

Information redundancy neglect versus overconfidence: a social learning experiment

Marco Angrisani
Antonio Guarino
Philippe Jehiel
Toru Kitagawa

The Institute for Fiscal Studies
Department of Economics, UCL

cemmap working paper CWP32/17

Information Redundancy Neglect versus Overconfidence: A Social Learning Experiment

Marco Angrisani, Antonio Guarino,
Philippe Jehiel and Toru Kitagawa*

Abstract

We study social learning in a continuous action space experiment. Subjects, acting in sequence, state their belief 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'.

1 Introduction

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 and Bikhchandani *et al.*, 1992), with a continuous action space

* Angrisani: Center for Economic and Social Research, University of Southern California (e-mail: angrisan@usc.edu); Guarino: Department of Economics, University College London (e-mail: a.guarino@ucl.ac.uk); Jehiel: Paris School of Economics and Department of Economics, University College London (e-mail: jehiel@enpc.fr); Kitagawa: Department of Economics, University College London (e-mail: t.kitagawa@ucl.ac.uk). We thank participants in seminars at various institutions for useful comments and suggestions. Guarino gratefully acknowledges the financial support of the ERC and of INET.

a sequence of rational decision makers are typically able to infer the 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. As a result, observing a predecessor’s action is equivalent to observing the information he has (if agents are fully rational). 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 modelling 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.

The purpose of this work is to study social learning in a continuous action space through a series of controlled experiments. Our interest is in understanding whether the Perfect Bayesian Equilibrium first proposed by Lee (1993) or any of the bounded rationality theories just mentioned is able to predict human subjects’ behavior.

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.^{1 2} 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 (2016)’s model generalizes Eyster and Rabin (2010)’s

¹Redundancy 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.

²Eyster and Rabin (2014) illustrate the implications of rational herding in an extended set up in which more than one action is taken at the same time. Eyster *et al.* (2015) study the one-action and the multiple-action cases in a laboratory experiment. The one-action case is related to our work and we will comment more on the different results in Section 4. The two experiments are independent, ours is antecedent, conducted in 2009-10-11.

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.

Of course, 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 et al. (2003) presents a very similar idea, by letting agents update as Bayesian 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 is new and which is repetition.³

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., 2016), 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).

In a multi-stage, multi-player 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 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 *et al.*, 1998; Odean, 1998; and Daniel *et al.*, 2001) 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

³We refer the reader to DeMarzo *et al.* (2003) and Golub and Jackson (2010) for further references and discussions of the links to the psychology and sociology literatures.

literature.⁴

We contribute to the understanding of the social learning process through some controlled experiments in which the different theoretical models just discussed can be carefully 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. 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.

Whereas in previous experiments subjects’ beliefs are hidden under a binary decision, in our experiment we elicit them. Moreover, we elicit a subject’s beliefs both before and after receiving the private signal, which allows us to observe how subjects learn from others (social learning) and, separately, from their private signal. Based on the same experimental data, De Filippis *et al.* (2016) study the updating rule of the second subject in the sequence and document an asymmetric form of updating depending on whether the signal realizations of the first two subjects are the same or not. The focus of this paper is, instead, on the process of social learning by all subjects in the sequence.

The results of our experiment are supportive of the 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 the 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 the overconfidence model.

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

⁴We refer to Odean (1998) for a discussion of the vast psychology literature on overconfidence about the own private information.

that subjects take is consistent with the Bayesian one, that is, the weights on all signals are the same, and not statistically different from one. 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. Section 2 describes the theoretical model of social learning and the different theoretical predictions. Section 3 presents the experiment. Section 4 contains the results. Section 5 offers a discussion and Section 6 concludes. An Appendix contains additional material.

2 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:

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

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 $E^S[-(V - a_t)^2 \mid h_t, s_t]$; therefore, his optimal action is $a_t^* = E^S(V \mid h_t, s_t)$.⁵

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

⁵The superscript S stands for subjective, since we want to allow for subjective expectations in some of the theories discussed below.

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 .

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 decisions 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.⁶ 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*

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

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:

⁶By this we mean that the agent at time 3 acts *as if* he had observed two time-1 signals.

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*

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

That is, the agent at time t acts as if, in addition to his own signal, he had observed time i signal $t - i$ 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 $\left(\frac{q_i}{1 - q_i} \right)^{(2s_i-1)k}$ rather than $\left(\frac{q_i}{1 - q_i} \right)^{(2s_i-1)}$, where $k \in (0, 1)$. We refer to this belief as “ k -overconfidence.”⁷ 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 $\frac{a_1^{OC}}{100 - a_1^{OC}} = \left(\frac{q_1}{1 - q_1} \right)$. 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 $\frac{a_2^{OC}}{100 - a_2^{OC}} = \left(\frac{q_1}{1 - q_1} \right)^k \left(\frac{q_2}{1 - q_2} \right)$. 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 $\left(\frac{q_2}{1 - q_2} \right) = \frac{a_2^{OC}}{100 - a_2^{OC}} \left(\frac{q_1}{1 - q_1} \right)^{-k}$. He then “discounts” agent 2's action and chooses

⁷Note 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 the Discussion section.

a_3^{OC} such that $\frac{a_3^{OC}}{100-a_3^{OC}} = \left(\frac{q_1}{1-q_1}\right)^k \left(\frac{q_2}{1-q_2}\right)^k \left(\frac{q_3}{1-q_3}\right)^k$.⁸ We refer to this model of learning as the “overconfidence model” (OC). The next proposition, proven in 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*

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

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 $\frac{\left(\frac{q_i}{1-q_i}\right)^k}{1 + \left(\frac{q_i}{1-q_i}\right)^k}$ 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.⁹

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

$$\begin{aligned} \ln \frac{a_t}{100 - a_t} &= \sum_{i=1}^{t-1} \beta_{t,i} (2s_i - 1) \ln \left(\frac{q_i}{1 - q_i}\right) + (2s_t - 1) \ln \left(\frac{q_t}{1 - q_t}\right), \\ \ln \frac{a_t}{100 - a_t} &= \sum_{i=1}^{t-1} \gamma_{t,i} \ln \left(\frac{a_i}{100 - a_i}\right) + (2s_t - 1) \ln \left(\frac{q_t}{1 - q_t}\right). \end{aligned}$$

We summarize the relations between $\beta_{t,i}$ ’s and $\gamma_{t,i}$ ’s in the next proposition that is proven in Appendix A.

