

Social Learning with Coarse Inference

Antonio Guarino and Philippe Jehiel*

3 July 2009

Abstract

We study social learning by boundedly rational agents. Agents take a decision in sequence, after observing their predecessors and a private signal. They are unable to understand their predecessors' decisions in their finest details: they only understand the relation between the aggregate distribution of actions and the state of nature. We show that, in a continuous action space, compared to the rational case, agents put more weight on early signals. Despite this behavioral bias, beliefs converge to the truth. In a discrete action space, instead, convergence to the truth does not occur even if agents receive signals of unbounded precisions.

1 Introduction

In many economic and social situations, people learn from observing the decisions of others. While they learn from others, they can even decide to imitate what others do, and follow the crowd. Indeed, a central message of the social learning literature (Banerjee, 1992; Bikhchandani *et al.*, 1992) is that perfectly rational agents can decide to neglect their private information (on which action is the most profitable) and simply herd on the decisions of previous decision makers. The rationale for this result is simple. Suppose agents must choose in a set of discrete actions. The actions taken by other agents may be based on the private information that these agents had.

*Guarino: Department of Economics and ELSE, University College London (e-mail: a.guarino@ucl.ac.uk); Jehiel: Paris School of Economics and Department of Economics and ELSE, University College London (e-mail: jehiel@enpc.fr). We are grateful to Martin Cripps for useful comments and discussions, and, in particular, for his help on the proof of convergence. We thank Nathan Larson for helpful comments. We also thank participants in the Workshop on Information Externalities, Social Learning and Financial Markets in Cambridge and in the PET 2009 conference in Galway for useful suggestions. We thank Minyu Chen for excellent research assistance. We gratefully acknowledge the financial support of the ERC and of the ESRC. We are responsible for any errors.

When a sufficient number of agents have made the same decision and consensus has arisen, the public information revealed by these actions overwhelms the next agent’s private information and this agent will find it optimal to conform to the majority.

While the social learning literature offers many insightful results, it also reaches some unsettling conclusions. The argument we have illustrated above goes through in the full rationality paradigm when agents’ action space is discrete and the private signals agents receive have bounded precision. With a continuous action space, the process of social learning is, instead, efficient when agents are fully rational (see Lee, 1993). The continuous action space allows the agents to fine tune their actions to their expectations, thereby allowing the following agents to infer perfectly the signals from the observation of the previous actions. As a result, all private information is perfectly aggregated. In this setting, history does not matter, in two senses: the action of the immediate predecessor already contains all the public information an agent needs to make the optimal decision; the actions of the early agents in the sequence do not have any long lasting effect on the decisions of the following decision makers.

The process of social learning is eventually efficient even in the discrete action space case, as long as signals can have unbounded precision and agents are fully rational (see Smith and Sørensen, 2000). In this case, even if a “herd” of one million people occurs, the decision of the next agent with a very precise signal to go against the herd overturns the weight of the long sequence of predecessors, thus allowing the followers to take advantage of his precise information. This is what Smith and Sørensen (2000) refer to as the overturning principle.

Some of the conclusions reached in the full rationality paradigm sound unintuitive. For example, it does not sound fully convincing that human subjects would make the right inference after seeing one breaking a long herd in Smith and Sørensen’s model, nor does it sound fully convincing that human subjects would perfectly infer the sequence of past signals just by observing past actions in Lee’s continuous action space model. There are several routes to approach this: either modify some aspects of the game while maintaining the rationality assumptions,¹ or stick to the original game and try to propose alternative approaches (say with bounded rationality) to model the interaction.

We follow the second route with the view that no matter what the most realistic social learning model is, it is likely that the type of inferences required in the full rationality paradigm goes beyond what real subjects can reasonably be expected to do.²

¹Smith and Sørensen (2008), for instance, to overcome the overturning principle, dispense with the assumption of perfect observability of the history of actions and propose a model in which agents only observe unordered samples from past history.

