

## Information Redundancy Neglect versus Overconfidence: A Social Learning Experiment<sup>†</sup>

By MARCO ANGRISANI, ANTONIO GUARINO, PHILIPPE JEHIEL,  
AND TORU KITAGAWA\*

*We study social learning in a continuous action space experiment. Subjects, acting in sequence, state their beliefs about the value of a good after observing their predecessors' statements and a private signal. We compare the behavior in the laboratory with the Perfect Bayesian Equilibrium prediction and the predictions of bounded rationality models of decision-making: the redundancy of information neglect model and the overconfidence model. The results of our experiment are in line with the predictions of the overconfidence model and at odds with the others'. (JEL C91, D12, D82, D83)*

Many economic decisions, from the most mundane ones, like the choice of a restaurant, to the most important ones, like the adoption of a new technology for a firm or the adoption of a new medical protocol for a physician, require making inferences about an underlying state of nature (e.g., which restaurant or technology or medical protocol is the best one). In many of these situations, economic agents have some private information about the state of nature and also have information about the choice of others (e.g., other diners, firms, doctors) who faced the same decision problem in the past. Being able to make inferences about the underlying state by using the information conveyed by others' decisions (which is referred to as “social learning”) may be very valuable but may also have some pathological effects, such as herding on incorrect choices.

An important aspect affecting how accurate the inference can be is the choice set from which agents can pick their actions. While with a discrete action space social learning typically leads to inefficient outcomes (“informational cascades”) (Banerjee 1992; Bikhchandani, Hirshleifer, and Welch 1992), with a continuous action space a sequence of rational decision-makers are typically able to infer the

\*Angrisani: Center for Economic and Social Research, University of Southern California (email: [angrisani@usc.edu](mailto:angrisani@usc.edu)); Guarino: Department of Economics, University College London (email: [a.guarino@ucl.ac.uk](mailto:a.guarino@ucl.ac.uk)); Jehiel: Paris School of Economics and Department of Economics, University College London (email: [jehiel@enpc.fr](mailto:jehiel@enpc.fr)); Kitagawa: Department of Economics, University College London (email: [t.kitagawa@ucl.ac.uk](mailto:t.kitagawa@ucl.ac.uk)). Michael Ostrovsky was coeditor for this article. We thank participants in seminars at various institutions for useful comments and suggestions. We are grateful to the Economic Research Council (ERC) for funding this research through grants 210772 (Guarino), 742816 (Jehiel), and 715940 (Kitagawa). Kitagawa also thanks the ESRC Centre for Microdata Methods and Practice (CeMMAP) for grant RES-589-28-0001.

<sup>†</sup>Go to <https://doi.org/10.1257/mic.20180394> to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

information received by the predecessors perfectly, and learning is efficient (Lee 1993).

The intuition for this result is very simple. When agents choose in a continuous action space, their action reflects their information very precisely, under the assumption of full rationality. As a result, observing a predecessor's action is equivalent to observing the information he has. As more agents make their decisions in a sequence, more information is aggregated, and eventually the state of nature is learned and the best decision is reached.

A crucial ingredient in this story is that agents are able to make inferences correctly. If agents are less sophisticated, the process of learning may be inefficient despite a continuous action space. A series of studies have questioned the ability of human subjects to make correct inferences and revisited standard models of social learning (both with a discrete and a continuous action space) by analyzing what happens when this is not the case. Essentially, these works depart from the assumption of full rationality used in the previous literature and propose alternative ways of modeling how economic agents form expectations.

Although various biases may affect the process of learning, two aspects have attracted particular attention: information redundancy neglect (also sometimes referred to as correlation neglect) on the one hand, the tendency of human subjects to put more weight on their private information than on the public information contained in the choices of other participants on the other. These biases, proposed not only in economics but also in psychology and in other social and cognitive sciences, seem particularly salient for social learning. When making inferences from the choices of others, agents may not fully take into account that others have themselves been influenced by the same sources of information (redundancy neglect). On the other hand, they may consider themselves better at making a particular decision than the others they observe (overconfidence). Of course, these biases may coexist in human behavior.

The purpose of this work is to study, through a series of controlled experiments, social learning in a continuous action space in which agents move in sequence one after the other and observe their predecessors' choices as well as a private signal (the private information). Our interest is in understanding how the private signals observed by the predecessors and the signal observed by the agent influence the agent's decision. In the case of the predecessors' signals, we are interested in understanding whether the influence depends on whether the signal was observed by an early mover or a late mover in the sequence. Our objective is to shed light on whether the rational paradigm or any of the bounded rationality theories of information redundancy neglect or overconfidence capture human subjects' behaviors.

Before illustrating the experiment, let us discuss the different bounded rationality approaches in more detail. Information redundancy neglect in social learning has been studied, for instance, by Eyster and Rabin (2010). In their work, while each agent uses his private information and learns from others, he is convinced that others only use their private information: as a result, he interprets a predecessor's action as if it simply reflected the agent's private information. Since agents fail to account that their predecessors have already incorporated earlier signals in their

decisions, early signals have an excessive impact on later decisions.<sup>1</sup> In a similar spirit, Bohren (2016) studies a social learning environment in which agents have a misspecified model about their predecessors. While in the economy there is a fraction  $p$  of agents who (in addition to their own signal) observe the actions of others, agents believe that this fraction is actually  $\hat{p}$ . Bohren's (2016) model generalizes Eyster and Rabin's (2010) model, which is obtained when  $p = 1$  and  $\hat{p} = 0$ . As in Eyster and Rabin (2010), when  $\hat{p} < p$ , there is an overweighting of early signals since agents read actions as if they were reflecting more private signals than they actually do.

As we said, information redundancy neglect has emerged in studies well beyond the specific topic of social learning and even well beyond the boundaries of economics. A vast literature in statistics, sociology, computer science, physics, and economics has adopted the DeGroot (1974) model of learning. In that model, when agents repeatedly communicate, they update their beliefs by taking a weighted average of their neighbors' beliefs and their own belief from the previous period. Clearly, in this model, agents do not adjust correctly for repetitions and dependencies in information that they observe multiple times. Golub and Jackson (2010) apply the DeGroot updating rule to the study of learning in networks. An earlier study by DeMarzo, Vayanos, and Zwiebel (2003) presents a very similar idea by letting agents update as Bayesians but not taking into account repetitions. They label the failure to adjust properly for information repetitions as persuasion bias. Under persuasion bias, individuals do not account accurately for which components of the information they receive are new and which are repetitions.<sup>2</sup> Information redundancy neglect in our context, no matter how it is modeled, implies that signals observed by earlier movers have more influence than signals observed by later movers.

An alternative paradigm some scholars have suggested to understand social learning is that agents tend to put more weight on their private information than on the public information contained in the choices of others. Such a tendency is documented in various studies (e.g., Nöth and Weber 2003, Çelen and Kariv 2004, Goeree et al. 2007, and De Filippis et al. 2021), and it is typically referred to as "overconfidence" since subjects seem to trust their own information (or own ability to learn from it) more than their predecessors' information (or ability to learn from it). One interpretation that we will favor is that agents tend to mistrust the ability of their predecessors to understand their private signals correctly.

In a multistage, multiplayer game, one also needs to specify what subjects think about the overconfidence of others. If the overconfidence bias is common knowledge (this is the simplest formulation of overconfidence also adopted in a number of applied works), one can show that in a continuous action space, agents are still able to infer correctly the signals of others. As a result, all previous signals still have the same weight in the inference process but lower than the agent's private signal's

<sup>1</sup>Redundancy neglect was not the focus of a study by Guarino and Jehiel (2013); nevertheless, they also find an overweighting of early signals in a model of social learning with a continuous action space. We will discuss this work in the next section.

<sup>2</sup>We refer the reader to DeMarzo, Vayanos, and Zwiebel (2003) and Golub and Jackson (2010) for further references and discussions of the links to the psychology and sociology literatures.

weight. This is in sharp contrast with the early signals overweighting prediction of the models discussed above.

The overconfidence bias has been studied in many areas of economics other than social learning. In the theory of asset pricing, for instance, many works (e.g., Kyle and Wang 1997; Daniel, Hirshleifer, and Subrahmanyam 1998, 2001; Odean 1998) model traders' overconfidence as their overestimation of the precision of their private signal about security values. This can be interpreted as traders' overconfidence about the information they receive or overconfidence about their own ability to interpret the information (Odean 1998). Note that this approach to overconfidence is closely related to the definition of overconfidence about the precision of the own signal versus the signals of others in the social learning literature.<sup>3</sup>

We contribute to the understanding of the social learning process through some controlled experiments in which the different theoretical models just discussed can be tested. Specifically, we replicate a simple theoretical model of social learning with a continuous action space in the laboratory. In our experiment, subjects have 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 choose a number between 0 and 100, which represents the subjective probability (expressed as a percentage) that the value of the good is 100.<sup>4</sup> To elicit his belief, we use a quadratic scoring rule. We then ask a second subject to make the same type of prediction based on the observation of the first subject's decision only. Then, we provide the second subject with another, conditionally independent signal about the value of the good and ask him to make a new prediction. We then ask a third subject to make his prediction based on the observation of the first subject's decision and of the second subject's second decision. We then provide the subject with another signal and ask him to make a new prediction. The procedure continues until all subjects have made their predictions.

The results of our experiment are supportive of our overconfidence model and at odds with the predictions of the Perfect Bayesian Equilibrium and of the redundancy of information neglect model. When we consider the first action taken by subjects, we observe that subjects do not put higher weight on early signals than on late signals. Early decisions in the sequence do not have an undue influence on later decisions. On the contrary, predecessors' signals have a weight lower than one and constant as predicted by our overconfidence model. Moreover, when we consider the second action taken by subjects, we observe that subjects put a weight not statistically different from one on their own signal, again in line with our overconfidence model. Remember that, in line with previous work, we model overconfidence as an

<sup>3</sup>We refer to Odean (1998) for a discussion of the vast psychology literature on overconfidence about the own private information.

<sup>4</sup>As we will discuss, almost all existing social learning experiments use a discrete (indeed, binary) action space, motivated by the theoretical work of Bikhchandani, Hirshleifer, and Welch (1992). At a theoretical level, the continuous action space has been studied by Lee (1993), Vives (1993), and by many other authors, in particular in the study of financial markets. The trade size, the amount to be invested are obvious, real life examples of choices better approximated by a continuous action space. Asking subjects to state their beliefs in the laboratory is a simple way to implement a social learning game with a continuous action space.

equilibrium phenomenon. In our model, it is common knowledge that all agents are overconfident, that is, that they all put a lower (and constant) weight on the signals of others. Our results are not just in favor of a broad notion of overconfidence; they are in favor of this more specific and parsimonious model.

As we will discuss in the next section, past experiments on social learning had documented some form of overconfidence but were unable to disentangle whether overconfidence meant putting too much weight on the own signal or too little on the predecessors'. With our design, we find that subjects weigh their signal correctly but attribute less information content to previous actions than they actually have. In other words, our message is not that subjects are overly optimistic about the precision of their own signal but that they are overly pessimistic about the precision of their predecessors' signals (or, in a different interpretation, about the predecessors' ability to read or act upon these signals correctly).<sup>5</sup>

In terms of negative results, it is important to note that, while we explicitly test how a specific model of information redundancy neglect fits with the data, the results of our experiment are at odds with any model of such a bias since information redundancy neglect essentially means that early decisions have a disproportionate effect on subsequent decisions. In the laboratory, instead, all actions have the same effect, and this effect is even lower than what is predicted by the Perfect Bayesian Equilibrium.

It is worth mentioning that we also ran another treatment in which the same subject received a sequence of signals. This treatment mimics the social learning treatment except that now the same subject observes directly all past signals. In sharp contrast with the social learning treatment, in this treatment, the action that subjects take is consistent with the Bayesian one; that is, the weights on all signals are the same and not statistically different from that corresponding to the objective precision of signals. Thus, our finding in the social learning treatment can safely be attributed to overconfidence as interpreted above, as opposed to being the result of some form of non-Bayesian updating or some form of recency effect, according to which more recent news would have more weight than older news (such alternative theories would imply that we should see departures from the Bayesian prediction in the individual decision-making treatment, which we do not).

The rest of the paper is organized as follows. After a brief discussion of the related literature, Section I describes the theoretical model of social learning and the different theoretical predictions. Section II presents the experiment. Section III contains the results. Section IV offers a discussion and concludes. An online Appendix contains additional material.

*Related Literature.*—Most of the papers in the experimental social learning literature have used setups with a discrete action space. A partial exception is the interesting work by Çelen and Kariv (2004), who aim to distinguish informational cascades from herd behavior. In their experiment, the action space is still binary, but subjects are asked to choose a threshold in an interval before receiving the

<sup>5</sup>Perhaps one may find it more natural to think of this model as expressing some distrust in others' ability rather than overconfidence in the own ability. We stick to the label of "overconfidence" since it is quite in use in the existing literature.

signal (one of the two actions is then taken depending on the signal realization, a procedure similar to a Becker-DeGroot-Marshack mechanism (Becker, DeGroot, and Marschak 1964) that allows the experimentalists to elicit the median belief). Their econometric model shows that subjects give too little weight to the information revealed by the predecessors' choices relative to their own private information. Since early decision-makers tend to rely a lot on their own private information, their actions are informative about their signals, but this is not well taken into account by their successors. Our findings about overconfidence share similarities with those of Çelen and Kariv (2004), and our analysis reveals that overconfidence in our setting is the result of underweighting predecessors' information rather than inflating the precision of one's own information. Note that, in addition, our setting allows us to estimate finely the relative weights attached to the various predecessors' signals as a function of where they lie in the sequence, and such estimations are of essential importance to be able to test different models of social learning as explained above.