⁸For 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).

⁹To 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 subjects to make two decisions: one after observing the predecessors’ only and another after observing the signal too. The theoretical predictions concerning the first action are given in our propositions by the first part of the formulas, excluding the multiplier $\left(\frac{q_t}{1-q_t}\right)^{2s_t-1}$.

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

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

Matrix $\mathbf{\Gamma}$ is the inverse of matrix \mathbf{B} : $\mathbf{\Gamma} = \mathbf{B}^{-1}$. 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 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 ABEE, agent t chooses action a_t^{ABEE} such that

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

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

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

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

In the PBE, an agent's action is just equal to his immediate predecessor's action (belief) just 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).

3 The Experiment and the Experimental Design

3.1 The Experiment

We ran the experiment in the ELSE Experimental Laboratory at the Department of Economics at University College London (UCL) in the fall 2009, winter 2010 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. 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 a built-on-purpose software.

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

1. Each session consisted of fifteen rounds. At the beginning of each round, the computer program randomly chose the value of a good. The value was equal to 0 or 100 with the same probability, independently of previous realizations.
2. In each round we asked all subjects to make decisions in sequence, one after the other. For each round, the sequence was randomly chosen by the computer software. Each subject had an equal probability of being chosen in any position in the sequence.
3. Participants were not told the value of the good. They knew, however, that they would receive information about the value, in the form of a symmetric

binary signal. If the value was equal to 100, a participant would receive a “green ball” with probability 0.7 and a “red ball” with probability 0.3; if the value was equal to 0, the probabilities were inverted. That is, the green signal corresponded to $s_t = 1$ and the red signal to $s_t = 0$, the signal precision q_t was equal to 0.7 at any time.

4. As we said, each round consisted of 10 periods. In the first period a subject was randomly chosen to make a decision. He received a signal and chose a number between 0 and 100, up to two decimal points.
5. The other subjects observed the decision made by the first subject on their screens. The identity of the subject was not revealed.
6. In the second period, a second subject was randomly selected. He was asked to choose a number between 0 and 100, having observed the first subject’s choice only.
7. After he had made that choice, he received a signal and had to make a second decision. This time, therefore, the decision was based on the observation of the predecessor’s action and of the private signal.
8. In the third period, a third subject was randomly selected and asked to make two decisions, similarly to the second subject: a first decision after observing the choice of the first subject and the second choice of the second subject; a second decision after observing the private signal too. The same procedure was repeated for all the remaining periods, until all subjects had acted. Hence, each subject, from the second to the tenth, made two decisions: one after observing the history of all (second) decisions made by the predecessors; the other after observing the private signal too.
9. At the end of the round, after all 10 subjects had made their decisions, subjects observed a feedback screen, in which they observed the value of the good and their own payoff for that round. The payoffs were computed 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.¹⁰

As should be clear from this description, compared to the existing experimental literature on social learning / informational cascades / herd behavior,

¹⁰This 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).

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.¹¹ Second, in previous experiments subjects made one decision after observing both the predecessors' actions and the signal. In our experiment, instead, they made two decisions, one based on public information only and one based on the private information as well.¹²

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

3.2 Experimental Design

Social Learning (SL). In addition to the social learning treatment (SL1) just described, we ran a second treatment (SL2) which only differed from the first because the signal had a precision which was randomly drawn in the interval $[0.7, 0.71]$ 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 10 subjects, we had sequences of 4 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 Appendix B). Therefore, we consider the three treatments as just

¹¹Within the discrete action space experiments, exceptions to the binary action space are the financial market experiments of Cipriani and Guarino (2005, 2009) and Drehman *et al.* (2005) where subjects can choose to buy, to sell or not to trade. In the interesting experimental design of Çelen and Kariv (2004), subjects choose a cut off value in a continuous signal space: depending on the realization of the signal, one of the two actions is implemented (as in a Becker, DeGroot and Marschak, 1964, mechanism). That design allows the authors to distinguish herd behavior from informational cascades.

¹²Cipriani and Guarino (2009) use a quasi strategy method, asking subject to make decisions conditional on either signal they might receive. Still, at each time, a subject never makes a decision based only on the predecessors' decisions.

one experimental condition.¹³ We will refer to it as the SL treatment.¹⁴

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 one decision.¹⁵ Specifically, a subject received a signal (as subject 1 in the SL treatments) and had to make a choice in the interval $[0, 100]$. Then, with a 50% probability, he received another signal and had to make a second decision (similarly to the second decision of subject 2 in the SL treatments). Then, he could make two more decisions, and the probability of moving from one decision to the next was always a 50%. 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 possible as possible to that of a subject in the SL treatments. In other words, we wanted the subject to make his first decision not knowing whether he would be asked to make a second one; 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: 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

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 one subject.

¹³Drawing the precision from the tiny interval $[0.7, 0.71]$, instead of having the simpler set up with fixed precision equal to 0.7, was 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 4 subjects was instead motivated by the opportuneness to collect more data for the first periods of the sequence.

¹⁴Since the results of SL1 and SL2 were not statistically different, we did not run more treatments with signals of different precision. 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 less possibilities of learning).

¹⁵This treatment was conducted in the fall 2014. The payment followed the same rules. The exchange rate was appropriately modified so that, in expectation, subjects could receive a similar amount of money per hour spent in the laboratory.

4 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 2 for the theoretical models (since in that section we only considered the theoretical action after observing the predecessors and the own signal).

4.1 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 propositions in Section 2 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.