²For example, if along with Smith and Sørensen (2008), we assume the order of moves is not observed, the rational

We note that the second route can itself be decomposed into several alternative approaches. For example, one can propose heuristic views on how subjects make their inferences in each specific social learning model, or one can view the subjects as making their inferences based on some partial yet correct understanding of the problem. It is the latter approach that we pursue here. Specifically, we develop an equilibrium approach in which agents make their inferences from what they observe based only on the knowledge of how the state of the world affects the distribution of actions. In simple words, this means that each agent understands the frequency with which each action is taken in a given state of the world, but does not understand how the frequency of the action depends on the specific history of decisions and on the private signals that agents receive. It is an equilibrium approach because the understanding of the agents even though partial is assumed to be correct.

We apply our approach both to the continuous and to discrete action space set ups. Key findings are that history does matter even in a continuous action space, and that the overturning principle does not hold when signals are of unbounded precision.

An important motivation for the type of partial understanding that we are assuming is that looking at previous social learning experiences it is probably easier to remember (or have access to data on) how actions are distributed as a function of the state of the world (which is often observable *ex post*) than to know how actions depend on the private history of decisions (which is often not even accessible).³ In other words, from a (repeated game) learning perspective, it seems plausible that agents can learn the relation between the state of the world and the frequency of actions, whereas they may find it difficult (or impossible, for lack of observation) to learn the relation between the actions and the private information available to agents. Besides, it may be argued that in a number of real life social learning interactions, subjects do not know the exact structure of information and preferences that applies to others, thereby making other forms of inferences either based on heuristic or introspective reasoning less reliable.

A simple, informal example will serve to clarify the concept. Imagine we run an experiment on the standard model of Bikhchandani *et al.* (1992). A group of subjects have to choose in sequence either action “a” or action “b.” Before making their decision, they receive a signal on whether the true state of the world is A or B. Then they choose in sequence, after observing the predecessors’ actions. The experiment is repeated many times. After each repetition, each agent is informed of the true state of the world and of his payoff (which depends only on the state of the world and on his chosen action). Suppose now at the end of the experiment, we consider all the repetitions in which

inferences in such a model are even more complex than in the standard model.

³Such a motivation is related to the theme of robust mechanism design (Bergmann and Morris, 2005) which explicitly acknowledges that beliefs of agents are not easily accessible.

the state of the world was A . We count the total number of “a” actions and of “b” actions (in all periods) and find out that action “a” was chosen 67% of the time, while action “b” was chosen 33% of the time. When the state of the world was B , instead, action “a” was chosen 35% of the time and action “b” 65%. Given these empirical frequencies, how would you play in this experiment if you had the possibility of participating in it? Our equilibrium concept assumes that agents expect that the probabilities of actions “a” and “b” being chosen are equal to these frequencies, independently of the time when the action is taken (and the specific sequence that the agent has observed until then). Moreover, agents best respond to these “frequencies,” by which we mean that they choose the best action assuming other’s actions in state $\omega = A$ or B are distributed according to these aggregate frequencies no matter what was observed. These actions in turn generate new data and assuming long run convergence of the overall process, the aggregate empirical frequencies provided to subjects should be correct, thereby motivating the equilibrium approach pursued in this paper.

Formally, our analysis is based on the concept of Analogy Based Expectation Equilibrium, developed by Jehiel (2005). We particularize the analogy-based expectation equilibrium by assuming that there are two analogy classes, one for each state of the world. This partition is referred to as the payoff-relevant analogy partition, given that each agent’s final payoff depends on the state of the world and his action and nothing else. In an Analogy Based Expectation Equilibrium, agents expect others to play according to the same mixed distribution in each state of the economy (irrespective of the history observed by the agent). Upon observing actions, agents update their beliefs (about the state of the economy) assuming other agents behave according to such state-dependent mixed distributions. Moreover, the state-dependent mixed distributions considered by the agents are assumed to coincide with the aggregate distribution of play in the corresponding state of the economy, a consistency requirement which is viewed as the outcome of a learning process. We will refer to this model of bounded rationality as the payoff-relevant reasoning model.