Among experiments with discrete action space, overweight of the private information relative to the information contained in the predecessors' actions was observed in an early experiment by Hung and Plott (2001), by use of a regression analysis. With a discrete action space, however, the weight put on the signal should change, also theoretically, depending on the specific sequence. This problem is not present in a setup with a continuous action space. Overconfidence is also documented in Nöth and Weber (2003). They conduct an experiment with binary actions and with binary signals of different precisions. They find that subjects tend to follow their signal even in circumstances in which it would be optimal to follow the herd (e.g., at time 2 upon receiving a signal of weak precision and different realization from the action taken at time 1). Goeree et al. (2007) present an experiment similar to the seminal work of Anderson and Holt (1997) but with long sequences of decisions. Their main conclusion is that "subjects tend to overweight their signals, or, alternatively, underweight the public prior generated by past publicly observed choices." Goeree et al. (2007) reach this conclusion by relying on the structural estimation of a Quantal Response Equilibrium (QRE) model. Thanks to the use of a continuous action space, we can estimate the influence of previous signals without relying on an equilibrium analysis and imposing weaker assumptions on the distribution of noise in the data. Moreover, thanks to our belief elicitation procedure, before and after receiving private information, we can identify the precise source of the overconfidence.<sup>6</sup>

It is also interesting to note that Goeree et al. (2007) consider a nonequilibrium model in the spirit of redundancy neglect in which expectations about error rates may be incorrect but find no strong evidence for such a model.<sup>7</sup> Different conclusions are reached by Kübler and Weizsäcker (2004). They conduct a binary action space experiment and study subjects' depth of reasoning through an error-rate model, which

<sup>6</sup>It should also be mentioned that in a discrete action space, subjects can observe deviations from equilibrium. With long sequences of decision-makers, cascades broke frequently in the experiment of Goeree et al. (2007). This may perhaps have affected subjects' beliefs about the errors made by others. In a continuous action space, there are no cascades, and it is impossible to know whether a predecessor acted "against his own signal." It is, therefore, even more remarkable that an overconfidence model still fits the data well.

<sup>7</sup>Similar conclusions are reached by Hung and Plott (2001).

differs from a Quantal Response Equilibrium model since no consistency between beliefs about error rates and actual error rates is imposed. They find that subjects underweigh the information contained in the predecessors' actions and also fail to realize that the predecessors' actions were themselves influenced by the history of actions they observed. In their words, subjects "behave as if disregarding the fact that their predecessors often use the information that is conveyed by third subjects' decisions" (Kübler and Weizsäcker 2004, 437). Note that this finding is an indication of information redundancy neglect. Essentially, subjects seem to use a simple "counting heuristic" in which they count the number of predecessors choosing action *A* or action *B* and attribute a lower information content than it actually has to each action (and independently of the specific sequence).

More recent papers related to our work are Eyster, Rabin, and Weizsäcker (2015); Cavatorta, Guarino, and Huck (2021); and Çelen, Geng, and Li (2019). In their work, Eyster, Rabin, and Weizsäcker (2015) present two treatments, one being similar to ours. In their experiment, though, subjects' task is simpler, in that they have to sum up the signal they receive and those received by the predecessors (not observed directly but inferable through the predecessors' choices). They find that a large fraction of choices is perfectly in line with the PBE predictions. Moreover, they do not find overconfidence, as we have defined it. Perhaps the difference in results is due to the difference in the difficulty of the task. When subjects face a simpler task, they are better able to behave rationally and trust that the predecessors did so too.<sup>8</sup> Cavatorta, Guarino, and Huck (2021) present some experiments with binary action space in which one of the two actions is never observed (the same way we do not observe people who consider the possibility of making an investment but do not make it); the other action is observable but with noise. In this more complex environment, they do not observe a clear bias by subjects to overweight the own private information or the public information contained in the observable action by predecessors. Çelen, Geng, and Li (2019) run an experiment with only two periods. While the first subject in the sequence has to predict the (binary) state of the world, the second subject, after receiving private information, can pay to receive extra private signal(s) (in one treatment) or to observe decision(s) made by first movers (in another treatment). If subjects believed the rationality of the first movers, the willingness to pay should be the same across treatments. It is, instead, lower for the social information, indicating that subjects do not fully trust the predecessor(s), a result in agreement with ours.

<sup>8</sup>Eyster, Rabin, and Weizsäcker (2015) also present a treatment in which more than one action is taken at each point in the sequence. This is based on work by Eyster and Rabin (2014), who illustrate the implications of rational herding in an extended setup in which more than one action is taken at the same time. The two experiments are independent; ours is antecedent, conducted in 2009, 2010, and 2011.

## I. The Theoretical Model

In our economy, there are  $T$  agents who make a decision in sequence. Time is discrete and indexed by  $t = 1, 2, \dots, T$ . Each agent, indexed by  $t$ , is chosen to take an action only at time  $t$  (in other words, agents are numbered according to their position in the sequence). The sequential order in which agents act is exogenously, randomly determined, with each sequence equally likely.

There is a good that can take two values,  $V \in \{0, 100\}$ . The two values are equally likely. Agent  $t$  takes an action  $a_t$  in the action space  $[0, 100]$ . The agent's payoff depends on his choice and on the value of the good. The payoff is quadratic and, in particular, equal to  $-(V - a_t)^2$ . Each agent  $t$  receives a private signal  $s_t \in \{0, 1\}$  correlated with the true value  $V$ . Specifically, he receives a symmetric binary signal distributed as follows:

$$\Pr(s_t = 1 | V = 100) = \Pr(s_t = 0 | V = 0) = q_t.$$

We assume that, conditional on the value of the good, the signals are independently distributed over time, with precision  $q_t \in (0.5, 1]$ . 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.

In addition to observing a private signal, each agent observes the sequence of actions taken by the predecessors. We denote the history of actions until time  $t - 1$  by  $h_t$ , that is,  $h_t = \{a_1, a_2, \dots, a_{t-1}\}$  (and  $h_1 = \emptyset$ ). Agent  $t$ 's information is then represented by the couple  $(h_t, s_t)$ . Given the information  $(h_t, s_t)$ , the agent chooses  $a_t$  to maximize his expected payoff  $E^S[-(V - a_t)^2 | h_t, s_t]$ ; therefore, his optimal action is  $a_t^* = E^S(V | h_t, s_t)$ .<sup>9</sup>

We now describe the different theoretical predictions.

Let us start with the Perfect Bayesian Equilibrium (PBE). Given that the action space is continuous, each action perfectly reveals the signal realization and its precision. Therefore, observing the sequence of actions is identical to observing the sequence of signals, and the process of learning is perfectly efficient. These observations lead to the following proposition.

**PROPOSITION 1** (Lee 1993): *In the PBE, after a sequence of signals  $\{s_1, s_2, \dots, s_t\}$ , agent  $t$  chooses action  $a_t^{PBE} = a_t^*(s_1, s_2, \dots, s_t)$  such that*

$$\frac{a_t^{PBE}}{100 - a_t^{PBE}} = \frac{a_t^*(s_1, s_2, \dots, s_t)}{100 - a_t^*(s_1, s_2, \dots, s_t)} = \prod_{i=1}^t \left( \frac{q_i}{1 - q_i} \right)^{2s_i - 1}.$$

*That is, the agent at time  $t$  acts as if he observed the sequence of all signals until time  $t$ .*

<sup>9</sup>The superscript  $S$  stands for subjective since we want to allow for subjective expectations in some of the theories discussed below.

Next is the “best response trailing naïve inference” (BRTNI) play proposed by Eyster and Rabin (2010). According to this theory, agents do not realize that predecessors’ actions already incorporate previous signals. Each agent learns from his own signals and from the actions of his predecessors but believes his predecessors choose their actions on the basis of their own signal only. Because of this, agent 3 in a sequence of decision-makers interprets agent 2’s decision as revealing his private information only. But in fact, agent 2’s action also reflects agent 1’s signal, which implies that agent 3 counts signal 1 twice, first through agent 1’s action and second through agent 2’s action.<sup>10</sup> By the same logic, as more agents make their decisions, early signals receive more and more weight. Indeed, the weight on predecessors’ signals increases exponentially with time. That leads to a severe overweighting of early signals.

**PROPOSITION 2** (Eyster and Rabin 2010): *If agents behave as in BRTNI, after a sequence of signals  $\{s_1, s_2, \dots, s_t\}$ , agent  $t$  chooses action  $a_t^{BRTNI} = a_t^{**}(s_1, s_2, \dots, s_t)$  such that*

$$\frac{a_t^{BRTNI}}{100 - a_t^{BRTNI}} = \frac{a_t^{**}(s_1, s_2, \dots, s_t)}{100 - a_t^{**}(s_1, s_2, \dots, s_t)} = \left[ \prod_{i=1}^{t-1} \left( \frac{q_i}{1 - q_i} \right)^{(2s_{i-1})(2^{t-i-1})} \right] \left( \frac{q_t}{1 - q_t} \right)^{2s_{t-1}}.$$

*That is, the agent at time  $t$  acts as if, in addition to his own signal, he had observed time- $i$  signal  $2^{t-i-1}$  times, for any  $i < t$ .*

Note that, while in the PBE each signal has an equal weight of 1 in the choice of the action, in BRTNI, the weights are exponentially decreasing.

Another study finding an overweighting of early signals in a model of social learning with a continuous action space is that by Guarino and Jehiel (2013). Their work is aimed at describing a steady state of an economy in which agents only understand the mapping between actions and the state of nature but not the map with the history of actions. In other words, agents do not take into account the history of decisions and their impact on subsequent agents’ decisions; they only consider the aggregate statistics conditional on a state of nature. Imposing a consistency condition on the aggregate statistics—in agreement with the Analogy Based Expectation Equilibrium (ABEE) of Jehiel (2005)—and requiring a genericity condition on the precisions, Guarino and Jehiel (2013) obtain the following equilibrium result.

**PROPOSITION 3** (Guarino and Jehiel 2013): *In the ABEE, after a sequence of signals  $\{s_1, s_2, \dots, s_t\}$ , agent  $t$  chooses action  $a_t^{ABEE} = a_t^{***}(s_1, s_2, \dots, s_t)$  such that*

$$\frac{a_t^{ABEE}}{100 - a_t^{ABEE}} = \frac{a_t^{***}(s_1, s_2, \dots, s_t)}{100 - a_t^{***}(s_1, s_2, \dots, s_t)} = \left[ \prod_{i=1}^{t-1} \left( \frac{q_i}{1 - q_i} \right)^{(2s_{i-1})(t-i)} \right] \left( \frac{q_t}{1 - q_t} \right)^{2s_{t-1}}.$$

<sup>10</sup>By this we mean that the agent at time 3 acts as if he had observed two time-1 signals.

*That is, the agent at time  $t$  acts as if, in addition to his own signal, he had observed time- $i$  signal  $t - i$  times, for any  $i < t$ .*

Note that in the ABEE, each agent  $t$ 's action in the sequence is taken as if the agent had observed a signal  $i$ ,  $i < t$ ,  $t - 1$  times, a much less severe overweighting than that in Eyster and Rabin (2010). It is also worth observing that while this work was not aimed at describing information redundancy neglect, nevertheless, the fact that agents do not take into account the impact of the sequence of decisions on successive actions leads to a form of redundancy neglect. In equilibrium, agents behave as if all actions had the same information content; this content is, however, determined in equilibrium by the aggregate frequencies since this is what agents focus on.

These predictions are immediately derived from theories already proposed in the literature. As we explained in the introduction, an alternative paradigm assumes that agents are “overconfident” in that they think others have a lower ability to understand their private signal. A simple interpretation is that agents think that others (and not themselves) may misread a good signal as a bad signal or vice versa. This is equivalent to thinking that others receive a signal of lower precision (which they correctly understand). Essentially, instead of having correct expectations on a predecessor  $i$ 's signal precision, agent  $t$  thinks that agent  $i$ 's signal precision is lower, so that the likelihood ratio after observing an action is  $(q_i/(1 - q_i))^{(2s_i-1)k}$  rather than  $(q_i/(1 - q_i))^{(2s_i-1)}$ , where  $k \in (0, 1)$ . We refer to this belief as “ $k$ -overconfidence.”<sup>11</sup> If  $k$ -overconfidence is common knowledge (i.e., each agent is  $k$ -overconfident and thinks others are  $k$ -overconfident, etc.), then it is easy to see that each predecessor's signal realization can be inferred from the predecessor's action. The agent, however, attributes to each such signal a lower precision. For instance, consider the case in which the first three signals are good. The first agent chooses  $a_1^{OC}$  such that  $a_1^{OC}/(100 - a_1^{OC}) = q_1/(1 - q_1)$ . Agent 2, however, “discounts” this action since he thinks agent 1 is less able to make the correct inference from the signal. Hence, after receiving his signal, agent 2 chooses  $a_2^{OC}$  such that  $a_2^{OC}/(100 - a_2^{OC}) = (q_1/(1 - q_1))^k(q_2/(1 - q_2))$ . Agent 3 agrees with agent 2 in reading agent 1's action (and so inferring his signal). Moreover, from observing action  $a_2^{OC}$ , he infers agent 2's signal since  $(q_2/(1 - q_2)) = (a_2^{OC}/(100 - a_2^{OC}))(q_1/(1 - q_1))^{-k}$ . He then “discounts” agent 2's action and chooses  $a_3^{OC}$  such that  $a_3^{OC}/(100 - a_3^{OC}) = (q_1/(1 - q_1))^k(q_2/(1 - q_2))^k(q_3/(1 - q_3))$ .<sup>12</sup> We refer to

<sup>11</sup>Note that, in agreement with much literature, we are defining overconfidence in relative terms (the agent is confident that he can interpret the signal better than his predecessors). An alternative definition would be that the agent is even overconfident in his own signal; that is, he thinks the signal is more informative than it is. We will come back to this point in Section IV.

