2S}) + (1 - q_2^{1S})(1 - q_2^{2S})] (1 - \mu_1(a_1))}. \end{aligned}$$

Similar computations show that

$$\begin{aligned} \frac{x_2^2(a_1, s_2 = 1)}{100} &= \Pr(\tau_r | a_1, s_2 = 1) = & (42) \\ &\frac{[(q_2^{1S} q_2^{2S}) + (1 - q_2^{1S})(1 - q_2^{2S})] \mu_1(a_1)}{[(q_2^{1S} q_2^{2S}) + (1 - q_2^{1S})(1 - q_2^{2S})] \mu_1(a_1) + [q_2^{1S}(1 - q_2^{2S}) + q_2^{2S}(1 - q_2^{1S})] (1 - \mu_1(a_1))}. \end{aligned}$$

An agent who believes that the predecessor is certainly rational ($\mu_1(a_1) = 1$) or certainly irrational ($\mu_1(a_1) = 0$) does not update x_2^1 (since $\frac{x_2^1}{100} = \mu_1(a_1)$), of course. Otherwise, the agent updates up this probability after observing a confirming signal and down after a contradicting signal.⁵³

The updating on the rationality of the predecessor can also be expressed as

$$\frac{\mu_2(a_1, s_2)}{1 - \mu_2(a_1, s_2)} =: \frac{\Pr(\tau_r | a_1, s_2)}{\Pr(\tau_i | a_1, s_2)} = \frac{\Pr(s_2 | a_1, \tau_r)}{\Pr(s_2 | a_1, \tau_i)} \frac{\mu_1(a_1)}{1 - \mu_1(a_1)}. \quad (43)$$

Although the agent updates on the rationality of the predecessor, still, as explained, the updating on the value of the good follows expression (4) in the main text. To see how the two updating procedures fit together, observe that

⁵²The analysis for the case in which $a_1 < 50$ is analogous.

⁵³To understand why, it is perhaps useful to consider the limit case in which the signal is almost perfectly informative. In such a case, after observing the signal, the agent knows (almost with certainty) the value of the good and, therefore, that the predecessor received his same signal. If the predecessor chose an action that the signal contradicts, it must be that he is irrational (since the probability of him receiving the other signal is almost zero). If the predecessor chose an action that the signal confirms, it must be that he is rational (since the probability of him receiving the same signal is almost 1).

the posterior likelihood ratio on the value of the good (expression (2) in the main text) can be expressed as

$$\frac{\Pr(V = 100|a_1, s_2)}{\Pr(V = 0|a_1, s_2)} = \frac{\Pr(V = 100|a_1, s_2, \tau_r) \Pr(\tau_r|a_1, s_2) + \Pr(V = 100|a_1, s_2, \tau_i) \Pr(\tau_i|a_1, s_2)}{\Pr(V = 0|a_1, s_2, \tau_r) \Pr(\tau_r|a_1, s_2) + \Pr(V = 0|a_1, s_2, \tau_i) \Pr(\tau_i|a_1, s_2)} = \quad (44)$$

$$\frac{1 - q_2^{2S} \Pr(V = 100|a_1, \tau_r) + \Pr(V = 100|a_1, \tau_i) \frac{\Pr(s_2|a_1, \tau_r) \frac{1 - \mu_2(a_1, s_2)}{\Pr(s_2|a_1, \tau_i) \mu_2(a_1, s_2)}}{q_2^{2S}}}{\Pr(V = 0|a_1, \tau_r) + \Pr(V = 0|a_1, \tau_i) \frac{\Pr(s_2|a_1, \tau_r) \frac{1 - \mu_2(a_1, s_2)}{\Pr(s_2|a_1, \tau_i) \mu_2(a_1, s_2)}}{q_2^{2S}}}.$$