We start our analysis with the case of a continuous action space in which agents seek to choose an action as close as possible to what they think the state is. In the standard (Perfect Bayesian Equilibrium) analysis, the process of social learning is efficient (Lee, 1993): each action reveals the signal realization with the result that agent at time t acts as if he had observed all previous $t - 1$ signals (plus, obviously, his own signal). We show that, instead, in our equilibrium, a behavioral bias arises in that agents give more weight to early signals than to later ones: in particular, the agent moving at time t chooses an action as if signal 1 had weight $t - 1$, signal 2 had weight $t - 2$, etc.⁴ Despite this bias, however, over time the beliefs converge almost surely to the truth and, as

⁴Signal 1 has weight $t - 1$ in the sense that the choice of action is the same as the one that would have been made

a result, actions converge to the correct one. Furthermore, convergence occurs at an exponential rate, exactly as in the standard analysis. In other words, even though the coarse inference implies that early actions have a disproportionate effect on subsequent decisions, our boundedly rational agents learn in the long run, as it occurs in the model with fully rational agents. We show that these basic insights are robust to various extensions of the model with continuous action space (e.g., the case in which only some predecessors' actions are observed, and that in which the order of play is unknown).

We then move to study the case of a discrete action space in which the signals can be of unbounded precision. In contrast to the insights found with full rationality, we show that, under the payoff-relevant reasoning model, the probability that the agents settle on the correct action and beliefs converge to the truth is bounded away from one. Such asymptotic inefficiencies occur because in our model when an agent breaks a herd, it is unlikely that subsequent agents will follow this deviator for they miss the inference that the deviator must have received a very precise signal.

A summary of our asymptotic results is as follows. While in the full rationality case, either the continuous action space or the unbounded precision of signals guarantee the convergence of beliefs to the truth, in the payoff-relevant reasoning model we find that convergence to the truth obtains in the continuous action space case but not in the unbounded precision case.

Almost all the literature on social learning assumes full rationality. Nevertheless, our approach to bounded rationality is obviously not the only one that can be applied to study social learning. If agents were fully cursed, as modelled in Eyster and Rabin (2005), they would base their decision solely on their own signals, as others' actions would be (wrongly) thought to be uninformative about the state of the world. The result would be that the beliefs would not converge to the truth and decisions would not settle on a particular action. More generally, if agents were partially cursed (as defined in Eyster and Rabin, 2005), early signals could not have more effects than later signals on subsequent actions. Such results should be contrasted with our finding in the continuous action space model that early signals have significantly more impact than later signals on current decisions when agents rely on the payoff relevant model of reasoning.⁵

Elaborating on their previous work, in a recent paper, Eyster and Rabin (2008) consider a

by a rational agent observing $t - 1$ independent draws of that signal realization. Similarly for the other signals.

⁵Note that the fully cursed equilibrium can be viewed as an analogy-based expectation equilibrium in which agents use the private information analogy partition (see Jehiel and Koessler, 2008). The contrast between the results reveals how much extra flexibility is allowed by varying the analogy partitions. We believe the payoff relevant analogy partition is more suited to the analysis of social learning given also that it allows for some form of non-trivial inference unlike the private information analogy partition.

framework in which agents wrongly believe that other players are cursed, whereas they are not, which combines ideas from the cursed equilibrium and the subjective prior paradigms. It turns out that in a framework with continuous action and signal space, this approach coincides with a heuristic approach in which subjects would interpret past actions as if they were signals of various precisions.⁶ A striking result they obtain is that early signals are overwhelmingly influential, leading to asymptotic inefficiencies. In our model with continuous action space, instead, despite the behavioral bias that assigns a higher weight to early signals, eventually beliefs converge to the true state of nature and actions settle on the correct one.

Other models of social learning with bounded rationality include Bala and Goyal (1998), De Marzo *et al.* (2003), Acemoglu *et al.* (2009) and Ellison and Fudenberg (1993). In Bala and Goyal (1998), agents in a network choose after observing their neighbors' actions and payoffs. There is private information in their model, but agents are assumed to ignore it to some extent. By assumption, each agent learns from his neighbor's actions (experiments) but does not ask what information might have led the neighbor to choose those actions. De Marzo *et al.* (2003) and Acemoglu *et al.* (2009) also focus on networks, but learning in these models is non-Bayesian. In Ellison and Fudenberg (1993) agents consider the experiences of their neighbors and learn using rules of thumb. In some cases, even naive rules can lead to efficient decisions, but adjustment to an innovation can be slow.