<sup>12</sup>For the previous models, we have not explicitly discussed whether the precisions  $q_i$  are private information or common knowledge. For the PBE, this is irrelevant since each action reveals both the signal realization and its precision. For the BRTNI and ABEE, the precisions do not have to be known either. For the OC model, one interpretation is that the objective precisions are common knowledge; nevertheless, agents use subjective precisions since they believe the predecessors misread the signal realization with some probability. However, for the same logic as for the PBE, the precisions do not need to be common knowledge (since, given that  $k$  is common knowledge, agents can infer them from the actions).

this model of learning as the “overconfidence model” (OC). The next proposition, proven in online Appendix A, describes the equilibrium of the OC model.

**PROPOSITION 4:** *Suppose agents are  $k$ -overconfident and this is common knowledge. In equilibrium, after a sequence of signals  $\{s_1, s_2, \dots, s_t\}$ , agent  $t$  chooses action  $a_t^{OC} = a_t^{****}(s_1, s_2, \dots, s_t)$  such that*

$$\frac{a_t^{OC}}{100 - a_t^{OC}} = \frac{a_t^{****}(s_1, s_2, \dots, s_t)}{100 - a_t^{****}(s_1, s_2, \dots, s_t)} = \left[ \prod_{i=1}^{t-1} \left( \frac{q_i}{1 - q_i} \right)^{(2s_i-1)k} \right] \left( \frac{q_t}{1 - q_t} \right)^{2s_t-1},$$

where  $k \in (0, 1)$ . That is, the agent at time  $t$  acts as if he had observed the sequence of all signals until time  $t$ , but attributing precision  $(q_i/(1 - q_i))^k / (1 + (q_i/(1 - q_i))^k)$  to all predecessors’ signals.

To conclude, it is worth making two observations. First, the theories we have presented differ in the way agents learn from others and are identical in the way they learn from their own signal. For all theories, agents update their beliefs upon observing their private signal in a Bayesian fashion. Social learning is, instead, different since the theories postulate various ways of forming expectations on the value of the good from others’ actions, that is, various ways of inferring signals from actions.<sup>13</sup>

Second, while the propositions above express the relation between an agent’s action and his predecessors’ signals, the theories also offer a prediction in terms of relations between an agent’s action and his predecessors’ actions. Let us denote agent  $t$ ’s weight on the predecessor  $i$ ’s signal by  $\beta_{t,i}$  and the weight on the predecessor  $i$ ’s action by  $\gamma_{t,i}$ . In other words, let

$$\ln \frac{a_t}{100 - a_t} = \sum_{i=1}^{t-1} \beta_{t,i} (2s_i - 1) \ln \left( \frac{q_i}{1 - q_i} \right) + (2s_t - 1) \ln \left( \frac{q_t}{1 - q_t} \right),$$

$$\ln \frac{a_t}{100 - a_t} = \sum_{i=1}^{t-1} \gamma_{t,i} \ln \left( \frac{a_i}{100 - a_i} \right) + (2s_t - 1) \ln \left( \frac{q_t}{1 - q_t} \right).$$

We summarize the relations between  $\beta_{t,i}$ s and  $\gamma_{t,i}$ s in the next proposition that is proven in online Appendix A.

**PROPOSITION 5:** *Consider the matrix  $\Gamma$  containing the weights  $\gamma_{t,i}$  that agent  $t$  puts on the predecessors’ actions and the matrix  $\mathbf{B}$  containing the weights  $\beta_{t,i}$  that he puts on the predecessors’ signals:*

$$\mathbf{\Gamma} = \begin{pmatrix} 1 & 0 & \dots & \dots & 0 \\ -\gamma_{2,1} & 1 & 0 & \dots & 0 \\ -\gamma_{3,1} & -\gamma_{3,2} & 1 & & 0 \\ \vdots & \vdots & & \ddots & \vdots \\ -\gamma_{t,1} & -\gamma_{t,2} & \dots & -\gamma_{t,t-1} & 1 \end{pmatrix} \text{ and } \mathbf{B} = \begin{pmatrix} 1 & 0 & \dots & \dots & 0 \\ \beta_{2,1} & 1 & 0 & \dots & 0 \\ \beta_{3,1} & \beta_{3,2} & 1 & & 0 \\ \vdots & \vdots & & \ddots & \vdots \\ \beta_{t,1} & \beta_{t,2} & \dots & \beta_{t,t-1} & 1 \end{pmatrix}.$$

<sup>13</sup>To have a neat understanding of how people learn from others and of how they use their own signal, as we discussed in the introduction, in the experiment, we ask each subject to make two decisions: one after observing the predecessors’ actions only and another after observing the own signal too. The theoretical predictions concerning the first action are given in our propositions by the first part of the formulas, excluding the multiplier  $(q_i/(1 - q_i))^{2s_i-1}$ .

Matrix  $\Gamma$  is the inverse of matrix  $\mathbf{B}$ :  $\Gamma = \mathbf{B}^{-1}$ . Specifically, for the equilibrium solutions we considered, this implies that

(a) In the PBE, agent  $t$  chooses action  $a_t^{PBE}$  such that

$$\frac{a_t^{PBE}}{100 - a_t^{PBE}} = \frac{a_{t-1}^{PBE}}{100 - a_{t-1}^{PBE}} \left( \frac{q_t}{1 - q_t} \right)^{2s_t-1}.$$

(b) In the BRTNI, agent  $t$  chooses action  $a_t^{BRTNI}$  such that

$$\frac{a_t^{BRTNI}}{100 - a_t^{BRTNI}} = \prod_{i=1}^{t-1} \frac{a_{t-i}^{BRTNI}}{100 - a_{t-i}^{BRTNI}} \left( \frac{q_t}{1 - q_t} \right)^{2s_t-1}.$$

(c) In the ABEE, agent  $t$  chooses action  $a_t^{ABEE}$  such that

$$\frac{a_t^{ABEE}}{100 - a_t^{ABEE}} = \left[ \prod_{i=1}^{t-1} \left( \frac{a_{t-i}^{ABEE}}{100 - a_{t-i}^{ABEE}} \right)^{\text{sign}(\sin(\frac{i\pi}{3}))} \right] \left( \frac{q_t}{1 - q_t} \right)^{2s_t-1},$$

where the function  $\text{sign}(x) = \begin{cases} +1 & \text{if } x > 0 \\ 0 & \text{if } x = 0 \\ -1 & \text{if } x < 0. \end{cases}$

(d) In the OC model, agent  $t$  chooses action  $a_t^{OC}$  such that

$$\frac{a_t^{OC}}{100 - a_t^{OC}} = \left[ \prod_{i=1}^{t-1} \left( \frac{a_{t-i}^{OC}}{100 - a_{t-i}^{OC}} \right)^{k(1-k)^{i-1}} \right] \left( \frac{q_t}{1 - q_t} \right)^{2s_t-1}.$$

In the PBE, an agent’s action is just equal to his immediate predecessor’s action (belief) updated on the basis of the private signal. All previous actions have no weight since all the information (i.e., private signals) until time  $t - 1$  is contained in agent  $(t - 1)$ ’s action. Agent  $t$  would make exactly the same inference if instead of observing the entire sequence of actions, he only observed the immediate predecessor’s action. In the case of BRTNI, instead, agent  $t$ ’s action depends on *all* predecessors’ actions, with equal weight equal to 1. This is in fact the way BRTNI is constructed: by assumption, agents believe that the predecessors’ actions are chosen on the basis of their private information only. For the ABEE, there is no simple intuition in the action space. The reason is that in the ABEE, an action is chosen on the basis of the aggregate frequencies conditional on the value of the good and not on the basis of the sequence. The formula in the proposition shows that actions are weighted in a cyclical way as implied by the trigonometric function. Finally, for the OC model, agent  $t$ ’s action depends on *all* predecessors’ actions, but the weights are increasing (so that the early actions have little weight). Intuitively, note that when  $k$  approaches 0, the weight on all predecessors’ actions goes to 0; when  $k$  approaches 1, the weight on the predecessors’ action goes to 0, except for the immediate predecessor for which it approaches 1, as in the PBE; for

intermediate values, the weights are increasing: early actions keep counting but less and less since the information they contain is already partially incorporated in subsequent actions (otherwise, the agent would essentially be inferring the same private signal from more than one action).

## II. The Experiment and the Experimental Design

### A. The Experiment

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

The sessions started with written instructions given to all subjects (provided in online Appendix E). We explained to participants that they were all receiving the same instructions. Subjects could ask clarification questions, which we answered privately. The experiment was programmed and conducted with 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 ten participants. The procedures were the following:

- (i) Each session consisted of 15 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.
- (ii) In each round, we asked all subjects to make decisions in sequence, one after the other. For each round, the sequence was randomly chosen by the computer software. Each subject had an equal probability of being chosen in any position in the sequence.
- (iii) Participants were not told the value of the good. They knew, however, that they would receive information about the value, in the form of a symmetric binary signal. If the value was equal to 100, a participant would receive a “green ball” with probability 0.7 and a “red ball” with probability 0.3; if the value was equal to 0, the probabilities were inverted. That is, the green signal corresponded to  $s_t = 1$  and the red signal to  $s_t = 0$ ; the signal precision  $q_t$  was equal to 0.7 at any time.
- (iv) As we said, each round consisted of ten periods. In the first period, a subject was randomly chosen to make a decision. He received a signal and chose a number between 0 and 100, up to 2 decimal points.
- (v) The other subjects observed the decision made by the first subject on their screens. The identity of the subject was not revealed.

- (vi) In the second period, a second subject was randomly selected. He was asked to choose a number between 0 and 100, having observed the first subject's choice only.
- (vii) After he had made that choice, he received a signal and had to make a second decision. This time, therefore, the decision was based on the observation of the predecessor's action and of the private signal.
- (viii) In the third period, a third subject was randomly selected and asked to make two decisions, similarly to the second subject: a first decision after observing the choice of the first subject and the second choice of the second subject, a second decision after observing the private signal too. The same procedure was repeated for all the remaining periods until all subjects had acted. Hence, each subject, from the second to the tenth, made two decisions: one after observing the history of all (second) decisions made by the predecessors, the other after observing the private signal too.
- (ix) At the end of the round, after all ten 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 on 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.<sup>14</sup>

As should be clear from this description, compared to the existing experimental literature on social learning, informational cascades and herd behavior, we made two important procedural changes. First, in previous experiments, subjects were asked to make a decision in a discrete (typically binary) action space, whereas we ask subjects to choose actions in a very rich space, which practically replicates the continuum. This allows us to elicit their beliefs rather than just observing whether they prefer one action to another.<sup>15</sup> Second, in previous experiments, subjects made one decision after observing both the predecessors' actions and the signal. In our

<sup>14</sup>This explanation helps the subjects since they do not have to solve the maximization problem by themselves (and to which extent they are able to do so is not the aim of this paper). For a discussion of methodological issues related to elicitation methods, see the recent survey by Schotter and Trevino (2014).

<sup>15</sup>Within the discrete action space experiments, exceptions to the binary action space are the financial market experiments of Cipriani and Guarino (2005, 2009) and Drehmann, Oechssler, and Roeder (2005), 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 cutoff 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.

experiment, instead, they made two decisions, one based on public information only and one based on the private information as well.<sup>16</sup>

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 pound sterling at the exchange rate of 100 liras for £7. 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.

### B. 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 that was randomly drawn in the interval  $[0.7, 0.71]$  as opposed to having a constant precision of 0.7 as in SL1. Each subject observed not only the ball color but also the exact precision of his own signal. A third treatment (SL3) was identical to SL2, with the exception that instead of having sequences of ten subjects, we had sequences of four subjects. Given the smaller number of subjects, each round lasted less time; for this reason, we decided to run 30 rounds per session rather than 15. We have no evidence that the outcomes from these three treatments are any different. In particular, for each period, we ran a Wilcoxon rank-sum test on the session-specific medians, separately for the first and the second decision taken by subjects. Except in one case, which we attribute to chance, we never reject the null hypothesis that outcomes come from the same distribution (the results of these tests are reported in online Appendix B). Therefore, we consider the three treatments as just one experimental condition.<sup>17</sup> We will refer to it as the SL treatment.<sup>18</sup>

*Individual Decision-Making (IDM).*—In the social learning treatments, subjects make decisions after observing private signals and the actions of others. Clearly, we may expect departures from the PBE even independently of the social learning aspect if subjects do not update in a Bayesian fashion. To control for this, we ran a treatment in which subjects observed a sequence of signals and made more than

<sup>16</sup>Cipriani and Guarino (2009) use a quasi strategy method, asking subjects to make decisions conditional on either signal they might receive. Still, at each time, a subject never makes a decision based only on the predecessors' decisions.

<sup>17</sup>Drawing the precision from the tiny interval  $[0.7, 0.71]$ , instead of having the simpler setup with fixed precision equal to 0.7, was in line with models such as Eyster and Rabin (2010) and Guarino and Jehiel (2013), where the precision is indeed different for each agent. Reducing the length of the sequence to four subjects was instead motivated by the opportuneness to collect more data for the first periods of the sequence.

<sup>18</sup>Since the results of SL1 and SL2 were not statistically different, we did not run more treatments with signals of different precisions. Moreover, as we will see, in the experiment we observed a lot of heterogeneity in subjects' updating after observing the signal, and adding more heterogeneity in precisions would have just made the experiment computationally more demanding for the subjects (and with fewer possibilities of learning).