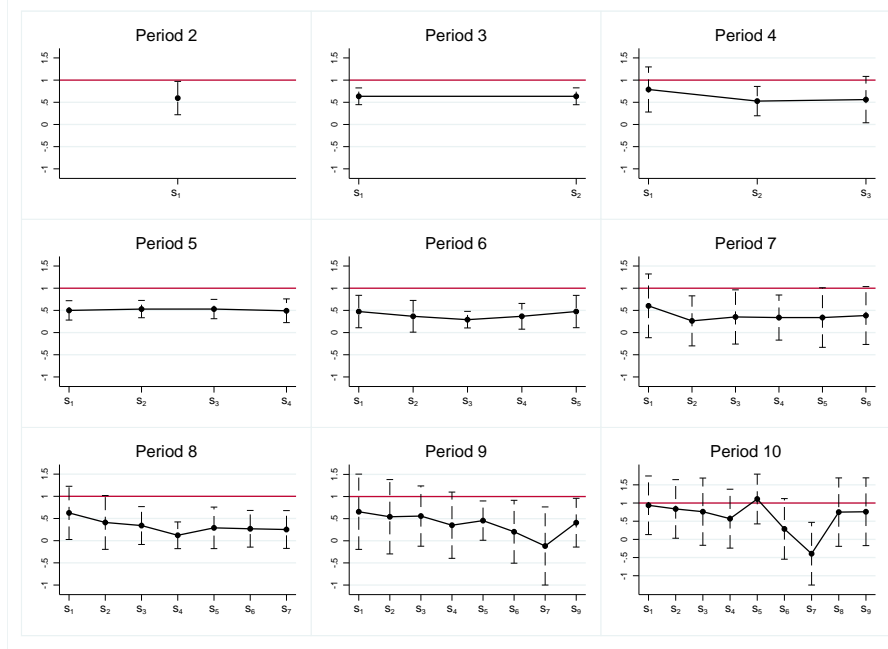
We use median regressions throughout the analysis. Specifically, for each period $t = 2, 3, \dots, 10$, we regress the loglikelihood ratio of a_t^1 on all the predecessors' signal likelihood ratios:

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

with $Med(\varepsilon_t | s_1, s_2, \dots, s_{t-1}) = 0$. Each coefficient $\beta_{t,i}$ is the weight given by the median agent t to signal s_i ($i < t$). 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, captured by our regression. In other words, we use a median regression rather than a linear 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 $t = 10$.¹⁶ To account for unobserved

¹⁶Recall that in one the session we had 9 rather than 10 subjects due to a last minute no show up by one subject.

Figure 1: Median Regressions of Action 1 on Predecessors' Signals
(Estimated Weights)



The figure shows the estimated coefficients from a median regression of action 1 on predecessors' signals. 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% 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.

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% 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% 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.

We estimate the “degree of overconfidence” k , under the hypothesis that at any time t the weight is constant for all signals 1, 2, $t - 1$, as predicted by the OC model. Table 3 reports the results. Estimates are significantly

Table 2: Hypothesis Testing: Weights on Predecessors' Signals (p-values)
Dependent Variable: Action 1

	$H_0^{PBE} :$ $\beta_{t,1} = \dots = \beta_{t,t-1} = 1$	$H_0^{BRTNI} :$ $\beta_{t,i} = 2^{t-i-1} \forall i = 1, \dots, t-1$
Period 2	0.035	0.035
Period 3	0.000	0.000
Period 4	0.023	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	0.003	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	0.035	.
Period 3	0.000	0.871
Period 4	0.000	0.098
Period 5	0.000	0.986
Period 6	0.000	0.857
Period 7	0.000	0.805
Period 8	0.000	0.830
Period 9	0.000	0.921
Period 10	0.000	0.262

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	0.261	0.975
Period 3	0.635	0.508	0.874
Period 4	0.674	0.503	0.997
Period 5	0.504	0.332	0.664
Period 6	0.416	0.142	0.706
Period 7	0.404	0.180	0.648
Period 8	0.358	0.200	0.554
Period 9	0.381	0.257	0.649
Period 10	0.489	0.301	0.997
All	0.488	0.327	0.706

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

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 0.09, that is we cannot reject the hypothesis that the parameters are equal (at 5%). 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 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.¹⁷ To tackle this issue and check the robustness 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

¹⁷This 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.

not change compared to the present ones. We refer the reader to Appendix C, where we discuss the precise methodology adopted and report the corresponding results.

4.2 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:

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

As one can see from Figure 2, 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 4 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 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 summary, our analysis shows that the median subject puts the correct weight of 1 on his own signal and approximately half the weight on his predecessors' signals.¹⁸

4.3 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

¹⁸Our results are somehow at odds with those in Eyster *et al.* (2015) in which the PBE is quite successful in predicting behavior in what they call the single-file treatment (one action per time). It should be noted though, that the subjects' task in that experiment is easier than in ours: subjects have to sum up signals, whereas in our case there is an inference problem. The different difficulty in the task is likely to be at the root of the different predictive power of the PBE.

Figure 2: Median Regressions of Action 2 on Own and Predecessors' Signals
(Estimated Weights)




The figure shows the estimated coefficients from a median regression of action 1 on own and predecessors' signals. 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% 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.

Table 4: Hypothesis Testing: Weights on Own and Predecessors' Signals (p-values)
Dependent Variable: Action 2

	$H_0^{PBE} :$ $\beta_{t,1} = \dots = \beta_{t,t} = 1$	$H_0^{BRTNI} :$ $\beta_{t,i} = 2^{t-i-1} \forall i = 1, \dots, t-1,$ $\beta_{t,t} = 1$
Period 1	0.977	0.977
Period 2	0.000	0.000
Period 3	0.000	0.000
Period 4	0.040	0.000
Period 5	0.006	0.000
Period 6	0.000	0.000
Period 7	0.301	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	0.977	0.977
Period 2	0.000	0.376 [†]
Period 3	0.000	0.414
Period 4	0.000	0.584
Period 5	0.000	0.994
Period 6	0.000	0.940
Period 7	0.000	0.941
Period 8	0.000	0.232
Period 9	0.000	0.887
Period 10	0.000	0.004

The table reports tests based on bootstrap standard errors (500 replications), clustering at the session level. †: 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% confidence interval of [0.322, 0.630].

Figure 3: Median Regressions of Action on Own Signals in IDM Treatment (Estimated Weights)



The figure shows the estimated coefficients from a median regression of subjects' action on own signals. For each period $t = 1, 2, 3, 4$, own signals, s_t , are on the x-axis; corresponding point estimates and 95% 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.

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 observed directly a sequence of signals (IDM treatment). Figure 3 shows the results. Subjects put the same weight on all signals, a result which 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 0.30).


Overall, the results of this treatment confirm that our findings in the SL treatment are indeed related to how they 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).

4.4 Social Learning

The results presented in the previous sections are a direct test of our first four theoretical propositions. Note that this is 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. In this section, we study how subject t 's action is related not to his predecessors' signals, but to his predecessors' actions. There are various reasons to do this. First, predecessors' actions are observed by a subject, whereas predecessors' signals are not. Second, if subjects behaved as in theory, the two analyses would lead to the same results, but if they do not (e.g., because of mistakes in choosing the action), an analysis of how subjects react to the predecessors' actions can reveal a behavior not detectable by focusing on the 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.¹⁹ 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 matter ($a_t^2 = a_{t-1}^2$), and BRTNI predicts that subjects take

¹⁹Of 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 actions and the predecessors' signals, whereas it would be by considering that between an action and the predecessors' actions.