Similarly,

$$\begin{aligned} \frac{\Pr(V = 100|a_1)}{\Pr(V = 0|a_1)} &= \frac{\Pr(V = 100|a_1, \tau_r) \mu_1(a_1) + \Pr(V = 100|a_1, \tau_i) (1 - \mu_1(a_1))}{\Pr(V = 0|a_1, \tau_r) \mu_1(a_1) + \Pr(V = 0|a_1, \tau_i) (1 - \mu_1(a_1))} = \\ &= \frac{\Pr(V = 100|a_1, \tau_r) + \Pr(V = 100|a_1, \tau_i) \frac{1 - \mu_1(a_1)}{\mu_1(a_1)}}{\Pr(V = 0|a_1, \tau_r) + \Pr(V = 0|a_1, \tau_i) \frac{1 - \mu_1(a_1)}{\mu_1(a_1)}}. \end{aligned} \quad (45)$$

After substituting (43) into (45) and then (45) into (44), one re-obtains expression (4) of the main text.

In other words, in a Bayesian setting, while it is true that agent 2 does update on the rationality of the predecessor after observing his private signal, his posterior belief on the value of the good follows expression (4), as results from the statistical property that conditional on the value of the good, the signal at $t = 2$ is independent of the action at $t = 1$

D.1 Results of the GC Treatment

We now present all the results for the GC treatment. Our goal is to study whether subjects change their mind about the predecessors' rationality in a Bayesian fashion (or not).

Figure D.1 reports the frequencies of the stated beliefs after observing the predecessor's action only (x_2^1). Similarly to Figure A.2, we have pooled together all the cases in which the observed choice at time 1 was greater than 50, all the cases in which it was lower than 50 and, for completeness, all the cases in which it was equal to 50.⁵⁴ In the vast majority of the cases subjects do believe the action reflects the signal the predecessor received. In 30% of the cases, they attribute probability 1 to the predecessor having received the signal $s_1 = 1$ (or $s_1 = 0$) after observing an action greater (lower) than 50. In more than 30% of the cases, the belief, although not 1, differs from certainty by at most 0.15 (i.e., upon observing an action greater than 50, the subject states at least 85). In less than 5% of the cases, subjects do not update at all, reporting 50, after observing an action different from 50, and in some rare instances they update in the wrong direction.

⁵⁴Recall that at time 1 the treatment was identical to the SL treatments. For this reason, in the interest of space, we refrain from reporting the results about action 1. Mann-Whitney U tests on the medians show that we cannot reject the hypothesis that the decisions at time 1 are the same as in the SL1, SL2, SL3 and IDM treatments (p -values: 0.42, 0.84, 0.11, 0.34).




Figure D.1: Distribution of first decisions (beliefs on the predecessor’s signal) at time 2 (the top panel refers to $a_1 > 50$, the middle to $a_1 = 50$ and the bottom to $a_1 < 50$).

In Table D.1 we describe how $x_2^1/100$ varies with a_1 . The median probability attached to the predecessor receiving the good signal increases with a_1 , in line with our findings from the SL treatments.

	1st Quartile	Median	3rd Quartile
$x_2^1/100$	0.70	0.90	1
$x_2^1/100$ (upon observing $50 < a_1 \leq 66.7$)	0.59	0.70	0.91
$x_2^1/100$ (upon observing $66.7 < a_1 \leq 83.4$)	0.75	0.90	1.00
$x_2^1/100$ (upon observing $a_1 > 83.4$)	0.88	0.95	1.00

Table D.1: Distribution of first decisions at time 2 in the GC Treatment.

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

For the reader’s convenience, we report here the analysis summarized in the main text about the distributions of beliefs after the subject has also received the private signal (x_2^2). Figure D.2 (equivalent to Figure 6.1 of the main text) reports these distributions of beliefs, distinguishing the cases of contradicting and confirming signals.⁵⁵ Analogously to the previous analysis, the figure is ob-

⁵⁵The signals are defined as confirming or contradicting with respect to a_1 . There are no relevant differences when the definition of confirming or contradicting is with respect to x_2^1 .