The paper is organized as follows. In Section 2 we present the economy with a continuous action space and we define our solution concept. In Section 3 we illustrate the main equilibrium analysis. In Section 4 we present the asymptotic properties of the economy. In Section 5 we elaborate on existence and uniqueness issues. In Section 6, we illustrate extensions of the model. In Section 7 we study the case of a discrete action space. Section 8 concludes. The Appendix contains the proofs not presented in the text.

2 The Model

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

Agent t takes an action a_t in the action space $[0, 1]$. The agent's payoff, depends on his choice and on the true state of the world $\omega \in \{0, 1\}$. We assume that the payoff is quadratic and, in

⁶With a continuous signal space, every action can be interpreted that way.

particular, equal to $-(\omega - a_t)^2$. Each agent t receives a private signal $s_t \in \{0, 1\}$ correlated with the true state ω . In particular, he receives a symmetric binary signal distributed as follows:

$$\Pr(s_t = 1 \mid \omega = 1) = \Pr(s_t = 0 \mid \omega = 0) = q_t \in \left(\frac{1}{2}, 1\right),$$

with $q_t \neq q_{t'}$ for any $t \neq t'$. This means that, conditional on the state of the world, the signals are independent but not identically distributed over time. We will refer to the ratio $\frac{\Pr(s_i \mid \omega=1)}{\Pr(s_i \mid \omega=0)}$ as the “multiplier” and denote it by $m(s_i)$.⁷ Note that the multiplier is equal to $\left(\frac{q_i}{1-q_i}\right)^{2s_i-1}$. We make the following (genericity) assumption:⁸

$$\text{For } (n_1, n_2, \dots, n_T) \in \mathbb{Z}^T, \text{ if } \prod_{i=1}^T \left(\frac{q_i}{1-q_i}\right)^{n_i} = 1, \text{ then } n_i = 0 \text{ for all } i. \quad (\text{A1})$$

Agents know the precision of their own signal, as well as the precision of the others’ signals.⁹

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, $h_t = \{a_1, a_2, \dots, a_{t-1}\}$.¹⁰ We denote the set of such histories by H_t . Agent t ’s information set is then represented by the couple (h_t, s_t) . Given the information (h_t, s_t) , the agent will choose a_t to maximize the expected payoff $E^{SUBJ}[-(\omega - a_t)^2 \mid h_t, s_t]$. Therefore, the optimal action will be $a_t^* = E^{SUBJ}[\omega \mid h_t, s_t]$. We use the superscript “SUBJ” in the expectation to emphasize that each agent forms a subjective expectation. We discuss the way in which this is done in the next subsection. For the time being, observe that the above considerations imply that we can restrict attention to pure strategies (given that $E^{SUBJ}[\omega \mid h_t, s_t]$ reduces to the choice of a single action).

2.1 Equilibrium Concept

We depart from the standard approach in one essential aspect, that is, in modelling how agents form expectations about their opponents’ strategies. We adopt the Analogy Based Expectation Equilibrium concept first introduced by Jehiel (2005) using, more specifically, the payoff-relevant analogy partition (see Jehiel and Koessler, 2008). We assume that agents are unable to understand the other agents’ strategies in their finest details, thereby making it impossible for them to assess how the choice of action depends on the public history and the private signal at every date t .

⁷To simplify the exposition, we use the same symbol to indicate a random variable and its realization.

⁸Such an assumption is met for almost every (q_1, \dots, q_T) , whatever the Lebesgue measure on $(\frac{1}{2}, 1)^T$.

⁹We make this assumption only to present our model in its simplest version, the closest to the standard social learning model. We will see later that agents need not be aware of the precision of others’ signals for our results to hold.

¹⁰Note that $h_1 = \emptyset$.

Instead, agents are assumed to make their inferences from past play based only on the knowledge of how the state of the world ω affects the distribution of actions. This is referred to as the payoff-relevant analogy partition given that final payoffs depend on the state of the world ω and the actions but not on the signals directly.