Figure 4: Median Regressions of Action 1 on Predecessors' Action 2
(Estimated Weights)



The figure shows the estimated coefficients from a median regression of action 1 on predecessors' action 2. 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% 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.

actions as signals. Essentially, by studying the relation between an action and the predecessors' actions, we do not need to rely on assumptions 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 regression equations:

$$\begin{aligned} \ln \left(\frac{a_t^1}{100 - a_t^1} \right) &= \gamma_{t,1} \ln \left(\frac{a_1^2}{100 - a_1^2} \right) + \gamma_{t,2} \ln \left(\frac{a_2^2}{100 - a_2^2} \right) + \dots \\ &+ \gamma_{t,t-1} \ln \left(\frac{a_{t-1}^2}{100 - a_{t-1}^2} \right) + \varepsilon_t. \end{aligned} \quad (3)$$

Table 5: Hypothesis Testing: Weights on Predecessors' Action 2 (p-values)
Dependent Variable: Action 1

	$H_0^{PBE} :$ $\gamma_{t,1} = \dots = \gamma_{t,t-2} = 0,$ $\gamma_{t,t-1} = 1$	$H_0^{BRNI} :$ $\gamma_{t,1} = \dots = \gamma_{t,t-1} = 1$
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
	$H_0^{ABEE} :$ $\gamma_{t,i} = \text{sign}(\sin(\frac{t-i}{3}\pi))$ $\forall i = 1, \dots, t-1$	$H_0^{OC} :$ $\gamma_{t,i} = \gamma_{t-1}(1 - \gamma_{t-1})^{t-i-1}$ $\forall i = 1, \dots, t-1$
Period 2	0.000	.
Period 3	0.000	0.111
Period 4	0.000	0.251
Period 5	0.000	0.234
Period 6	0.000	0.523
Period 7	0.000	0.134
Period 8	0.000	0.867
Period 9	0.000	0.028
Period 10	0.000	0.591

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

Table 6: Median Regressions of Action 1 on Predecessors' Action 2: Estimation of k under $H_0^{OC} : \gamma_{t,i} = \gamma_{t-1}(1 - \gamma_{t-1})^{t-i-1} \forall i = 1, \dots, t-1$

	\hat{k}	95% Confidence Interval	
		lower limit	upper limit
Period 3	0.469	0.204	0.598
Period 4	0.580	0.248	0.840
Period 5	0.154	0.127	0.394
Period 6	0.154	0.077	0.801
Period 7	0.248	0.104	0.717
Period 8	0.180	0.083	0.745
Period 9	0.324	0.102	0.430
Period 10	0.211	0.132	0.336
All	0.320	0.156	0.526

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

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

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 4 and Tables 5-6. 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.²⁰

Finally, in Figure 5 and Table 7 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.²¹ This model is never rejected,

²⁰In the case of constant k , the p-value is 0.08.

²¹It is worth mentioning that whereas in the social learning literature, as in much psychological literature, researchers have talked about “overconfidence,” in other experimental studies subjects show “underconfidence.” In particular, in experiments on decision making with naive advice, it has been observed that “when given a choice between getting advice or the information upon which the advice is based, subjects tend to opt for the advice, indicating a kind of underconfidence in their decision making abilities [...]” (Schotter, 2003). Our result is in

Figure 5: Median Regressions of Action 2 on Predecessors' Action 2 and Own Signal
(Estimated Weights)



The figure shows the estimated coefficients from a median regression of action 2 on predecessors' action 2 and own signals. 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% 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.

Table 7: Hypothesis Testing: Weights on Predecessors' Action 2 and Own Signal (p-values)
Dependent Variable: Action 2

	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	0.002	0.000
Period 6	0.000	0.000
Period 7	0.016	0.000
Period 8	0.000	0.000
Period 9	0.000	0.000
Period 10	0.000	0.000
	H_0^{ABEE} : $\gamma_{t,i} = \text{sign}(\sin(\frac{t-i}{3}\pi))$ $\forall i = 1, \dots, t-1$ $\beta_{t,t} = 1$	H_0^{OC} : $\gamma_{t,i} = \gamma_{t-1}(1 - \gamma_{t-1})^{t-i-1}$ $\forall i = 1, \dots, t-1$ $\beta_{t,t} = 1$
Period 2	0.000	.
Period 3	0.000	0.170
Period 4	0.000	0.143
Period 5	0.000	0.402
Period 6	0.000	0.944
Period 7	0.000	0.477
Period 8	0.000	0.862
Period 9	0.000	0.164
Period 10	0.000	0.017

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

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

4.5 Efficiency and Convergence

At last, 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' action 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 the other theoretical models we have presented, in contrast, 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 sequences of 4 or 10 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 6 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.²² In round 1, subjects earn 96% of what is potentially obtainable. Over rounds, this percentage declines but only slightly, to over around 90% in the final periods of the experiment. The figure also reports the ratio between the average OC, ABEE and BRTNI payoffs and the average PBE payoffs.²³ 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 heterogenous use of the own signal, the OC model 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 7 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.,

favour of overconfidence and at odds with underconfidence.

²²In toehr words, for each period t , the PBE payoff is computed as $100 - 0.01 (V - a_{ti}^{PBE})^2$ for each observation i and averaging across all observations.


²³The average theoretical payoffs are computed analogously to what explained in the previous footnote for the PBE. For the OC model, we used the estimated value $k = 0.320$.

Figure 6: Ratio of Average Payoff to Average PBE Payoff



To compute payoffs under the OC model, we assume a value of the overconfidence k parameter equal to 0.320 (last row in Table 6).

Figure 7: Distance between Stated and PBE Beliefs



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% of the cases in period 4 (SL3) and in 44% 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.

5 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, 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 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 depends on just one parameter, k) and it explains the data well.