Figure D.2: Distribution of second decisions (beliefs on the predecessor’s signal) at time 2. The top panel refers to confirming signals. The bottom panel refers to contradicting signals. In each panel, we report the frequency of the actual decisions (left bars) and of the simulated, theoretical ones (right bars), using objective signal precisions.

tained after transforming an action $x_2^1 < 50$ into $100 - x_2^1$ and the corresponding signal s_1 into $1 - s_1$. For each case, we plot two frequencies. The left bar in each bin shows the empirical frequency of x_2^2 ; the right bar, instead, shows the frequency of the theoretical beliefs (denoted by \hat{x}_2^2) that subjects would have reported, had they updated in a Bayesian way (with signal precisions equal to the objective precisions). In other words, for each reported first belief x_2^1 we have computed the corresponding updated belief, \hat{x}_2^2 , from expressions (41) and (42) using $q_2^{1S} = q_2^{2S} = 0.7$. We also computed the updated beliefs using subjective precisions q_2^{1S} and q_2^{2S} , and the results are very similar: see Figure D.3.

The empirical and theoretical frequencies are remarkably different in the contradicting signal case, whereas they are remarkably similar in the confirming signal case. For the contradicting case, the theoretical belief distribution first-order stochastically dominates the empirical belief distribution: subjects’ beliefs about their predecessors’ signals (equivalent to their predecessors’ rationality, according to the definitions of Section 2) shift more drastically than what Bayesian updating predicts. Note, in particular, that the mode of the empirical distribution is 50, that is, after a contradicting signal, subjects give equal chance




Figure D.3: Distribution of second decisions (beliefs on the predecessor's signal) at time 2. The top panel refers to confirming signals. The bottom panel refers to contradicting signals. In each panel, we report the frequency of the actual decisions (left bars) and of the simulated ones (right bars). The simulated decisions are based on subjective signal precisions.

to the predecessor having received a green or red ball.

For confirming signals, the theoretical, Bayesian, belief distribution matches the empirical belief distribution remarkably well, indicating that subjects update their beliefs in line with Bayesian updating. We performed two-sample Kolmogorov-Smirnov tests and we reject the null hypothesis of equivalence of the distributions (p -value = 0) for the contradicting signal case, whereas we cannot reject the null at 10% significance level for the confirming signal case (p -value = 0.13).

To quantify the departure from Bayesian updating in the GC Treatment, we also defined β_2^2 , a subject's weight on the signals, as in the following expressions:

$$\frac{x_2^2}{100 - x_2^2} = \left(\frac{0.7^2 + 0.3^2}{2(0.7)(0.3)} \right)^{\beta_2^2} \frac{x_2^1}{100 - x_2^1}, \quad (46)$$

when the subject observed $s_2 = 1$ and, analogously,

$$\frac{x_2^2}{100 - x_2^2} = \left(\frac{2(0.7)(0.3)}{0.7^2 + 0.3^2} \right)^{\beta_2^2} \frac{x_2^1}{100 - x_2^1}, \quad (47)$$

when he observed $s_2 = 0$.

For a given x_2^1 , the two formulas are equivalent to Bayesian updating when $\beta_2^2 = 1$ (and indicate an overweight of the signals for $\beta_2^2 > 1$ and underweight for $\beta_2^2 < 1$).⁵⁶ Table D.2 reports the results.

	1st Quartile	Median	3rd Quartile
β_2^2	0	2.31	5.48
β_2^2 (upon observing confirming signal)	0	1.33	4.63
β_2^2 (upon observing contradicting signal)	0	2.86	6.07

Table D.2: Distribution of β in the GC treatment.

The table shows the quartiles of the distribution of β in the GC treatment. The data refer to all cases in which the first decision at time 2 was different from 0 or 100.

In the case of a confirming signal the median action coincides with the Bayesian one, confirming the finding highlighted in Figure 4.3. In the case of a contradicting signal, the weight for the median action is considerably higher, 2.63.

Overall, the results of this treatment show how subjects change their view about the predecessors' rationality in a direct way. The changes are not in line with Bayesian updating, but in line with the LRTU, the theoretical model we proposed to explain the asymmetric updating observed in the SL treatments.

Finally, for the sake of comparison with the SL treatments, we simulated the a_2^1 and a_2^2 resulting from this treatment, using the same subjective precisions, as

⁵⁶Clearly, the parameter β_2^2 does not have the same interpretation as α_2^2 . Moreover, note that the value of β_2^2 is undetermined when $x_2^1 = 100$, therefore we exclude these cases. When $x_2^2 = 100$ we use the same approximation as previously discussed.