Formally, let us denote a strategy profile by σ , that is, $\sigma = (\sigma_1, \sigma_2, \dots, \sigma_T)$, where an agent's strategy σ_t maps (h_t, s_t) into an action, that is, $\sigma_t : H_t \times \{0, 1\} \rightarrow [0, 1]$. With a slight abuse of notation, $\sigma_t(a|h_t, s_t)$ denotes the probability that agent t picks action a when the history is h_t and the signal is s_t . Since, as already mentioned, equilibria in pure strategies will be considered, $\sigma_t(a|h_t, s_t)$ will either be equal to 0 or to 1. Given a particular strategy profile σ in pure strategies, we denote by $\mu^\sigma(h_t, s_t|\omega)$ the probability that history h_t is realized and s_t is the signal at t when ω is the state of the world. Observe that because s_t takes values in $\{0, 1\}$ and the strategies σ_t are pure, $\mu^\sigma(h_t, s_t|\omega) > 0$ only for finitely many (h_t, s_t) .

The aggregate distribution of actions as a function of the state of the world ω is given by:

$$\bar{\sigma}(a|\omega) = \frac{\sum_{t=1}^T \sum_{h_t, s_t} \sigma_t(a|h_t, s_t) \mu^\sigma(h_t, s_t|\omega)}{\sum_{t=1}^T \sum_{h_t, s_t} \mu^\sigma(h_t, s_t|\omega)}.$$

This is a mixed distribution that assigns positive weight to those actions that are played after some (h_t, s_t) , and the weight assigned to $\sigma_t(a|h_t, s_t)$ is proportional to $\mu^\sigma(h_t, s_t|\omega)$, as required from a learning story in which agents would have learned from past interactions the distribution of actions as a function of the state ω .¹¹

Given that, for each t , $\sum_{h_t, s_t} \mu^\sigma(h_t, s_t|\omega) = 1$, this expression simplifies into

$$\bar{\sigma}(a|\omega) = \frac{\sum_{t=1}^T \sum_{h_t, s_t} \sigma_t(a|h_t, s_t) \mu^\sigma(h_t, s_t|\omega)}{T}. \quad (1)$$

In an Analogy Based Expectation Equilibrium, every agent t assumes that other agents choose action a with probability $\bar{\sigma}(a|\omega)$ when the state is ω , and that these behaviors are randomized independently across periods. Furthermore, agent t believes that the signal s_t is independent of the actions of his predecessors. He chooses a best-response accordingly:¹²

Definition 1 *An Analogy Based Expectation Equilibrium with payoff-relevant analogy partitions (ABEE) is a strategy profile σ such that for every t , σ_t is a best response to the conjecture that*

¹¹Note that the aggregate distribution is obtained considering all periods, not only the preceding periods, since agents bundle *all* decision nodes in the two analogy classes.

¹²Compared to the framework developed in Jehiel and Koessler (2008), there are a few differences. First, we consider a multi-stage, multi-player setup, whereas Jehiel and Koessler consider two-person, simultaneous move games. Second, the analogy partitions as defined above include the decision nodes of all players and not just those of the *other* players.

other agents follow the strategy $\bar{\sigma}$ as defined in (1), and that, conditional on ω , agent t 's signal s_t is independent of his predecessors' actions.

At an interpretative level, we see the ABEE as representing a steady state of a learning process and not as a result of introspective reasoning. The consistency required by the equilibrium concept should thus be viewed as the outcome of a dynamic process in which agents would eventually know how actions are distributed as a function of the state the economy. Such a learning process only requires that agents be informed of the state of the world as well as of the actions chosen in previous plays (together with the structure of their own payoffs and the precision of their own signal). Agents need not have a prior knowledge about the payoffs or information structure of other players, nor of their ways of reasoning.

3 Equilibrium Analysis

At a first glance, finding an ABEE seems rather complicated. Definition 1 involves a fixed point argument (σ_t is a best-response to $\bar{\sigma}$ and $\bar{\sigma}$ is derived from the various σ_t through (1)), which could make the analysis difficult. One could also suspect that there are multiple ABEE. As we will see (in Section 5), generically there is only one ABEE and it corresponds to the strategy profile constructed below.