Second, as we pointed out in Section 2, 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 (QRE; 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% (a percentage approximately constant across periods). In this respect, our OC model is similar in spirit to the Subjective Quantal Response Equilibrium (SQRE; Goeree *et al.*, chapter 3), in which agents may have a misconceived (or subjective) view about the noise parameter defining the distribution of mistakes

of the other agents. This feature is also present in the experimental work on the formation of informational cascades by Kübler and Weizsäcker (2004). 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.* (2016) use the same experimental data 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.* (2016) 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, the OC model we have presented makes the PBE more flexible, by letting agents have a subjective belief on the predecessors' signal precisions. One may contemplate the possibility of having such a flexibility for the BRTNI and the ABEE models too. It would, however, be difficult to give a reasonable interpretation of these amended models. The founding idea of BRTNI is that agents take another agent's action as a signal. It is not so natural to think that they then discount this mis-constructed signal. Similarly, the ABEE is obtained assuming that agents only consider the aggregate frequencies of actions given a state of the world (the value of the good, in our set up). It is unnatural to think that then they add an error rate to the decisions they observe. In other words, for both these theories of bounded rationality, a process of introspection would probably let the agents to cast doubts on the theories in the first place and rethink their process of decision making, rather than adding error rates to the actions they observe.

Sixth, 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.

Appendix (for online publication)

A Proofs

Proof of Proposition 4

The proposition can be proven in a recursive way. Agent 1 only observes his signal and chooses

$$\frac{a_1^{OC}}{100 - a_1^{OC}} = \frac{a_1^{****}(s_1)}{100 - a_1^{****}(s_1)} = \left(\frac{q_1}{1 - q_1} \right)^{2s_1 - 1}.$$

Agent 2 observes a_1^{OC} and infers the signal realization, since an action greater (lower) than 50 can only be taken after observing a good (bad) signal. By the assumption of “k-overconfidence,” he has subjective expectations on the predecessor’s signal precision, and the likelihood ratio after observing the action is $\left(\frac{q_1}{1 - q_1} \right)^{(2s_1 - 1)k}$ rather than $\left(\frac{q_1}{1 - q_1} \right)^{(2s_1 - 1)}$. Hence,

$$\frac{a_2^{OC}}{100 - a_2^{OC}} = \frac{a_2^{****}(s_1, s_2)}{100 - a_2^{****}(s_1, s_2)} = \left(\frac{q_1}{1 - q_1} \right)^{(2s_1 - 1)k} \left(\frac{q_2}{1 - q_2} \right)^{2s_2 - 1}.$$

Note that this is equivalent to attributing precision $\frac{\left(\frac{q_1}{1 - q_1} \right)^k}{1 + \left(\frac{q_1}{1 - q_1} \right)^k}$ to the predecessor’s signals. Since k-overconfidence is common knowledge, agent 3 infers the signal realizations from the observation of a_1^{OC} and a_2^{OC} (since $a_2^{OC} > a_1^{OC}$ is only possible after observing a signal $s_2 = 1$, and $a_2^{OC} < a_1^{OC}$ after observing a signal $s_2 = 0$) and again uses subjective expectations for the precision of both, thus choosing a_3^{OC} such that

$$\frac{a_3^{OC}}{100 - a_3^{OC}} = \frac{a_3^{****}(s_1, s_2, s_3)}{100 - a_3^{****}(s_1, s_2, s_3)} = \prod_{i=1}^2 \left(\frac{q_i}{1 - q_i} \right)^{(2s_i - 1)k} \left(\frac{q_3}{1 - q_3} \right)^{2s_3 - 1}.$$

The same steps apply to any further agent $t = 4, 5, \dots, T$.

Proof of Proposition 5

Let us define $l(x) := \log \frac{x}{1-x}$. First, observe that, for each $t \geq 2$, the β coefficients are determined by the following equations:

$$\begin{aligned} l(a_1^2) &= (2s_1 - 1)l(q_1), \\ l(a_2^2) &= \beta_{2,1}(2s_1 - 1)l(q_1) + (2s_2 - 1)l(q_2) \\ &\vdots \\ l(a_t^2) &= \beta_{t,1}(2s_1 - 1)l(q_1) + \dots + \beta_{t,t-1}(2s_t - 1)l(q_{t-1}) + (2s_t - 1)l(q_t). \end{aligned}$$

In matrix notation,

$$\begin{pmatrix} l(a_1^2) \\ l(a_2^2) \\ \vdots \\ l(a_t^2) \end{pmatrix} = B \cdot \begin{pmatrix} (2s_1 - 1)l(q_1) \\ (2s_2 - 1)l(q_2) \\ \vdots \\ (2s_t - 1)l(q_t) \end{pmatrix}, \quad (\text{A.1})$$

where B is the $t \times t$ lower triangular matrix

$$B = \begin{pmatrix} 1 & 0 & \cdots & \cdots & 0 \\ \beta_{2,1} & 1 & 0 & \cdots & 0 \\ \beta_{3,1} & \beta_{3,2} & 1 & & 0 \\ \vdots & \vdots & & \ddots & \vdots \\ \beta_{t,1} & \beta_{t,2} & \cdots & \beta_{t,t-1} & 1 \end{pmatrix}.$$

Similarly, the γ coefficients are defined by the following equations:

$$\begin{aligned} l(a_1^2) &= (2s_1 - 1)l(q_1), \\ l(a_2^2) &= \gamma_{2,1}l(a_1^2) + (2s_2 - 1)l(q_2), \\ &\vdots \\ l(a_t^2) &= \gamma_{t,1}l(a_1^2) + \cdots + \gamma_{t,t-1}(2s_t - 1)l(a_{t-1}^2) + (2s_t - 1)l(q_t). \end{aligned}$$

In matrix notation,

$$\Gamma \begin{pmatrix} l(a_1^2) \\ l(a_2^2) \\ \vdots \\ l(a_t^2) \end{pmatrix} = \begin{pmatrix} (2s_1 - 1)l(q_1) \\ (2s_2 - 1)l(q_2) \\ \vdots \\ (2s_t - 1)l(q_t) \end{pmatrix},$$

where Γ is the $t \times t$ lower triangular matrix containing γ coefficients,

$$\Gamma = \begin{pmatrix} 1 & 0 & \cdots & \cdots & 0 \\ -\gamma_{2,1} & 1 & 0 & \cdots & 0 \\ -\gamma_{3,1} & -\gamma_{3,2} & 1 & & 0 \\ \vdots & \vdots & & \ddots & \vdots \\ -\gamma_{t,1} & -\gamma_{t,2} & \cdots & -\gamma_{t,t-1} & 1 \end{pmatrix}. \quad (\text{A.2})$$

By comparing (A.1) with (A.2), one can see that, since B is nonsingular, $\Gamma = B^{-1}$ must hold. Hence, for $l < t$, $-\gamma_{t,l}$ is given by the $[t, l]$ -element of B^{-1} .