Figure D.4: Distribution of simulated first actions (a_2^1) at time 2 (the top panel refers to $a_1 > 50$, the middle to $a_1 = 50$ and the bottom to $a_1 < 50$).

elicited in the SL and IDM treatments. Let us start with a_2^1 . For, e.g., $a_1 > 50$, we can write it as

$$\begin{aligned}
 a_2^1 &= \Pr(V = 100|a_1) = & (48) \\
 \Pr(V = 100|a_1, s_1 = 1) \Pr(s_1 = 1|a_1) + \Pr(V = 100|a_1, s_1 = 0) \Pr(s_1 = 0|a_1) &= \\
 q_2^{1S} \frac{x_2^1}{100} + (1 - q_2^{1S}) \left(1 - \frac{x_2^1}{100}\right). &
 \end{aligned}$$

To simulate these actions, we used the subjective precisions q_2^{1S} as in the SL and IDM treatments (we described the eliciting procedure in detail in the Econometrics section above) and independent of x_2^1 . The simulated distributions, conditional on $a_1 < 50$, $a_1 = 50$, $a_1 > 50$ are shown in Figure D.4. They are remarkably similar to those in Figure A.2.

Similarly, in Figure D.5 we report the frequencies of the simulated a_2^2 by using the subjective precisions (q_2^{1S}, q_2^{2S}). The figure is again remarkably similar to Figure 4.3.




Figure D.5: Distribution of simulated a_2^2 for confirming (top panel) or contradicting (bottom panel) s_2 .

E Treatment with stochastic signal precision

We now present an extra SL treatment we also ran as a robustness check. This treatment differs from the SL treatments presented in the text in that the signal precision is random. Each time, the computer draws a signal precision from a uniform distribution on $[0.5, 1]$ in addition to the signal realization, that is, each agent t ($t = 1, 2$) receives a symmetric binary signal $s_t \in \{0, 1\}$ distributed as follows:

$$\Pr(s_t = 1 \mid V = 100) = \Pr(s_t = 0 \mid V = 0) = q_t,$$

with $q_t \sim U[0.5, 1]$. Conditional on the value of the good, the signals are independently (but not identically) distributed over time. The subject receiving the signal is informed of the signal precision. Both signal precision and signal realization are private information. We used groups of 4 subjects and randomly matched them in pairs at the beginning of each round (similarly to the procedure for the GC treatment). The experiment was repeated for 30 rounds and lasted, on average, more than 2 hours. The payoff for each round and the final payment were computed as in all other treatments. Table E.1 summarizes the main feature of the study.

Treatment	Signal Precision	Sequence	Subjects in a group	Groups	Participants	Rounds
SL4	(0.5, 1]	2	4	11	44	30

Table E.1: Treatment's features.

E.1 Results

At time 1, a Bayesian agent would choose $a_1 = 100q_1$ upon receiving a good signal and $a_1 = 100(1 - q_1)$ upon receiving a bad signal. To evaluate how subjects did in the laboratory, we used the following model:

$$a_{1i} = 100 \left(s_{1i} \frac{q_{1i}^{\alpha_{1i}}}{q_{1i}^{\alpha_{1i}} + (1 - q_{1i})^{\alpha_{1i}}} + (1 - s_{1i}) \frac{(1 - q_{1i})^{\alpha_{1i}}}{q_{1i}^{\alpha_{1i}} + (1 - q_{1i})^{\alpha_{1i}}} \right). \quad (49)$$

Table E.2 shows the quartiles of the distribution of weights on the signal subjects received. The median weight is 1, indicating that the median action was in line with Bayesian updating.

	1st Quartile	Median	3rd Quartile
α_{1i}	0.89	1.00	1.85

Table E.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.