TABLE 1—TREATMENTS' FEATURES

	SL1	SL2	SL3	IDM
Signal precision	0.7	[0.7, 0.71]	[0.7, 0.71]	0.7
Sequence	10	10	4	1 or 2 or 3 or 4
Subjects in a group	10	10	4	—
Groups	5	5	5	—
Participants	50	49	20	30
Rounds	15	15	30	30

Notes: SL: Social Learning; IDM: Individual Decision-Making. In SL2, there are 49 subjects since one session was run with 9 participants rather than 10 due to a last-minute unavailability of 1 subject.

one decision.<sup>19</sup> Specifically, a subject received a signal (as subject 1 in the SL treatments) and had to make a choice in the interval  $[0, 100]$ . Then, with a 50 percent probability, he received another signal and had to make a second decision (similarly to the second decision of subject 2 in the SL treatments). Then, he could make 2 more decisions, and the probability of moving from one decision to the next was always 50 percent. Note that, at the cost of collecting less data, we decided not to ask subjects to make more than one decision in all rounds. Our purpose was to make the task of the subject as close as possible to that of a subject in the SL treatments. In other words, we wanted the subject to make his first decision not knowing whether he would be asked to make a second one, the second without knowing whether he could make a third, and so on. This way, his decisions were made in conditions as close as possible to the SL treatments. Table 1 shows the salient characteristics of all treatments.

### III. Results

Our main objective is to understand how human subjects learn from the observation of their predecessors' choices (social learning). To this purpose, we first focus on the first action taken by subjects at any time  $t > 1$  (denoted by  $a_t^1$ ). We will later study how subjects learn from their own signal together with social learning, by focusing on their second action at any time  $t$  (denoted by  $a_t^2$ ). Note that we use the notation  $a_t^2$  for mnemonic purposes;  $a_t^2$  coincides with  $a_t$ , as defined in Section I for the theoretical models (since in that section we only considered the theoretical action after observing the predecessors' actions and the own signal).

#### A. Inferring Others' Signals

Let us start by considering how the first action chosen by a subject at time  $t > 1$  ( $a_t^1$ ) is influenced by the signals received by the subject's predecessors. Of course, the subject does not observe these signals, but he does observe the actions the predecessors have chosen upon receiving these signals. The first four

<sup>19</sup>This treatment was conducted with the software z-Tree (Fischbacher 2007) in 2014. The payment followed the same rules. The exchange rate was modified so that, in expectation, subjects could receive a similar amount of money per hour spent in the laboratory.

propositions in Section I give very different predictions on how these signals are weighted. According to the PBE, each signal is correctly inferred and given an equal weight of 1. According to the redundancy of information neglect model, early signals have a much higher weight. According to the OC model, the weights are all equal but lower than 1. To test these different predictions, we run a regression analysis, focusing on the first action at each time. Using the first action has the advantage that we do not need extra assumptions on how subjects update upon receiving their private signal.<sup>20</sup> We use median regressions throughout the analysis.

Specifically, for each period  $t = 2, 3, \dots, 10$ , we regress the loglikelihood ratio of  $a_t^1$  on all the predecessors' signal likelihood ratios:

$$(1) \quad \ln\left(\frac{a_t^1}{100 - a_t^1}\right) = \beta_{t,1}\ln\left(\frac{q_1}{1 - q_1}\right)^{2s_1-1} + \beta_{t,2}\ln\left(\frac{q_2}{1 - q_2}\right)^{2s_2-1} + \dots + \beta_{t,t-1}\ln\left(\frac{q_{t-1}}{1 - q_{t-1}}\right)^{2s_{t-1}-1} + \varepsilon_t^1,$$

with  $\text{med}(\varepsilon_t^1 | s_1, s_2, \dots, s_{t-1}) = 0$ . Each coefficient  $\beta_{t,i}$  is the weight given by the median agent  $t$  to signal  $s_i$  for  $i < t$ . Subjects in the experiment sometimes choose the extreme values 0 and 100; for the dependent variable to be well defined in these cases, we rewrite  $a_t^1 = 100$  as  $a_t^1 = 100 - 0.1$  and  $a_t^1 = 0$  as  $a_t^1 = 0.1$ . Clearly, the choice of 0.1 is arbitrary. This choice, however, does not affect the median of the distribution as far as the proportion of the boundary actions is less than 1/2. In other words, we use a median regression rather than a least-square regression since the results of the latter are sensitive to extremely large or small values of the dependent variable and, hence, to how extreme values of the action are treated in the analysis. Given the experimental design, we have 300 observations for  $t = 2, 3, 4$ ; 150 observations for  $t = 5, \dots, 9$ ; and 135 for  $t = 10$ .<sup>21</sup> To account for unobserved correlations among subjects within each session, we compute bootstrap standard errors (using 500 replications) clustering at the session level (Hahn 1995).

Figure 1 shows the estimated coefficients  $\hat{\beta}_{t,i}$ , and their 95 percent confidence intervals. The estimated coefficients are systematically below 1 and do not exhibit any tendency to decrease from early to late periods. As shown in Table 2, the null hypothesis that the weights are equal to 1 (as in the PBE) is rejected at conventional significance levels. We find even stronger evidence against the hypothesis that the  $\beta_{t,i}$  coefficients take values according to the BRTNI and the ABEE predictions. On the other hand, the null hypothesis that weights are constant across periods is never rejected at the 5 percent significance level. These results reject the PBE as well as the redundancy of information neglect model predictions. Instead, they do not falsify the OC model.

<sup>20</sup>To give an example, suppose subjects, after observing others' actions, make inferences as in BRTNI. After receiving the own signal, though, they totally neglect the information coming from others' actions and only use the information contained in their own signal. If we had elicited the belief only once, after the subjects receive the private signal, we would conclude that subjects only use their signals and would be unable to observe this type of inference. We discuss this issue in greater detail in Section IIIC.

<sup>21</sup>Recall that in one session we had nine rather than ten subjects due to a last minute no-show by one subject.

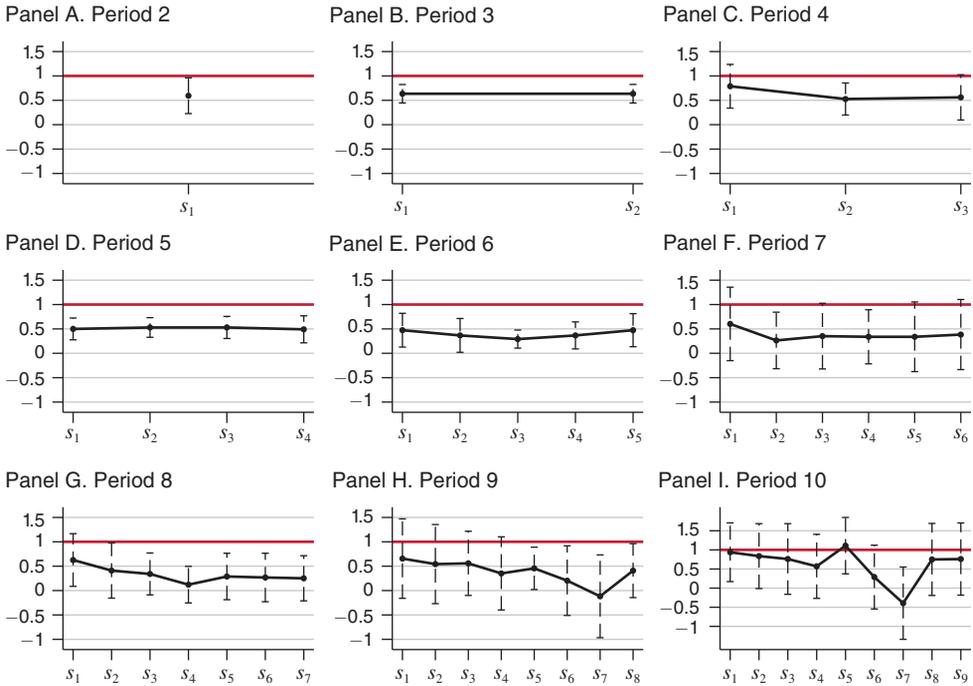


FIGURE 1. MEDIAN REGRESSIONS OF ACTION 1 ON PREDECESSORS' SIGNALS (ESTIMATED WEIGHTS)

*Notes:* The figure shows the estimated coefficients from a median regression of first action loglikelihood ratios on predecessors' signal loglikelihood ratios. For each period  $t = 1, \dots, 10$ , predecessors' signals,  $s_i, i = 1, \dots, t - 1$ , are on the x-axis; corresponding point estimates and 95 percent confidence intervals are on the y-axis, represented by black dots and dashed capped lines, respectively. Confidence intervals are computed by bootstrap (500 replications), clustering at the session level.

We estimate the “degree of overconfidence”  $k$ , under the hypothesis that at any time  $t$  the weight is constant for all signals  $1, 2, \dots, t - 1$  as predicted by the OC model. Table 3 reports the results. Estimates are significantly lower than 1 for all periods, approximately between 0.4 and 0.6. Note that our theoretical OC model imposes a further restriction, that is, that the parameter  $k$  is the same across periods. When we impose the further restriction that the parameter  $k$  is the same across periods, we obtain an estimate of 0.49 (last row of Table 3). Hence, subjects put on a predecessor's signal approximately half the weight that a Bayesian agent would put on a signal he would directly observe. Testing the null hypothesis that the parameters  $k$  are all equal across periods gives a  $p$ -value of 0.16; that is, we cannot reject the hypothesis that the parameters are equal at any conventional significance level. It is rather remarkable that the degree of overconfidence remains constant over time. One may suspect that the inference problem becomes more complicated for later decision-makers in the sequence and that, as a consequence, subjects attribute a lower information content to later predecessors' actions. This is not what the experimental data indicate. Perhaps subjects just form an expectation on how signals are reflected in each action and attribute it to all the actions they observe.

So far, we have estimated equation (1) and the parameter  $k$  using all predecessors' signals. One could observe that in some cases, subjects did not have a chance

TABLE 2—HYPOTHESIS TESTING: WEIGHTS ON PREDECESSORS' SIGNALS (*p*-VALUES)  
DEPENDENT VARIABLE: ACTION 1 (LOGLIKELIHOOD RATIO)

	$H_0^{PBE}$ : $\beta_{t,1} = \dots = \beta_{t,t-1} = 1$	$H_0^{BRTNI}$ : $\beta_{t,i} = 2^{t-i-1}, \forall i = 1, \dots, t-1$
Period 2	0.031	0.031
Period 3	0.000	0.000
Period 4	0.013	0.000
Period 5	0.000	0.000
Period 6	0.000	0.000
Period 7	0.000	0.000
Period 8	0.000	0.000
Period 9	0.000	0.000
Period 10	0.007	0.000

	$H_0^{ABEE}$ : $\beta_{t,i} = t-i, \forall i = 1, \dots, t-1$	$H_0^{OC}$ : $\beta_{t,1} = \dots = \beta_{t,t-1}$
Period 2	0.031	—
Period 3	0.000	0.864
Period 4	0.000	0.064
Period 5	0.000	0.987
Period 6	0.000	0.828
Period 7	0.000	0.872
Period 8	0.000	0.849
Period 9	0.000	0.903
Period 10	0.000	0.400

Note: The table reports tests based on bootstrap standard errors (500 replications), clustering at the session level.

TABLE 3—MEDIAN REGRESSIONS OF ACTION 1 ON PREDECESSORS' SIGNALS: ESTIMATION OF *k* UNDER  $H_0^{OC}$ :  $\beta_{t,1} = \dots = \beta_{t,t-1}$

	$\hat{k}$	95% confidence interval	
		Lower limit	Upper limit
Period 2	0.653	0.448	0.966
Period 3	0.635	0.508	0.847
Period 4	0.674	0.512	0.997
Period 5	0.504	0.287	0.668
Period 6	0.416	0.168	0.643
Period 7	0.404	0.183	0.651
Period 8	0.358	0.225	0.554
Period 9	0.381	0.251	0.679
Period 10	0.489	0.253	0.997
All	0.488	0.324	0.724

Note: The table reports 95% confidence intervals obtained with bootstrap (500 replications), clustering at the session level.

to infer the signal from the action. Consider, for instance, a subject in period 2 who observed  $a_1 = 50$ . Since the belief stated by subject 1 is identical to the prior, it was impossible to infer his signal.<sup>22</sup> To tackle this issue and check the robustness

<sup>22</sup>This is true even looking at the frequencies. Empirically, the choice of action 50 in period 1 was only slightly more frequent upon receiving a bad signal than a good one. Knowing these frequencies, the posterior belief upon observing action 50 at time 1 would have been 0.54.

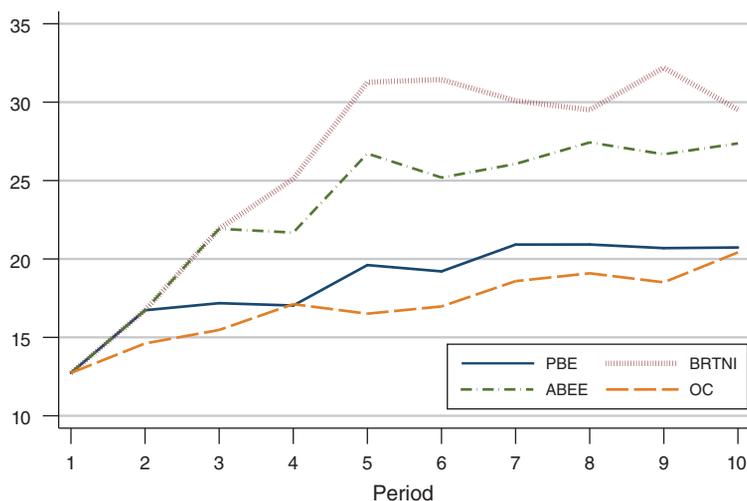


FIGURE 2. DISTANCE BETWEEN STATED BELIEFS AND MODELS' THEORETICAL VALUES

*Notes:* For each period, the figure shows mean distances between subjects' action 2 and theoretical predictions from PBE, BRTNI, ABEE, and OC models. Theoretical values for the OC model are computed using the value of the "degree of overconfidence"  $k = 0.488$  estimated in Table 3.

of our findings, we repeat our entire analysis after excluding the cases in which an action was, presumably, uninformative. For instance, we exclude the cases in which  $a_1 = 50$  and  $a_t^2 = a_{t-1}^2$  for  $t \geq 2$ . The results do not change compared to the present ones. We refer the reader to online Appendix C, where we discuss the precise methodology adopted and report the corresponding results.