The closed form solutions for $\gamma_{y,i}$ for each theory can also be obtained in a recursive way.

For the PBE, note that agent t chooses action a_t^{PBE} such that

$$\begin{aligned} \frac{a_t^*(s_1, s_2, \dots, s_t)}{100 - a_t^*(s_1, s_2, \dots, s_t)} &= \prod_{i=1}^t \left(\frac{q_i}{1 - q_i} \right)^{2s_i - 1} = \\ &= \prod_{i=1}^{t-1} \left(\frac{q_i}{1 - q_i} \right)^{2s_i - 1} \left(\frac{q_t}{1 - q_t} \right)^{2s_t - 1} = \\ &= \frac{a_{t-1}^*(s_1, s_2, \dots, s_{t-1})}{100 - a_{t-1}^*(s_1, s_2, \dots, s_{t-1})} \left(\frac{q_t}{1 - q_t} \right)^{2s_t - 1}. \end{aligned}$$

1. For BRTNI, observe that, by assumption, agent t chooses action a_t^{BRTNI} such that

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

(Indeed, Eyster and Rabin (2009) derive the β coefficients from this formula).

In the OC model, agent 2 chooses action a_2^{OC} such that

$$\begin{aligned} \frac{a_2^{OC}}{100 - a_2^{OC}} &= \left(\frac{q_1}{1 - q_1} \right)^{(2s_1 - 1)k} \left(\frac{q_2}{1 - q_2} \right)^{2s_2 - 1} = \\ &= \left(\frac{a_1^{OC}}{100 - a_1^{OC}} \right)^k \left(\frac{q_2}{1 - q_2} \right)^{2s_2 - 1}. \end{aligned}$$

Agent 3 chooses action a_3^{OC} such that

$$\begin{aligned} \frac{a_3^{OC}}{100 - a_3^{OC}} &= \prod_{i=1}^2 \left(\frac{q_i}{1 - q_i} \right)^{(2s_i - 1)k} \left(\frac{q_3}{1 - q_3} \right)^{2s_3 - 1} = \\ &= \left(\frac{a_1^{OC}}{100 - a_1^{OC}} \right)^k \left(\frac{q_2}{1 - q_2} \right)^{(2s_2 - 1)k} \left(\frac{q_3}{1 - q_3} \right)^{2s_3 - 1} = \\ &= \left(\frac{a_1^{OC}}{100 - a_1^{OC}} \right)^k \left(\frac{a_2^{OC}}{100 - a_2^{OC}} \right)^k \left(\frac{a_1^{OC}}{100 - a_1^{OC}} \right)^{-k^2} \left(\frac{q_3}{1 - q_3} \right)^{2s_3 - 1} = \\ &= \left(\frac{a_1^{OC}}{100 - a_1^{OC}} \right)^{k(1-k)} \left(\frac{a_2^{OC}}{100 - a_2^{OC}} \right)^k \left(\frac{q_3}{1 - q_3} \right)^{2s_3 - 1}. \end{aligned}$$

The same steps apply to any further agent $t = 4, 5, \dots, T$.

Finally, let us consider the ABEE. First of all, recall that in the ABEE $\beta_{t,i} = t - i$, that is, $\beta_{t,t-k} = k$ for all $t = 2, 3, \dots$, and $k = 1, 2, \dots, (t - 1)$.

Consider now the system of equations $\Gamma B = I$. For $t = 2, 3, 4, \dots$, the product of the t -th row vector of Γ and the $(t - 1)$ -th column vector of B gives

$$-\gamma_{t,t-1} + \beta_{t,t-1} = 0,$$

from which we obtain that $\gamma_{t,t-1} = 1$. For $t = 3, 4, 5, \dots$, the product of the t -th row vector of Γ and the $(t-2)$ -th column vector of B gives

$$-\gamma_{t,t-2} - \gamma_{t,t-1}\beta_{t-1,t-2} + \beta_{t,t-2} = 0,$$

from which we obtain that $\gamma_{t,t-2} = 1$. For $t = 4, 5, 6, \dots$, the product of the t -th row vector of Γ and the $(t-3)$ -th column vector of B gives

$$-\gamma_{t,t-3} - \gamma_{t,t-2}\beta_{t-2,t-3} - \gamma_{t,t-1}\beta_{t-1,t-3} + \beta_{t,t-3} = 0,$$

from which we obtain that $\gamma_{t,t-3} = 0$.

Now, let us consider all $t = 5, 6, 7, \dots$, and $k = 4, 5, 6, \dots, (t-1)$. The product of the t -th row vector of Γ and the $(t-k)$ -th column vector of B gives

$$\begin{aligned} \gamma_{t,t-k} &= -\sum_{j=1}^{k-1} \gamma_{t,t-k+j}\beta_{t,t-j} + \beta_{t,t-k} \\ &= -\sum_{j=1}^{k-1} \gamma_{t,t-k+j}j + k. \end{aligned}$$

On the basis of this equation, observe that the difference of $\gamma_{t,t-k-1}$ and $\gamma_{t,t-k}$ gives

$$\gamma_{t,t-k-1} - \gamma_{t,t-k} = -\gamma_{t,t-k} - \gamma_{t,t-k+1} - \gamma_{t,t-k+2} - \dots - \gamma_{t,t-1} + 1.$$

Similarly, the difference between $(\gamma_{t,t-k-2} - \gamma_{t,t-k-1})$ and $(\gamma_{t,t-k-1} - \gamma_{t,t-k})$ gives

$$\gamma_{t,t-k-2} = \gamma_{t,t-k-1} - \gamma_{t,t-k}.$$

Moreover, the sum of $\gamma_{t,t-k-2}$ and $\gamma_{t,t-k-3}$ gives

$$\gamma_{t,t-k-3} = -\gamma_{t,t-k}.$$

Hence, starting from the three initial values, $\gamma_{t,t-1} = \gamma_{t,t-2} = 1$ and $\gamma_{t,t-3} = 0$, this equation iteratively pins down the whole sequence of $(\gamma_{t,t-1}, \gamma_{t,t-2}, \dots, \gamma_{t,1})$. Specifically, $(\gamma_{t,t-4}, \gamma_{t,t-5}, \gamma_{t,t-6}) = (-1, -1, 0)$, $(\gamma_{t,t-7}, \gamma_{t,t-8}, \gamma_{t,t-9}) = (1, 1, 0)$, $(\gamma_{t,t-10}, \gamma_{t,t-11}, \gamma_{t,t-12}) = (-1, -1, 0)$, and so on. For instance, for subject 10, the weights are $(\gamma_{10,1}, \gamma_{10,2}, \gamma_{10,3}, \dots, \gamma_{10,9}) = (0, 1, 1, 0, -1, -1, 0, 1, 1)$.