As for time 2, let us consider the first action taken by the subjects. Similarly to the analysis in the main text, we illustrate the subjects' behavior by means

of the following simple model: when the subject observed $a_1 > 50$, he chose a_{2i}^1 such that

$$a_{2i}^1 = 100 \frac{q_{1i}^{\alpha_{2i}^1}}{q_{1i}^{\alpha_{2i}^1} + (1 - q_{1i})^{\alpha_{2i}^1}}; \quad (50)$$

analogously, when he observed $a_1 < 50$, he chose a_{2i}^1 such that

$$a_{2i}^1 = 100 \frac{(1 - q_{1i})^{\alpha_{2i}^1}}{q_{1i}^{\alpha_{2i}^1} + (1 - q_{1i})^{\alpha_{2i}^1}}. \quad (51)$$

Table E.3 reports the results. The weights for all the three quartiles are not dissimilar from those at time 1, suggesting that subjects typically followed the predecessor's action.⁵⁷

	1st Quartile	Median	3rd Quartile
α_2^1	0.77	1.21	2.81

Table E.3: Distribution of weights on private signal for actions at time 1.

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 equal to that of the signal received by subject 1) for the subject at time 2.

Finally, let us illustrate how subjects at time 2 updated their beliefs after receiving their private signal. We compute the weight that the subject puts on his signal by using the usual model of updating:

$$a_{2i}^2 = 100 \frac{q_{2i}^{\alpha_{2i}^2} \frac{a_{2i}^1}{100}}{q_{2i}^{\alpha_{2i}^2} \frac{a_{2i}^1}{100} + (1 - q_{2i})^{\alpha_{2i}^2} \left(1 - \frac{a_{2i}^1}{100}\right)}, \quad (52)$$

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

$$a_{2i}^2 = 100 \frac{(1 - q_{2i})^{\alpha_{2i}^2} \frac{a_{2i}^1}{100}}{(1 - q_{2i})^{\alpha_{2i}^2} \frac{a_{2i}^1}{100} + q_{2i}^{\alpha_{2i}^2} \left(1 - \frac{a_{2i}^1}{100}\right)}, \quad (53)$$

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

Table E.4 reports the results.⁵⁸

For all quartiles, the weight is higher for contradicting than for confirming signals. The median weight for the contradicting signal case, in particular, is

⁵⁷We ran a Mann-Whitney U test (Wilcoxon rank-sum test) on the median weight for the action by subject 1 and the first action by subject 2; we cannot reject the null hypothesis that their distribution is the same (p -value = 0.70).

⁵⁸The value of α_{2i}^2 is undetermined when $a_{2i}^1 = 100$, therefore we exclude these cases. When $a_{2i}^1 = 100$ we use the same approximation as previously discussed.

	1st Quartile	Median	3rd Quartile
α_2^2	0.33	1.27	3.65
α_2^2 (upon observing confirming signal)	0.00	0.78	3.22
α_2^2 (upon observing contradicting signal)	0.85	2.03	3.98

Table E.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. The data refer to all cases in which the first action at time 2 was different from 50.

2.03, as if the subject double counted the private signal.⁵⁹ This confirms the asymmetric updating result we illustrated in the main text.

⁵⁹Mann-Whitney U test (Wilcoxon rank-sum test) on the null hypothesis that the median weight for the confirming and contradicting signal follow the same distribution gives p -value = 0.10.

F Instructions for the SL Treatment

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.

F.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.

F.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.

F.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.

F.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.

F.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.

G Instructions for the GC Treatment

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

You are about to take part in a study on decision making. 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 and some luck you will earn a different amount of money. You will receive the money immediately after the experiment.

G.1 The experiment

The experiment consists of 15 rounds of decision making. There are several participants in the laboratory. In each round you will be randomly matched with one other participant only. You have the same chance of being matched with any other participant. You will be randomly re-matched with another participant at the beginning of every round.

A good with two possible values

In each round, you will have to make some predictions, but your task will be different depending on whether you happen to be the first or the second in the pair to make a decision. Before we explain this, let us start by saying that in each round there will be a good whose value is either 0 or 100 units of a fictitious currency called “lira.” Whether the good is worth 0 or 100 liras will be determined randomly at the beginning of each round: there will be a chance of 50% that the value is 0 and a chance of 50% that it is 100 liras, like in the toss of a fair coin. The computer chooses the value of the good in each round randomly. 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 receive some information about which value is more likely to have been chosen. The computer will use one of two “virtual urns” both containing green and red balls. The proportion of the two types of balls in each urn, however, is different. 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 be precise:

- If the value is 100, the proportion of green balls is 70% and the proportion of red balls is 30%, like in an urn with 70 GREEN balls and 30 RED balls.
- If the value is 0, the proportion of green balls is 30% and the proportion of red balls is 70%, like in an urn with 30 GREEN balls and 70 RED balls.