Our regression results indicate that the OC model fits our data best. In Figure 2, we offer a simple and different illustration of this result. For each period  $t$ , we report the distances between the observed actions and each model's theoretical predictions, given the sequences of signals. In period 1, the different models make the same prediction, and the distance is, therefore, identical. For all following periods, the OC model has the smallest distance, except for periods 4 and 10, when the difference with the PBE model is negligible.<sup>23</sup>

### B. Inferring Others' Signals: A Heuristic Approach

The previous results suggest that the data from the laboratory are consistent with the OC model. One may wonder how subjects whose way of reasoning is similar to that of the OC model infer signals from actions. The question is relevant since in the previous analysis we used the true signal realizations, which are observed by us but not by the subjects themselves.

<sup>23</sup>Theoretical values for the OC model are computed using the value of the "degree of overconfidence"  $k = 0.488$  estimated in Table 3. In Section III E, we study how subjects' actions are related to predecessors' actions rather than signals. In this case, we estimate a "degree of overconfidence"  $k = 0.320$ . The OC model keeps exhibiting the smallest distance from stated beliefs across periods when its theoretical predictions are computed using the estimate  $k = 0.320$ .

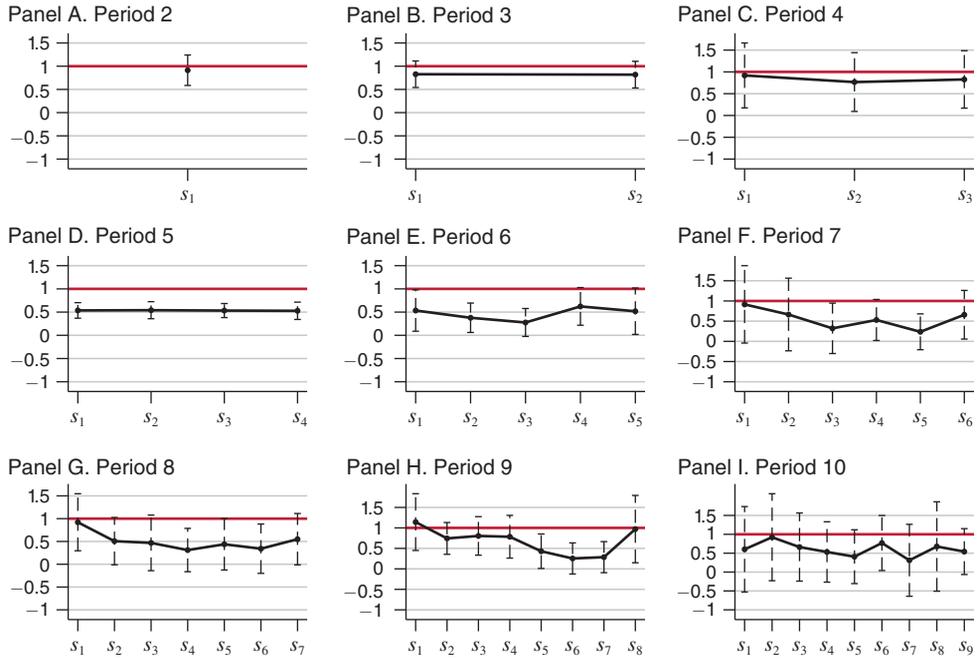


FIGURE 3. MEDIAN REGRESSIONS OF ACTION 1 ON PREDECESSORS' SIGNALS: A HEURISTIC APPROACH (ESTIMATED WEIGHTS)

Notes: The figure shows the estimated coefficients from a median regression of first action loglikelihood ratios on predecessors' signal loglikelihood ratios. For each period  $t = 1, \dots, 9$ , a predecessor's signal is  $s_i$  if  $|a_t^2 - a_i^{OC}(s_i)| < |a_t^2 - a_i^{OC}(1 - s_i)|$  or  $1 - s_i$  if  $|a_t^2 - a_i^{OC}(s_i)| > |a_t^2 - a_i^{OC}(1 - s_i)|$ , where  $a_i^{OC}(s_i)$  is the theoretical action  $a_i^{OC}$  conditional on the realization of the signal. For each period  $t = 1, \dots, 10$ , predecessors' signals,  $s_i$ ,  $i = 1, \dots, t - 1$ , are on the x-axis; corresponding point estimates and 95 percent confidence intervals are on the y-axis, represented by black dots and dashed capped lines, respectively. Confidence intervals are computed by bootstrap (500 replications), clustering at the session level.

While in the laboratory there is noise, it seems natural to think that if the OC model well represents subjects' way of thinking, subjects reconstructed the sequence of signals by inferring a signal  $s_i = 0$  ( $= 1$ ) whenever the observed action  $a_i^2$  was closer to the theoretical action  $a_i^{OC}$  conditional on a signal 0 (1) than to the theoretical action  $a_i^{OC}$  conditional on a signal 1 (0).<sup>24</sup> If that were the case, is it still true that the OC model is consistent with the data? To test this hypothesis, we now repeat the regression in equation (1) by substituting the true signal realizations with the signals constructed according to this heuristic, using the estimated value of the "degree of overconfidence"  $k = 0.488$ .<sup>25</sup> Figure 3 and Table 4 report the results.

As one can observe from Table 4, again the OC model is not rejected by the data. Figure 3, while not perfectly identical to Figure 1, still indicates that for most periods the weight on the predecessors' signals is systematically lower than 1. We will

<sup>24</sup> It is worth mentioning that subjects choose actions not coinciding with the PBE even at time 1, when they only have to make the inference from their own signal. Our heuristic approach takes this into account.

<sup>25</sup> The results reported in Figure 3 and Table 4 remain essentially unchanged when we use the estimate  $k = 0.320$  for the OC model (see footnote 23).

TABLE 4—HYPOTHESIS TESTING: WEIGHTS ON PREDECESSORS’ SIGNALS (*p*-VALUES)  
 DEPENDENT VARIABLE: ACTION 1 (LOGLIKELIHOOD RATIO) PREDECESSORS’ SIGNALS  
 INFERRED HEURISTICALLY

	$H_0^{OC}$ : $\beta_{t,1} = \dots = \beta_{t,t-1}$
Period 2	—
Period 3	0.648
Period 4	0.537
Period 5	0.999
Period 6	0.269
Period 7	0.734
Period 8	0.448
Period 9	0.068
Period 10	0.942

Note: The table reports tests based on bootstrap standard errors (500 replications), clustering at the session level.

come back to the relation between a subject’s action and the predecessors’ actions. For the time being, we notice that this heuristic lends more credibility to the fact that the OC model well summarizes subjects’ behavior in the laboratory since this model is not rejected by the data even when we reconstruct the signals on the basis of the model itself.

*C. Inferring Others’ Signals and Learning from the Own Signal*

While the previous results are compatible with the OC model, we still have to verify how subjects update their belief upon receiving their own signal. Recall that according to the OC model, agents, whereas attributing a lower weight to the predecessors’ signals, weigh their own signal correctly. To investigate this issue, we now study how subjects chose their second action,  $a_t^2$ . Specifically, for each period  $t = 2, 3, \dots, 10$ , we regress the loglikelihood ratio of  $a_t^2$  on all the predecessors’ signal likelihood ratios and on the own signal likelihood ratio:

$$(2) \quad \ln\left(\frac{a_t^2}{100 - a_t^2}\right) = \beta_{t,1} \ln\left(\frac{q_1}{1 - q_1}\right)^{2s_{t-1}} + \beta_{t,2} \ln\left(\frac{q_2}{1 - q_2}\right)^{2s_{t-2}} + \dots$$

$$+ \beta_{t,t-1} \ln\left(\frac{q_{t-1}}{1 - q_{t-1}}\right)^{2s_{t-1}-1} + \beta_{t,t} \ln\left(\frac{q_t}{1 - q_t}\right)^{2s_t-1} + \varepsilon_t^2.$$

As one can see from Figure 4, for each period  $t$ , the weight on the own signal is very close to 1 (and statistically not different from 1), with the only exception of period 10, for which it is actually higher than 1. The results of the hypothesis testing reported in Table 5 reveal that the predictions of the PBE, BRTNI, and ABEE are again strongly rejected by the data—with the trivial exception of period 1, in which there is no social learning, and that of period 7 for the PBE, which, in the absence of a clear pattern, may well be attributed to chance. We find, instead, support for the predictions of the OC model, which are not rejected by the data at any period, except for period 10. In view of the results presented in Table 3 for the weights on

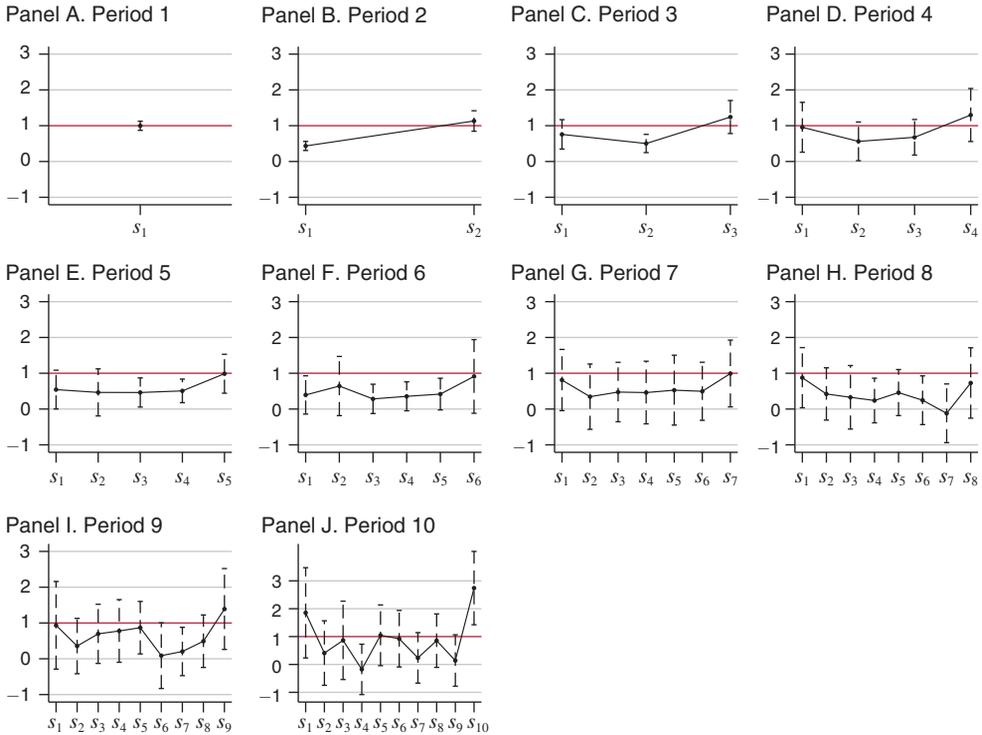


FIGURE 4. MEDIAN REGRESSIONS OF ACTION 2 ON OWN AND PREDECESSORS' SIGNALS (ESTIMATED WEIGHTS)

Notes: The figure shows the estimated coefficients from a median regression of second action loglikelihood ratios on own and predecessors' signal loglikelihood ratios. For each period  $t = 1, \dots, 10$ , predecessors' signals,  $s_i$ ,  $i = 1, \dots, t - 1$ , and own signal,  $s_t$ , are on the x-axis; corresponding point estimates and 95 percent confidence intervals are on the y-axis, represented by black dots and dashed capped lines, respectively. Confidence intervals are computed by bootstrap (500 replications), clustering at the session level.

predecessors' signals, the rejection of the OC model in period 10 is presumably due to the weight on the own signal being even higher than 1.

In the interest of space, we do not report the “degree of overconfidence”  $k$  estimated using action 2. We simply note that the results are in line with those in Table 3 obtained using action 1. In particular, when we impose the restriction that the parameter  $k$  is the same across periods, we obtain an estimated value of 0.57.

An interesting and important question is whether the weight given by subjects to predecessors' signals is the same before and after observing their own signal. To answer this question, we test whether the coefficients of the regression equations (1) and (2) are the same. That is, for each period, we estimate the median regressions (1) and (2) and then test whether pairwise differences between coefficients on predecessors' signals across the two regressions are statistically different from zero. Again we adopt a bootstrap procedure with standard errors clustered at the session level. We find strong evidence in favor of the null, which is never rejected for any pair of coefficients across periods.