Finally note that, given its cyclical feature, the sequence of weights can be expressed as $\gamma_{t,t-k} = \text{sign}(\sin(\frac{k}{3}\pi))$, or $\gamma_{t,i} = \text{sign}(\sin(\frac{t-i}{3}\pi))$.

B Testing Differences across Treatments

Table B.1: Differences across Treatments:
Median Rank-sum Test for Action 1 (p-value)

	SL1 vs. SL2	SL1 vs. SL3	SL2 vs. SL3
Period 1	0.999	0.999	0.999
Period 2	0.136	0.520	0.738
Period 3	0.317	0.881	0.317
Period 4	0.738	0.597	0.829
Period 5	0.881		
Period 6	0.911		
Period 7	0.316		
Period 8	0.289		
Period 9	0.435		
Period 10	0.420		

For each period, the test is performed using session-specific medians.

Table B.2: Differences across Treatments:
Median Rank-sum Test for Action 2 (p-value)

	SL1 vs. SL2	SL1 vs. SL3	SL2 vs. SL3
Period 1	0.459	0.834	0.751
Period 2	0.220	0.999	0.243
Period 3	0.218	0.345	0.914
Period 4	0.281	0.244	0.117
Period 5	0.911		
Period 6	0.599		
Period 7	0.023		
Period 8	0.590		
Period 9	0.529		
Period 10	0.805		

For each period, the test is performed using session-specific medians.

C Factoring Out Uninformative Actions

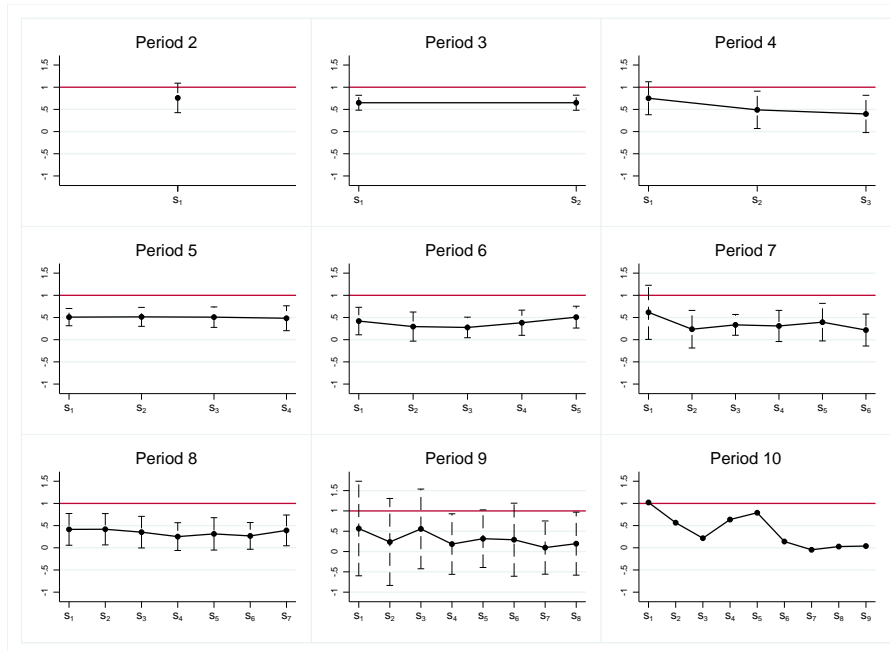
In this section we offer a robustness check, by factoring out actions that are, presumably, uninformative. In particular, as a first step, we define an action at time i as uninformative according to the criterion:

$$(a_{i=1}^2 = 50) \text{ or } (a_i^2 = a_{i-1}^2) \text{ for } i = 2, \dots, t-1.$$

To factor out these actions, we eliminate them and renumber the entire sequence (e.g., if action 3 is uninformative, then action 3 is eliminated, period 4 becomes period 3, period 5 becomes period 4 and so on).

It is worth noting that this procedure implies a loss of observations for later periods. In particular, the available observations for $t = 10$ are 47. Coefficients for this period are not reliably estimated. We report them without confidence intervals and only for the sake of completeness. For the same reason, we do not report hypothesis testing p-values and estimates of k for this period.

Figure C.1: Quantile Regressions of Action 1 on Predecessors' Signals
Eliminating Uninformative Periods (Estimated Weights)



The figure shows the estimated coefficients from a median regression of action 1 on predecessors' signals after eliminating uninformative periods. 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% 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.

Table C.1: Hypothesis Testing: Weights on Predecessors' Signals (p-values)
 Dependent Variable: Action 1
 Eliminating Uninformative Periods

	$H_0^{PBE} :$ $\beta_{t,1} = \dots = \beta_{t,t-1} = 1$	$H_0^{BRTNI} :$ $\beta_{t,i} = 2^{t-i-1} \forall i = 1, \dots, t-1$
Period 2	0.155	0.155
Period 3	0.000	0.000
Period 4	0.004	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
	$H_0^{ABEE} :$ $\beta_{t,i} = t - i \forall i = 1, \dots, t-1$	$H_0^{OC} :$ $\beta_{t,1} = \dots = \beta_{t,t-1}$
Period 2	0.155	.
Period 3	0.000	0.999
Period 4	0.000	0.094
Period 5	0.000	0.993
Period 6	0.000	0.583
Period 7	0.000	0.620
Period 8	0.000	0.995
Period 9	0.000	0.995

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

We have repeated the analysis using a more stringent criterion according to which an action i is classified as uninformative if and only if

$$(a_{i=1}^2 = 50) \text{ or } (a_i^2 = a_{i-1}^2) \text{ or } (a_{i-1}^2 = 0 \text{ or } a_{i-1}^2 = 100) \text{ for } i = 2, \dots, t-1.$$

The results are similar to those presented here and available upon request.

We have also used a different methodology, by attributing the value $s_i = 0.5$ (uninformative signal) to any uninformative action. The results are again broadly similar to those presented here and available upon request.