What you have to do

As we said, you will have to make some predictions.

If you are the first person to make a decision, we will ask you with which chance you believe the value of the good is 100. You will first observe a ball from the urn and then you will have to state the chance you attach to the value of the good being 100.

If you are the second person to make the decision, we will ask you to make two predictions. You will observe the decision made by the first person, and then we will ask you with which chance you believe the first person observed a green ball. After you have done this, you will observe a ball from the urn and then we will ask you again with which chance you believe the first person observed a green ball.

Let us now describe the precise procedures for each round.

G.2 Procedures for each round

The precise sequence of events is the following:

First: the computer program will randomly decide 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 now choosing 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. You have the same (50%) 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.

Fourth: after the person sees this information, (s)he has to choose a number between 0 and 100 (the chance that the value is 100). The decision made will be visible to the second person. Note that there will be two cells, one to input integers and one to input decimal points.

Fifth: the second person will be asked to make her/his prediction about the colour observed by the first person, choosing a number between 0 and 100 (the chance that the first person saw a green ball).

Sixth: after the decision, the computer will draw a ball from the “virtual urn” (always with 70% of one colour and 30% of the other) and inform (only) the second person of its color.

Seventh: the second person, after observing the ball color, will now make a new prediction about the colour, choosing again a number between 0 and 100 (the chance that the first person saw a green ball).

Eighth: the computer will reveal the value of the good for the round and the payoff you earned (for each decision you made) in the round.

G.3 Your per-decision payoff

Your earnings depend on how well you make your prediction. When you are the first decision maker, this means how well you predict the value of the good. 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 by 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 25.30. In this case you made a mistake of 74.70. We will give you $100 - 0.01 * 74.70^2 = 44.19$ 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. In this case, your payoff would be 0.

Exactly the same rules apply in case you are the second decision maker, except that, of course, now you will earn money depending on the prediction on the ball color. If you predict the color exactly, you will earn 100 liras. If you make a mistake 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 by a small amount; if your mistake is big, you will be penalized more than proportionally.

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

Example A: Suppose the true color observed by the predecessor is Green. Suppose you predict 65 (that is, you attach a chance of 65% to the green ball). In this case, you made a mistake of 35. We will give you $100 - 0.01 * 35^2 = 87.75$ liras.

Example B: Suppose the true value is red. Suppose you predict 10 (that is, you attach a chance of 10% to the green ball). In this case you made a mistake of 10. We will give you $100 - 0.01 * 10^2 = 99$ 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 (for decision maker 1) or the true color (for decision maker 2) is. Any other prediction will decrease the amount you can expect to earn. For instance, suppose you are the first decision maker and think there is a 90% chance that the value of the good is 100 and, hence, a 10% chance that the 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. Similarly, suppose you are the second decision maker and think there is a 45% chance that the color observed by the first decision maker is green and, hence, a 55% chance that the color is red. If this is your belief, to maximize your expected payoff, choose 45 as your prediction. Similarly, if you think the ball color is green with chance 33% and red with chance 67%, select 33.

G.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.

G.5 The final payment

To compute your payment, we will randomly choose three rounds. The computer will randomly choose one round among the first five, one among the rounds 6 – 10, and one among the last five rounds. If in the round you have made one choice, we will use the payoff for that decision. If instead you have made two decisions, we will randomly pick one of the two decisions and consider that for your payoff. We will sum the payoffs that you have obtained for these three decisions. We will then convert your payoffs into pounds at the exchange rate of 100 liras = £6. That is, for every 100 liras you earn, you will obtain 6 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.

H Instructions for the IDV2 Treatment

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

You are about to take part in a study on decision making. 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 and some luck you will earn a different amount of money. You will receive the money immediately after the experiment.