TABLE 5—HYPOTHESIS TESTING: WEIGHTS ON OWN AND PREDECESSORS’ SIGNALS (*p*-VALUES)  
DEPENDENT VARIABLE: ACTION 2 (LOGLIKELIHOOD RATIO)

	$H_0^{PBE}$ : $\beta_{t,1} = \dots = \beta_{t,t} = 1$	$H_0^{BRTNI}$ : $\beta_{t,i} = 2^{t-i-1}, \forall i = 1, \dots, t-1,$ $\beta_{t,t} = 1$
Period 1	0.976	0.976
Period 2	0.000	0.000
Period 3	0.000	0.000
Period 4	0.019	0.000
Period 5	0.020	0.000
Period 6	0.000	0.000
Period 7	0.427	0.000
Period 8	0.000	0.000
Period 9	0.000	0.000
Period 10	0.000	0.000

	$H_0^{ABEE}$ : $\beta_{t,i} = t - i, \forall i = 1, \dots, t - 1,$ $\beta_{t,t} = 1$	$H_0^{OC}$ : $\beta_{t,1} = \dots = \beta_{t,t-1},$ $\beta_{t,t} = 1$
Period 1	0.976	0.976
Period 2	0.000	0.366 <sup>a</sup>
Period 3	0.000	0.467
Period 4	0.000	0.469
Period 5	0.000	0.994
Period 6	0.000	0.923
Period 7	0.000	0.947
Period 8	0.000	0.165
Period 9	0.000	0.878
Period 10	0.000	0.005

Note: The table reports tests based on bootstrap standard errors (500 replications), clustering at the session level.  
<sup>a</sup>The reported *p*-value refers to the null hypothesis that  $\beta_{2,2} = 1$ , while  $\beta_{2,1}$  can take any value. We also compute the value of  $\beta_{2,1}$  that minimizes the quantile regression criterion function, under the constraint that  $\beta_{2,2} = 1$ . We obtain a value of  $\beta_{2,1} = 0.476$  with a bootstrap 95 percent confidence interval of [0.308, 0.548].

Under the assumption (backed by the data) that the median regression coefficients in equations (1) and (2) are the same, subtracting (2) from (1) gives

$$\ln\left(\frac{a_t^2}{100 - a_t^2}\right) - \ln\left(\frac{a_t^1}{100 - a_t^1}\right) = \beta_{t,t} \ln\left(\frac{q_t}{1 - q_t}\right)^{2s_t-1} + \varepsilon_t^2 - \varepsilon_t^1.$$

If  $\varepsilon_t^2 - \varepsilon_t^1$  is median zero conditional on  $(s_1, \dots, s_t)$ —a sufficient condition for it is that  $\varepsilon_t^1$  and  $\varepsilon_t^2$  are i.i.d., symmetrically and unimodally distributed around zero—then the median regression of the update from the first to the second belief depends only on the current signal and not on the past signals. We have tested this implication by regressing the difference between the loglikelihood ratio of  $a_t^2$  and the loglikelihood ratio of  $a_t^1$  on predecessors’ and own signal likelihood ratios. We find that the weight on the own signal is very close to 1 and not statistically different from 1 across periods (except for period 4, where the estimated coefficient is 0.76 with a bootstrap 95 percent confidence interval ranging from 0.51 to 0.94). Importantly, the weights on predecessors’ signals are small in magnitudes and rarely statistically different from zero. Overall, predecessors’ signals have no predictive power for the update from action 1 to action 2, as the joint null hypothesis that all coefficients on predecessors’ signals are zero is not rejected. The only exceptions are period 5 and

period 9. In both cases, however, while we reject the hypothesis that predecessors' signals have no explanatory power, the estimated coefficients appear relatively small in value.

In summary, our analyses show that subjects put the correct weight of 1 on their own signal and approximately half the weight on their predecessors' signals.<sup>26,27</sup>

#### D. A Control: Inferring from a Sequence of Signals

In our experiment, the private signal also happens to be the latest piece of information subjects receive before making their decision. One may wonder whether a tendency of human subjects to put more weight on the most recent piece of information could explain (or, at least, affect) our results. It could be that, independently of social learning, subjects put more weight on the latest signal, compared to previous information.

To check for this possibility and, more generally, for deviation from equilibrium due to behavioral departures from Bayesian updating, we ran a treatment in which subjects observe a sequence of signals directly (IDM treatment). Figure 5 shows the results. Subjects put the same weight on all signals, a result that holds at any period. While the weight is not significantly different from 1 for the first three periods, it is significantly higher in period 4 (the  $p$ -value for the hypothesis that all weights are equal to 1 in period 4 is 0.00). The result is, however, affected by subjects choosing extreme actions (0 or 100) after four signals of the same color. If we exclude these cases, the weights are again not significantly different from 1 ( $p$ -value of 0.90).

Overall, the results of this treatment confirm that our findings in the SL treatment are indeed related to how individuals learn from the actions of others and not to subjective beliefs on signal precisions or to behavioral departures from Bayesian updating (which would emerge even when subjects directly receive information rather than having to infer it from others' actions).

#### E. Social Learning

In Section IIIA, we have offered a direct test of our first four theoretical propositions. In Section IIIB, we have proposed a heuristic for inferring signals from actions consistent with the OC model and shown that it lends further support to the results previously obtained. Now, we take one more step and study how subject  $t$ 's action is related not to his predecessors' signals (true or as inferred through a heuristic) but to his predecessors' actions. There are various reasons to do this. First, the

<sup>26</sup>One may think that, by eliciting beliefs twice, we make subjects pay more attention to the information content of the predecessors' actions. Also, psychological biases, such as the confirmation bias, may induce subjects to discount the own signal, which is received only after forming a belief by observing others' actions. If this were the case, our experimental design would overestimate the weight put on the predecessors' actions, that is, bias the results against the OC model.

<sup>27</sup>A further implication of the OC model is that the difference between the first action at time  $t > 1$  ( $a_t^1$ ) and the second action at time  $t - 1$  is constant over time (and negative, considering the case  $s_{t-1} = 1$ ). In contrast, the PBE model predicts a difference of zero at any time  $t$ , and the other models predict it to be increasing over time. In online Appendix D, we show that the median difference (of actions' loglikelihood ratios) is negative and roughly constant, again in agreement with the OC model.

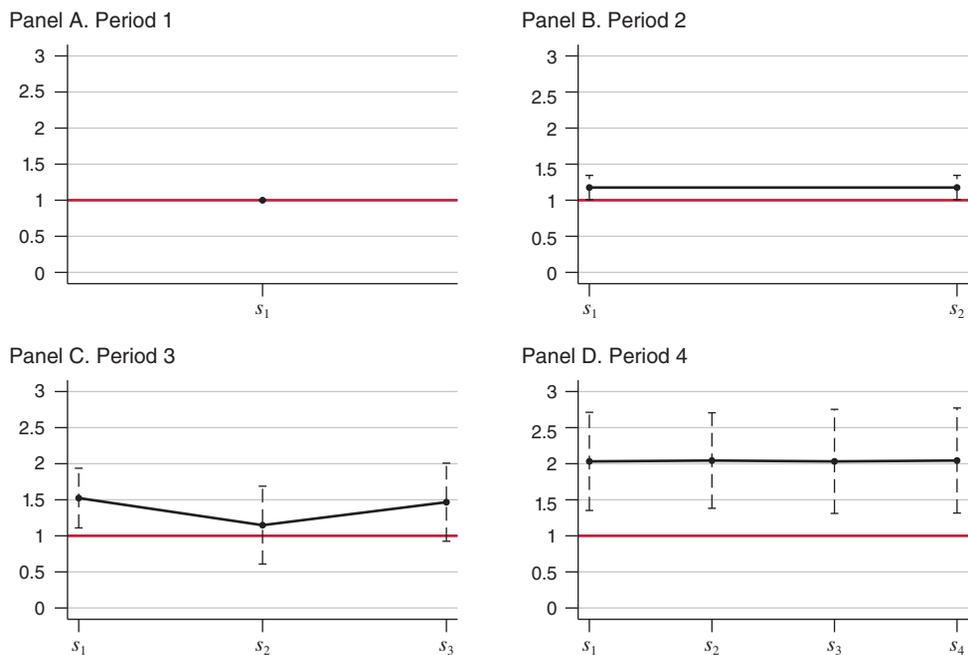


FIGURE 5. MEDIAN REGRESSIONS OF ACTION ON OWN SIGNALS IN IDM TREATMENT (ESTIMATED WEIGHTS)

*Notes:* The figure shows the estimated coefficients from a median regression of subjects' action loglikelihood ratios on own signal loglikelihood ratios. For each period  $t = 1, 2, 3, 4$ , own signals,  $s_t$ , are on the x-axis; corresponding point estimates and 95 percent confidence intervals are on the y-axis, represented by black dots and dashed capped lines, respectively. Confidence intervals are computed by bootstrap (500 replications), clustering at the session level.

predecessors' actions are what a subject actually observes. Second, this analysis can reveal a behavior not detectable by focusing on signals. As an example, suppose subject 4 forms expectations as in the PBE and chooses the action as in the PBE. Suppose he observes a sequence of actions 70, 83, 70. He then chooses 70, which is consistent with the PBE. If, however, subject 3 had received the good signal and not the bad signal (and not used his signal as in the PBE), by studying the relation between actions and predecessors' signals we would not classify subject 4's action as PBE.<sup>28</sup> For PBE and BRTNI, this analysis offers a more immediate test of the theories since the PBE predicts that only the immediate predecessor's action matters ( $a_t^2 = a_{t-1}^2$ ), and BRTNI predicts that subjects take actions as signals. Essentially, focusing on the signals is equivalent to a joint test: a test that subjects form expectations from predecessors' actions as the theoretical models predict and that predecessors use their own signals as the models predict. By studying the relation between an action and the predecessors' actions, instead, we do not need to rely on assumptions

<sup>28</sup>Of course, there is a merit in studying the relation with signals, as we did before. Consider the case in which the first three signals are all good and the sequence of actions is 70, 99, 93. Subject 4 could infer that all three signals were good (the third subject corrected the overreaction by subject 2) and act as in a PBE. The PBE prediction would not be falsified by considering the relation between an action and the predecessors' signals, whereas it would be by considering that between an action and the predecessors' actions.

or heuristics on how subjects relate predecessors’ actions to their signals. This analysis offers a test of the predictions described in Proposition 5.

Recall that in the PBE, only the immediate predecessor’s action is relevant for the choice of action  $a_t^1$  (all other actions have zero weight). According to BRTNI, instead, all previous actions have an equal weight of 1. The OC model predicts a specific relation between the action taken at time  $t$  and the predecessors’ actions, with early actions having less weight than late ones.

To test these predictions, we estimate the following equations by median regressions:

$$(3) \ln\left(\frac{a_t^1}{100 - a_t^1}\right) = \gamma_{t,1}\ln\left(\frac{a_1^2}{100 - a_1^2}\right) + \gamma_{t,2}\ln\left(\frac{a_2^2}{100 - a_2^2}\right) + \dots + \gamma_{t,t-1}\ln\left(\frac{a_{t-1}^2}{100 - a_{t-1}^2}\right) + \varepsilon_t^1;$$

$$(4) \ln\left(\frac{a_t^1}{100 - a_t^1}\right) = \gamma_{t,1}\ln\left(\frac{a_1^2}{100 - a_1^2}\right) + \gamma_{t,2}\ln\left(\frac{a_2^2}{100 - a_2^2}\right) + \dots + \gamma_{t,t-1}\ln\left(\frac{a_{t-1}^2}{100 - a_{t-1}^2}\right) + \gamma_{t,t}\ln\left(\frac{q_t}{1 - q_t}\right)^{2s_t-1} + \varepsilon_t^1.$$

Note that for the right-hand-side variables, we approximate  $a_t^2 = 100$  as  $a_t^2 = 100 - 0.1$  and  $a_t^2 = 0$  as  $a_t^2 = 0.1$  (the same approximation remains true for the dependent variable, as explained in the previous subsection).

The results of the estimation of equation (3) are shown in Figure 6 and Tables 6–7. They confirm our previous findings: while the data are at odds with the PBE, the BRTNI, and the ABEE, they support the predictions of the OC model. Specifically, the estimated coefficients on the predecessors’ actions exhibit a somewhat increasing pattern, as suggested by the OC model, which is more marked in periods 2–5. The estimates of the  $k$  parameter are again lower than 1 and actually typically lower than the estimates from the regressions of actions on predecessors’ signals, although for most periods, and overall (when we assume that  $k$  is constant across periods), the estimates obtained from the two different regressions are not statistically different from each other.<sup>29</sup>

Finally, in Figure 7 and Table 8, we report the results of the estimation of equation (4). Once again, the data are in contrast with the PBE, the BRTNI, and the ABEE but not with the OC model.<sup>30</sup> This model is never rejected, except for

<sup>29</sup>In the case of constant  $k$ , the  $p$ -value is 0.09.

<sup>30</sup>It is worth mentioning that whereas in the social learning literature, as in much psychological literature, researchers have talked about “overconfidence,” in other experimental studies subjects show “underconfidence.” In particular, in experiments on decision-making with naïve 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 in favor of overconfidence and at odds with underconfidence.