Table C.2: Quantile Regressions of Action 1 on Predecessors' Signals:
 Estimation of k under $H_0^{OC} : \beta_{t,1} = \dots = \beta_{t,t-1}$
 Eliminating Uninformative Periods

	\hat{k}	95% Confidence Interval	
		lower limit	upper limit
Period 2	0.752	0.460	0.996
Period 3	0.650	0.518	0.821
Period 4	0.559	0.438	0.997
Period 5	0.508	0.332	0.648
Period 6	0.422	0.250	0.547
Period 7	0.332	0.200	0.511
Period 8	0.272	0.097	0.436
Period 9	0.351	0.001	0.997
All	0.463	0.327	0.622

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

References

- [1] Banerjee, A., 1992. A Simple Model of Herd Behavior. *Quarterly Journal of Economics*, 107(3), 797-817.
- [2] Bikhchandani, S., Hirshleifer, D., Welch, I., 1992. A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades. *Journal of Political Economy*, 110(5), 992-1026.
- [3] Becker, G., DeGroot, M., Marschak, J., 1964. Measuring Utility by a Single-Response Sequential Method. *Behavioral Science*, 9(3), 226-232.
- [4] Bohren, A., 2016. Informational Herding with Model Misspecification. *Journal of Economic Theory*, 163, 222-247.
- [5] Cipriani, M., Guarino, A., 2005. Herd Behavior in a Laboratory Financial Market. *American Economic Review*, 95(5), 1427-1443.
- [6] Cipriani, M., Guarino, A., 2009. Herd Behavior in Financial Markets: An Experiment with Financial Market Professionals. *Journal of the European Economic Association*, 7(1), 206-233.
- [7] Çelen, B., Kariv, S., 2004. Distinguishing Informational Cascades from Herd Behavior in the Laboratory. *American Economic Review*, 94(3), 484-498.
- [8] Daniel, K., Hirshleifer, D., Subrahmanyam, A., 1998. Investor Psychology and Security Market Under- and Overreactions. *Journal of Finance*, 53(6), 1839-1885.
- [9] Daniel, K., Hirshleifer, D., Subrahmanyam, A., 2001. Overconfidence, Arbitrage, and Equilibrium Asset Pricing. *Journal of Finance*, 56(3), 921-965.
- [10] De Filippis, R., Guarino, A., Jehiel, P., Kitagawa, T., 2016. Updating Ambiguous Beliefs in a Social Learning Experiment. *CEMMAP Working Paper 18/16*
- [11] DeGroot, M., 1974. Reaching a Consensus. *Journal of American Statistical Association*, 69, 118-121.
- [12] DeMarzo, P., Vayanos, D., Zwiebel, J., 2003. Persuasion Bias, Social Influence, and Unidimensional Opinions. *Quarterly Journal of Economics*, 118(3), 909-968.
- [13] Drehmann, M., Oechssler, J., Roeder, A., 2005. Herding and Contrarian Behavior in Financial Markets: An Internet Experiment. *American Economic Review*, 95(5), 1403-1426.
- [14] Eyster, E., Rabin, M., 2010. Naïve Herding in Rich-Information Settings. *American Economic Journal: Microeconomics*, 2(4), 221-243.

- [15] Eyster, E., Rabin, M., 2014. Extensive Imitation is Irrational and Harmful. *Quarterly Journal of Economics*, 129 (4). 1861–1898.
- [16] Eyster, E., Rabin, M., Weizsäcker, G., 2015 An Experiment on Social Mislearning. SSRN Working Paper
- [17] Fischbacher, U. , 2007. Z-Tree: Zurich Toolbox for Ready-Made Economic Experiments. *Experimental Economics*, 10(2), 171-178.
- [18] Goeree, J., Palfrey, T., Rogers, B., McKelvey, R., 2007. Self-Correcting Information Cascades. *Review of Economic Studies*, 74(3), 733-762.
- [19] Goeree, J., Holt, C., Palfrey, T., 2016, Quantal Response Equilibrium. *Princeton University Press*
- [20] Golub, B., Jackson, M., 2010. Naïve Learning in Social Networks and the Wisdom of Crowds. *American Economic Journal: Microeconomics*, 2(1), 112-149.
- [21] Guarino, A., Jehiel, P., 2013. Social Learning with Coarse Inference. *American Economic Journal: Microeconomics*, 5(1), 147-174.
- [22] Hahn, J., 1995. Bootstrapping Quantile Regression Estimators. *Econometric Theory*, 11(1), 105–121.
- [23] Jehiel, P., 2005. Analogy-Based Expectation Equilibrium. *Journal of Economic Theory*, 123(2), 81-104.
- [24] Kübler, D., Weizsäcker, G., 2004. Limited Depth of Reasoning and Failure of Cascade Formation in the Laboratory. *Review of Economic Studies*, 71, 425-441.
- [25] Kyle, A., Wang, A., 1997. Speculation Duopoly with Agreement to Disagree: Can Overconfidence Survive the Market Test? *Journal of Finance*, 52(5), 2073-2090.
- [26] Lee, I., 1993. On the Convergence of Informational Cascades. *Journal of Economic Theory*, 61(2), 395-411.
- [27] McKelvey, R., Palfrey, T., 1995. Quantal Response Equilibria for Normal Form Games. *Games and Economic Behavior*, 10(1), 6-38.
- [28] Nöth, M., Weber, M., 2003. Information Aggregation with Random Ordering: Cascades and Overconfidence. *The Economic Journal*, 113(84), 166-189.
- [29] Nyarko, Y., Schotter, A., 2002. An Experimental Study of Belief Learning Using Elicited Beliefs. *Econometrica*, 70(3), 971–1005.
- [30] Odean, T., 1998. Volume, Volatility, Price, and Profit When All Traders Are Above Average. *Journal of Finance*, 53(6), 1887-1934.

- [31] Schotter, A. 2003. Decision Making with Naive Advice. *American Economic Review PP*, 93(2), 196-201.
- [32] Schotter, A., Trevino, I., 2014. Belief Elicitation in the Lab. *Annual Review of Economics*, 6, 103-128.
- [33] Yildiz M., 2003. Bargaining without a Common Prior: an Immediate Agreement Theorem. *Econometrica*, 71, 793–811
- [34] Yildiz, M., 2011. Bargaining with Optimism. *Annual Review of Economics*, 3, 451-478.