H.1 The experiment

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

A good with two possible values

In each round there will be a good whose value is either 0 or 100 units of a fictitious currency called “lira.” Whether the good is worth 0 or 100 liras will be determined randomly at the beginning of each round: there will be a chance of 50% that the value is 0 and a chance of 50% that it is 100 liras, like in the toss of a fair coin. The computer chooses the value of the good in each round randomly. The value of the good in one round never depends on the value of the good in any of the previous rounds.

Two urns

You will not be told the value of the good, but you will receive some information about which value is more likely to have been chosen. The computer will use one of two “virtual urns” both containing green and red balls. The proportion of the two types of balls in each urn, however, is different. If the value of the good is 0, the computer will draw a ball from an urn containing more red balls. If the value is 100, instead, the computer will draw a ball from an urn containing more green balls. To be precise:

- If the value is 100, the proportion of green balls is 70% and the proportion of red balls is 30%, like in an urn with 70 GREEN balls and 30 RED balls.
- If the value is 0, the proportion of green balls is 30% and the proportion of red balls is 70%, like in an urn with 30 GREEN balls and 70 RED balls.

Information about the value and decisions

As we said, you will have to make two decisions in each round. You will receive information before you have to make them.

Before your first decision, the computer will draw a ball from one of the two urns, as explained above. After drawing the ball, the computer will announce a color, either Green or Red. The computer will announce the color of the drawn ball, but sometimes, with a 25% chance, will announce the incorrect color. For instance, imagine the computer draws a green ball. You will see

the announcement “Green ball,” but there is also a 25% chance you will see the incorrect announcement “Red ball”. The announcement is, therefore, more likely than not to coincide with the true ball color, but sometimes it differs from the true ball color.

After you have observed this announcement, you will be asked with which chance you believe the value of the good is 100.

After your first decision, the computer will draw a ball from the same “virtual urn” as before (always with 70% chance of one color and 30% of the other) and show this ball to you on the screen. You will observe the ball color and be asked again with which chance you believe the value of the good is 100. Note that in this case you will directly observe the ball color, there is no possibility of an incorrect announcement.

Let us now describe the precise procedures for each round.

H.2 Procedures for each round

The precise sequence of events is the following:

First: the computer program will randomly decide 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 now choosing the value of the good by flipping a coin.”

Second: the computer will draw a ball from the “virtual urn.” Based on the ball color, it will make an announcement. You will see this announcement on the screen.

Third: after you see this information, you have to choose a number between 0 and 100 (the chance that the value is 100). Note that there will be two cells, one to input integers and one to input decimal places (e.g., if you want to state 34.05 input 34 in the first cell and 5 in the second; if you want to state 34.50 input 34 in the first cell and 50 in the second).

Fourth: after the decision, the computer will draw a ball from the same “virtual urn” (again with 70% chance of one colour and 30% of the other) and inform you of its color. You will see the ball color on the screen.

Fifth: you will be asked to choose a number between 0 and 100, the chance that the value of the good is 100.

Sixth: the computer will reveal the value of the good for the round and the payoff you earned (for each decision you made) in the round.

H.3 Your per-decision payoff

Your earnings depend on how well you make your prediction. If you predict the value of the good exactly (i.e., you state that there is a 100% chance the value is

100 and the value is indeed 100; or you state that there is a 0% chance the value is 100 and the value is indeed 0), 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 by 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 25.30. In this case you made a mistake of 74.70. We will give you $100 - 0.01 * 74.70^2 = 44.19$ 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. In this case, your payoff would be 0.

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 the 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.

H.4 The other rounds

The procedures described for the first round will be repeated 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.

H.5 The final payment

To compute your payment, we will randomly choose three rounds. The computer will randomly choose one round among the first five, one among the rounds 6 – 10, and one among the last five rounds. The computer will randomly pick one of the two decisions in the round and consider that for your payoff. We will sum the payoffs that you have obtained for these three decisions. We will then convert your payoffs into pounds at the exchange rate of 100 liras = £6. That is, for every 100 liras you earn, you will obtain 6 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.