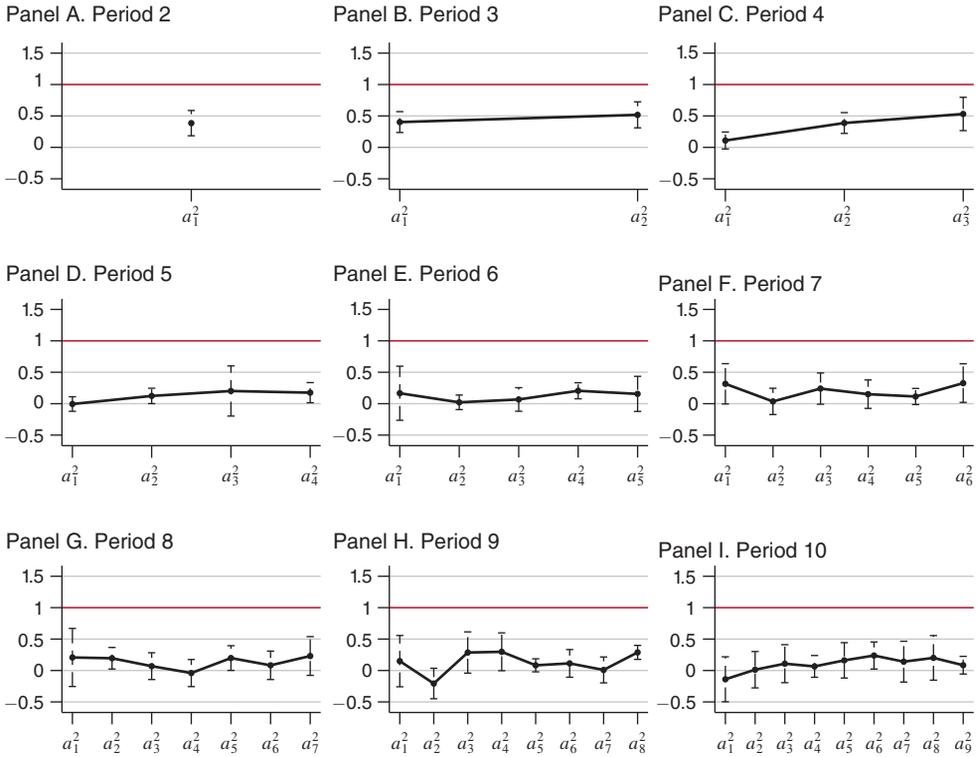


FIGURE 6. MEDIAN REGRESSIONS OF ACTION 1 ON PREDECESSORS' ACTION 2 (ESTIMATED WEIGHTS)

Notes: The figure shows the estimated coefficients from a median regression of first action loglikelihood ratios on predecessors' second action loglikelihood ratios. For each period  $t = 1, \dots, 10$ , predecessors' actions,  $a_i^2$ ,  $i = 1, \dots, t - 1$ , are on the x-axis; corresponding point estimates and 95 percent confidence intervals are on the y-axis, represented by black dots and dashed capped lines, respectively. Confidence intervals are computed by bootstrap (500 replications), clustering at the session level.

period 10, presumably because the estimated coefficient on the own signal takes a value significantly larger than 1.

### F. Efficiency and Convergence

Finally, we want to study the consequences of the observed behavior in the laboratory in terms of learning (in)efficiency. The PBE offers a benchmark for efficient learning: in the PBE, each agent perfectly infers the signals from the predecessors' actions and uses the information to choose the optimal action. The private information is aggregated, and, eventually, agents learn the true value of the good almost surely. In contrast, in the other theoretical models we have presented, there are inefficiencies due to the incorrect beliefs agents form. Asymptotic convergence to the realized value occurs in the other theoretical models, with the exception of BRTNI, where, given the extreme overweight of early actions, beliefs can converge to the incorrect value of the good (Eyster and Rabin 2010). While in our experiment, with

TABLE 6—HYPOTHESIS TESTING: WEIGHTS ON PREDECESSORS’ ACTION 2 (*p*-VALUES)  
DEPENDENT VARIABLE: ACTION 1 (LOGLIKELIHOOD RATIO)

	$H_0^{PBE}$ : $\gamma_{t,1} = \dots = \gamma_{t,t-2} = 0,$ $\gamma_{t,t-1} = 1$	$H_0^{BRTNI}$ : $\gamma_{t,1} = \dots = \gamma_{t,t-1} = 1$
Period 2	0.000	0.000
Period 3	0.000	0.000
Period 4	0.000	0.000
Period 5	0.000	0.000
Period 6	0.000	0.000
Period 7	0.000	0.000
Period 8	0.000	0.000
Period 9	0.000	0.000
Period 10	0.000	0.000

	$H_0^{ABEE}$ : $\gamma_{t,i} = \text{sign}\left(\sin\left(\frac{t-i}{3}\pi\right)\right)$ $\forall i = 1, \dots, t-1$	$H_0^{OC}$ : $\gamma_{t,i} = \gamma_{t,t-1}(1 - \gamma_{t,t-1})^{t-i-1}$ $\forall i = 1, \dots, t-1$
Period 2	0.000	—
Period 3	0.000	0.078
Period 4	0.000	0.209
Period 5	0.000	0.299
Period 6	0.000	0.504
Period 7	0.000	0.187
Period 8	0.000	0.854
Period 9	0.000	0.087
Period 10	0.000	0.760

Note: The table reports tests based on bootstrap standard errors (500 replications), clustering at the session level.

TABLE 7—MEDIAN REGRESSIONS OF ACTION 1 ON PREDECESSORS’ ACTION 2: ESTIMATION OF *k* UNDER  $H_0^{OC}$ :  $\gamma_{t,i} = \gamma_{t,t-1}(1 - \gamma_{t,t-1})^{t-i-1}, \forall i = 1, \dots, t-1$

	$\hat{k}$	95% confidence interval	
		Lower limit	Upper limit
Period 3	0.469	0.279	0.598
Period 4	0.580	0.309	0.861
Period 5	0.154	0.133	0.415
Period 6	0.154	0.081	0.844
Period 7	0.248	0.105	0.754
Period 8	0.180	0.092	0.752
Period 9	0.324	0.101	0.429
Period 10	0.211	0.133	0.336
All	0.320	0.155	0.564

Note: The table reports 95 percent confidence intervals obtained with bootstrap (500 replications), clustering at the session level.

sequences of four or ten signals, we cannot study asymptotic convergence to the realized value, still we can compare the stated beliefs with the PBE ones. We proceed in two ways. First, to study efficiency, we compare, period by period, the realized average payoffs with the expected theoretical payoffs under PBE. Second, we look at the distance between stated and PBE beliefs at the end of the sequence of decision-making (i.e., period 4 in treatment SL3 and period 10 in treatments SL1 and SL2).

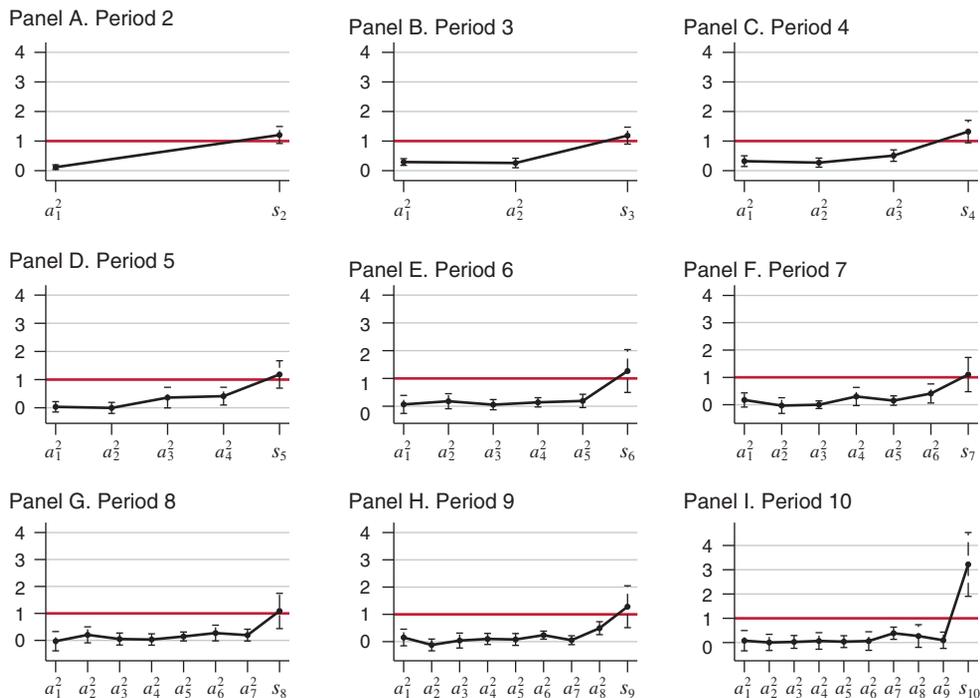


FIGURE 7. MEDIAN REGRESSIONS OF ACTION 2 ON PREDECESSORS' ACTION 2 AND OWN SIGNAL (ESTIMATED WEIGHTS)

Notes: The figure shows the estimated coefficients from a median regression of second action loglikelihood ratios on predecessors' second action loglikelihood ratios and own signal loglikelihood ratios. For each period  $t = 1, \dots, 10$ , predecessors' actions,  $a_i^2, i = 1, \dots, t - 1$ , and own signal,  $s_t$ , are on the x-axis; corresponding point estimates and 95 percent confidence intervals are on the y-axis, represented by black dots and dashed capped lines, respectively. Confidence intervals are computed by bootstrap (500 replications), clustering at the session level.

Figure 8 reports the average realized per-period payoff as a ratio of the average payoff subjects would have obtained had they played as in the PBE. In period 1, subjects earn 96 percent of what is potentially obtainable. Over periods, this percentage declines but only slightly, to over around 90 percent in the final periods of the experiment.<sup>31</sup> The figure also reports the ratio between the average OC, ABEE, and BRTNI payoffs and the average PBE payoffs. Note that the OC model predicts a slightly higher efficiency than the realized one. This is true even at time 1, however, when according to the model learning is fully efficient (since the inference from the signal is always correct). Although, as we have seen, the median action at time 1 is in line with the PBE (and so is efficient), there is heterogeneity in the actions, and this determines the loss in efficiency (as highlighted by the realized payoffs). This loss is approximately constant over periods, as one can notice by comparing the blue solid line and the orange dashed line. In other words, once we take into account this loss of efficiency due to the heterogeneous use of the own signal, the OC model

<sup>31</sup> It is important to observe that these deviations from the PBE may look smaller than they are for the simple reason that they are *average* deviations.

TABLE 8—HYPOTHESIS TESTING: WEIGHTS ON PREDECESSORS’ ACTION 2 AND OWN SIGNAL (*p*-VALUES)  
DEPENDENT VARIABLE: ACTION 2 (LOGLIKELIHOOD RATIO)

	$H_0^{PBE} :$ $\gamma_{t,1} = \dots = \gamma_{t,t-2} = 0$ $\gamma_{t,t-1} = 1$ $\beta_{t,t} = 1$	$H_0^{BRTNI} :$ $\gamma_{t,1} = \dots = \gamma_{t,t-1} = 1$ $\beta_{t,t} = 1$
Period 2	0.000	0.000
Period 3	0.000	0.000
Period 4	0.000	0.000
Period 5	0.008	0.000
Period 6	0.000	0.000
Period 7	0.024	0.000
Period 8	0.000	0.000
Period 9	0.000	0.000
Period 10	0.000	0.000
	$H_0^{ABEE} :$ $\gamma_{t,i} = \text{sign}\left(\sin\left(\frac{t-i}{3}\pi\right)\right)$ $\forall i = 1, \dots, t-1$ $\beta_{t,t} = 1$	$H_0^{OC} :$ $\gamma_{t,i} = \gamma_{t,t-1}(1 - \gamma_{t,t-1})^{t-i-1}$ $\forall i = 1, \dots, t-1$ $\beta_{t,t} = 1$
Period 2	0.000	—
Period 3	0.000	0.169
Period 4	0.000	0.093
Period 5	0.000	0.383
Period 6	0.000	0.942
Period 7	0.000	0.405
Period 8	0.000	0.783
Period 9	0.000	0.273
Period 10	0.000	0.006

Note: The table reports tests based on bootstrap standard errors (500 replications), clustering at the session level.

predicts the inefficiency in the data remarkably well. The BRTNI model, instead, predicts a marked reduction in efficiency (red dotted line) that we do not observe in the data.

Figure 9 reports the histogram of the distance between stated beliefs (as measured by  $a_t^2$ ) and PBE beliefs in the last period of decision-making (i.e., period 4 in SL3 and period 10 in SL1 and SL2). The distance is computed as the absolute value of the difference between the two beliefs. A distance lower than 5 occurred in 33 percent of the cases in period 4 (SL3) and in 44 percent of the cases in period 10 (SL1 and SL2), indicating also a process of convergence over time. Overall, while the underweight of the predecessors’ signals poses a limit to a perfect convergence, the cases in which the inference is strongly incorrect are few, as shown by the low occurrence of instances in which the distance is higher than 50.

#### IV. Discussion and Conclusions

To conclude, it is worth discussing some aspects of our experiment and our results.

First, our OC model assumes the overconfidence parameter  $k$  to be common knowledge among players. Such assumption, while strong, is shared by a number of other papers, for example about bargaining (see, e.g., Yildiz 2003, 2011). Other assumptions would sound plausible as well, for example, allowing the  $k$  parameter to be heterogeneous among agents or allowing agents to entertain subjective

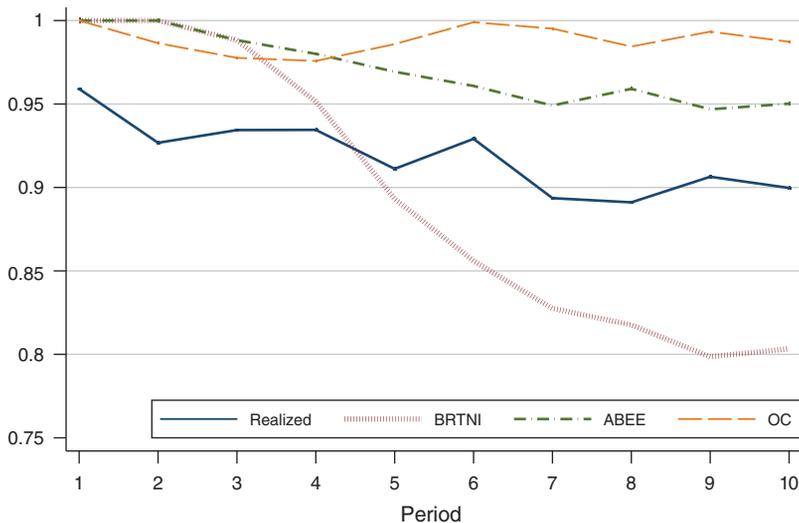
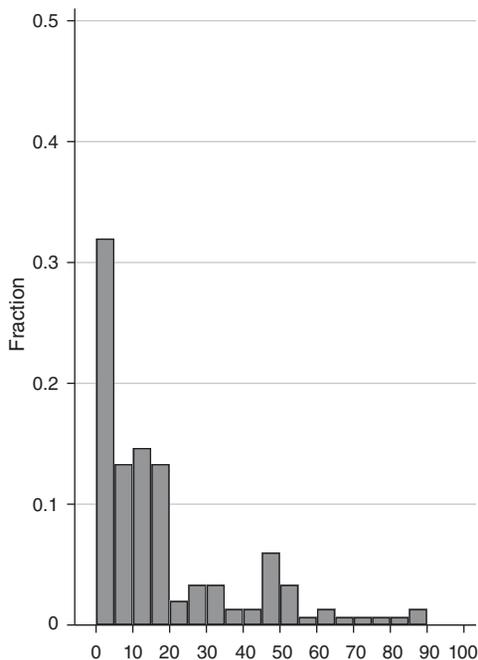


FIGURE 8. RATIO OF AVERAGE PAYOFF TO AVERAGE PBE PAYOFF

Notes: For each period  $t$  and model  $M = \{PBE, BRTNI, ABEE, OC\}$ , the payoff is computed as  $100 - 0.01 (V - a_{ii}^M)^2$  for each observation  $i$  and averaging across all observations. To compute payoffs under the OC model, we assume a value of the overconfidence  $k$  parameter equal to 0.488 (last row in Table 3).

Panel A. SL3: Period 4



Panel B. SL1&2: Period 10

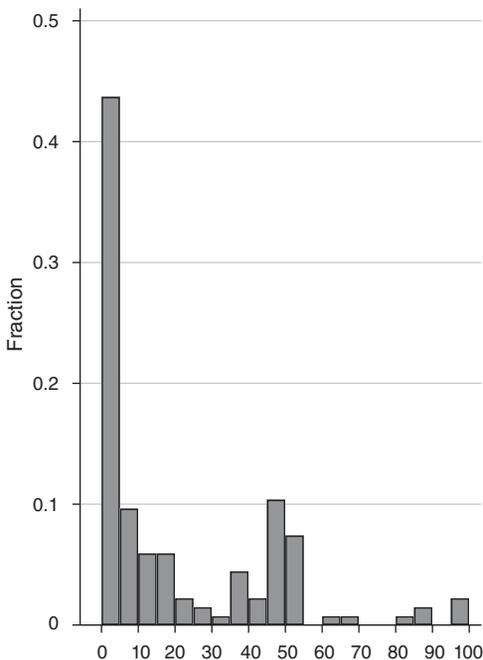


FIGURE 9. DISTANCE BETWEEN STATED AND PBE BELIEFS

beliefs about the ability of others to correctly observe (or interpret) the actions of their predecessors. The main rationales for our OC specification are that the model is simple (it has just one parameter,  $k$ ) and it explains the data well.

Second, as we pointed out in Section I, our OC model describes agents who are overconfident in a relative form: they believe they have a higher ability to understand the private signal (or to act upon it) than their predecessors. An alternative definition of overconfidence is in absolute value, that is, the agent is overconfident in his own signal, attributing to it a precision higher than the objective one. The results of our experiments support relative overconfidence and not overconfidence in the own signal. The clearest evidence that overconfidence in the own signal is rejected is at time 1 since the subject only observes his signal and does not have to weigh his signal relative to other information: as we have shown in Figure 2, the estimated coefficient for time 1 is 1, indicating that the median action is perfectly in line with the Bayesian one (and there is no overconfidence). The estimated coefficients at later periods confirm this finding.

Third, our OC model shares some similarity with the Quantal Response Equilibrium (McKelvey and Palfrey 1995) in that both allow agents to believe that others make mistakes. In the QRE, however, there is the extra restriction that beliefs about the error rates are correct, a restriction not imposed in our analysis. As a matter of fact, the restriction is also rejected by the data. The cases in which subjects updated in the wrong direction (i.e.,  $a_t^2 > a_t^1$  after observing a bad signal or  $a_t^2 < a_t^1$  after observing a good signal) amount to only 6 percent (a percentage approximately constant across periods). In this respect, our OC model is similar in spirit to the Subjective Quantal Response Equilibrium (SQRE) (Goeree, Holt, and Palfrey 2016, chap. 3), in which agents may have a misconceived (or subjective) view about the noise parameter defining the distribution of mistakes of the other agents. Our OC model is in line with such an extension of QRE. Yet, it is simpler while providing a good fit for the observed data.

Fourth, De Filippis et al. (2021) use some of the experimental data analyzed in this paper (and data from other experiments) to show that at time 2 subjects update their private information in an asymmetric way, depending on whether it confirms or contradicts the belief formed on the observation of time 1's action only. De Filippis et al. (2021) study this issue at a considerable level of detail in that paper since it is an important aspect of the updating used by subjects. In the present analysis, in which we study decisions at any time, we have abstracted from this issue. Our focus here is on social learning, that is, on how subjects learn from others (and form their "first belief,"  $a_t^1$ ) rather than on how they update on their private information (and form their "posterior belief,"  $a_t^2$ ). Moreover, while at time 2 the meaning of contradicting and confirming signal is well defined (since there is just one predecessor), at later periods it becomes less clear. In an attempt to fit the data better, one could, perhaps, incorporate asymmetric updating in our framework, but our results seem already to be clearly supporting one model and rejecting others.

Fifth, while we have formally tested BRTNI, our results are clearly at odds with any model of information redundancy neglect. Our analysis shows that early actions do not have a disproportionate effect on later decisions (early signals have the same weight as later signals). Therefore, our results should not be considered as a

rejection of the specific way redundancy neglect is modeled in BRTNI but rather as an illustration that at least in the context studied in our experiment subjects do not seem to show redundancy neglect.

Finally, for the ABEE, it is perhaps not so surprising that it does not offer good predictions in this type of experiment. The assumption that agents only consider the aggregate frequencies of actions seems plausible in real-world contexts, in which access to other statistics may even be difficult, but is less plausible in a laboratory experiment in which a subject stares at the sequence of actions other participants have chosen before him. It could be interesting to study whether presenting subjects with aggregate statistics changed their behavior and made it more in line with the ABEE. This is left for future research.

## REFERENCES

- Anderson, Lisa R., and Charles A. Holt.** 1997. "Information Cascades in the Laboratory." *American Economic Review* 87 (5): 847–62.
- Angrisani, Marco, Antonio Guarino, Philippe Jehiel, Toru Kitagawa.** 2021. "Replication data for: Information Redundancy Neglect versus Overconfidence: A Social Learning Experiment." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. <https://doi.org/10.3886/E111281V1>.
- Banerjee, Abhijit V.** 1992. "A Simple Model of Herd Behavior." *Quarterly Journal of Economics* 107 (3): 797–817.
- Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch.** 1992. "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades." *Journal of Political Economy* 110 (5): 992–1026.
- Becker, G.M., M.H. DeGroot, and J. Marschak.** 1964. "Measuring Utility by a Single-Response Sequential Method." *Behavioral Science* 9 (3): 226–32.
- Bohren, J. Aislinn.** 2016. "Informational Herding with Model Misspecification." *Journal of Economic Theory* 163: 222–47.
- Cavatorta, E., A. Guarino, and S. Huck.** 2021. "Learning from Aggregate and Partial Information." Unpublished.
- Çelen, Boğaçhan, Sen Geng, and Huihui Li.** 2019. "Belief Error and Non-Bayesian Social Learning: Experimental Evidence." Xiamen University, Wang Yanan Institute for Studies in Economics (WISE) Working Paper 2019-07-08.
- Çelen, Boğaçhan, and Shachar Kariv.** 2004. "Distinguishing Informational Cascades from Herd Behavior in the Laboratory." *American Economic Review* 94 (3): 484–98.
- Cipriani, Marco, and Antonio Guarino.** 2005. "Herd Behavior in a Laboratory Financial Market." *American Economic Review* 95 (5): 1427–43.
- Cipriani, Marco, and Antonio Guarino.** 2009. "Herd Behavior in Financial Markets: An Experiment with Financial Market Professionals." *Journal of the European Economic Association* 7 (1): 206–33.
- Daniel, Kent, David Hirshleifer, and Avanidhar Subrahmanyam.** 1998. "Investor Psychology and Security Market Under- and Overreactions." *Journal of Finance* 53 (6): 1839–85.
- Daniel, Kent D., David Hirshleifer, and Avanidhar Subrahmanyam.** 2001. "Overconfidence, Arbitrage, and Equilibrium Asset Pricing." *Journal of Finance* 56 (3): 921–65.
- De Filippis, Roberta, Antonio Guarino, Philippe Jehiel, and Toru Kitagawa.** 2021. "Non-Bayesian Updating in a Social Learning Experiment." *Journal of Economic Theory*. <https://doi.org/10.1016/j.jet.2021.105188>.
- DeGroot, M.** 1974. "Reaching a Consensus." *Journal of American Statistical Association* 69 (345): 118–21.
- DeMarzo, Peter M., Dimitri Vayanos, and Jeffrey Zwiebel.** 2003. "Persuasion Bias, Social Influence, and Unidimensional Opinions." *Quarterly Journal of Economics* 118 (3): 909–68.
- Drehmann, Mathias, Jörg Oechssler, and Andreas Roider.** 2005. "Herding and Contrarian Behavior in Financial Markets: An Internet Experiment." *American Economic Review* 95 (5): 1403–26.
- Eyster, Erik, and Matthew Rabin.** 2010. "Naïve Herding in Rich-Information Settings." *American Economic Journal: Microeconomics* 2 (4): 221–43.

- Eyster, Erik, and Matthew Rabin.** 2014. "Extensive Imitation Is Irrational and Harmful." *Quarterly Journal of Economics* 129 (4): 1861–98.
- Eyster, Erik, Matthew Rabin, and Georg Weizsäcker.** 2015. "An Experiment on Social Mislearning." CEPR Discussion Paper 11020.
- Fischbacher, Urs.** 2007. "z-Tree: Zurich Toolbox for Ready-Made Economic Experiments." *Experimental Economics* 10: 171–78.
- Goeree, Jacob K., Charles Holt, and Thomas R. Palfrey.** 2016. *Quantal Response Equilibrium—A Stochastic Theory of Games*. Princeton, NJ: Princeton University Press.
- Goeree, Jacob K., Thomas R. Palfrey, Brian W. Rogers, and Richard D. McKelvey.** 2007. "Self-Correcting Information Cascades." *Review of Economic Studies* 74 (3): 733–62.
- Golub, Benjamin, and Matthew O. Jackson.** 2010. "Naïve Learning in Social Networks and the Wisdom of Crowds." *American Economic Journal: Microeconomics* 2 (1): 112–49.
- Guarino, Antonio, and Philippe Jehiel.** 2013. "Social Learning with Coarse Inference." *American Economic Journal: Microeconomics* 5 (1): 147–74.
- Hahn, Jinyong.** 1995. "Bootstrapping Quantile Regression Estimators." *Econometric Theory* 11 (1): 105–21.
- Hung, Angela A., and Charles R. Plott.** 2001. "Information Cascades: Replication and an Extension to Majority Rule and Conformity-Rewarding Institutions." *American Economic Review* 91 (5): 1508–20.
- Jehiel, Philippe.** 2005. "Analogy-Based Expectation Equilibrium." *Journal of Economic Theory* 123 (2): 81–104.
- Kübler, Dorothea, and Georg Weizsäcker.** 2004. "Limited Depth of Reasoning and Failure of Cascade Formation in the Laboratory." *Review of Economic Studies* 71 (2): 425–41.
- Kyle, Albert S., and F. Albert Wang.** 1997. "Speculation Duopoly with Agreement to Disagree: Can Overconfidence Survive the Market Test?" *Journal of Finance* 52 (5): 2073–90.
- Lee, In Ho.** 1993. "On the Convergence of Informational Cascades." *Journal of Economic Theory* 61 (2): 395–411.
- McKelvey, Richard D., and Thomas R. Palfrey.** 1995. "Quantal Response Equilibria for Normal Form Games." *Games and Economic Behavior* 10 (1): 6–38.
- Nöth, Markus, and Martin Weber.** 2003. "Information Aggregation with Random Ordering: Cascades and Overconfidence." *Economic Journal* 113 (484): 166–89.
- Nyarko, Yaw, and Andrew Schotter.** 2002. "An Experimental Study of Belief Learning Using Elicited Beliefs." *Econometrica* 70 (3): 971–1005.
- Odean, Terrance.** 1998. "Volume, Volatility, Price, and Profit When All Traders Are above Average." *Journal of Finance* 53 (6): 1887–1934.
- Schotter, Andrew.** 2003. "Decision-making with Naive Advice." *American Economic Review* 93 (2): 196–201.
- Schotter, Andrew, and Isabel Trevino.** 2014. "Belief Elicitation in the Laboratory." *Annual Review of Economics* 6: 103–28.
- Vives, Xavier.** 1993. "How Fast do Rational Agents Learn?" *Review of Economic Studies* 60 (2): 329–47.
- Yildiz, Muhamet.** 2003. "Bargaining without a Common Prior: An Immediate Agreement Theorem." *Econometrica* 71 (3): 793–811.
- Yildiz, Muhamet.** 2011. "Bargaining with Optimism." *Annual Review of Economics* 3: 451–78.