Let us start the construction of an ABEE. Letting $\alpha(s_1, s_2, \dots, s_k)$ denote the equilibrium action taken after the sequence of signal realizations (s_1, s_2, \dots, s_k) , we conjecture that the agents' strategies are such that, in equilibrium, for any two different sequences of signal realizations, $\{s_1, s_2, \dots, s_k\} \neq \{s'_1, s'_2, \dots, s'_l\}$, an agent chooses two distinct actions in $[0, 1]$, that is, $\alpha(s_1, s_2, \dots, s_k) \neq \alpha(s'_1, s'_2, \dots, s'_l)$.¹³ Note that this requirement comprises both the case in which $k = l$ and $s_t \neq s'_t$ for at least one $t = 1, 2, \dots, k$, and the case in which $k \neq l$. Given this conjecture, we construct an equilibrium and, then, we verify that it satisfies the conjecture.

Specifically, consider a sequence of signals (s_1, s_2, \dots, s_T) received by agents $1, 2, \dots, T$. At time 1, agent 1 has only his signal s_1 , and he forms the likelihood ratio

$$\frac{\Pr(\omega = 1|s_1)}{\Pr(\omega = 0|s_1)} = m(s_1) = \left(\frac{q_1}{1 - q_1} \right)^{2s_1 - 1}.$$

Therefore, agent 1 will choose the action $\alpha(s_1)$ such that $\frac{\alpha(s_1)}{1 - \alpha(s_1)} = m(s_1)$. That is, he chooses action $a_1 = q_1$ if $s_1 = 1$, and $a_1 = 1 - q_1$ if $s_1 = 0$.

¹³Note that $\alpha(s_1, s_2, \dots, s_k)$ is well defined, given that the equilibrium involves pure strategies.

Consider now agent 2. After observing action $a_1 = \alpha(s_1)$ and the private signal s_2 , agent 2 forms the following (subjective) likelihood ratio:

$$\frac{\Pr(\omega = 1|s_2, a_1)}{\Pr(\omega = 0|s_2, a_1)} = \frac{\bar{\sigma}(a_1|\omega = 1)\Pr(s_2|\omega = 1)}{\bar{\sigma}(a_1|\omega = 0)\Pr(s_2|\omega = 0)}.$$

Recall that $\bar{\sigma}(a|\omega) = \frac{\sum_{t=1}^T \sum_{h_t, s_t} \sigma_t(a|h_t, s_t) \mu^\sigma(h_t, s_t|\omega)}{T}$. Given our conjecture that $\alpha(s_1)$ is different from any other $\alpha(s'_1, \dots, s'_t)$, the probability of action $a = q_1$ ($a = 1 - q_1$) is equal to one for the sequence $\{s_1 = 1\}$ ($\{s_1 = 0\}$) and zero for any other sequence of signal realizations. This implies that

$$\begin{aligned} \bar{\sigma}(a = q_1|\omega) &= \frac{\mu^\sigma(h_1, s_1 = 1|\omega)}{T} = \frac{\Pr(s_1 = 1|\omega)}{T}, \text{ and} \\ \bar{\sigma}(a = 1 - q_1|\omega) &= \frac{\mu^\sigma(h_1, s_1 = 0|\omega)}{T} = \frac{\Pr(s_1 = 0|\omega)}{T}. \end{aligned}$$

Therefore, we obtain that

$$\frac{\bar{\sigma}(a_1|\omega = 1)}{\bar{\sigma}(a_1|\omega = 0)} = \frac{\Pr(s_1|\omega = 1)}{\Pr(s_1|\omega = 0)} = m(s_1), \quad (2)$$

and, thus, agent 2 chooses action $\alpha(s_1, s_2)$ such that

$$\frac{\alpha(s_1, s_2)}{1 - \alpha(s_1, s_2)} = m(s_1)m(s_2).$$

The actions chosen by agents 1 and 2 are the same as in the standard case (as we will see in the next subsection). This ceases to be the case for the subsequent actions chosen in periods $t \geq 3$. Consider agent 3: after receiving the signal s_3 and observing the history $a_1 = \alpha(s_1)$ and $a_2 = \alpha(s_1, s_2)$, the agent forms the (subjective) likelihood ratio

$$\frac{\Pr(\omega = 1|s_3, a_1, a_2)}{\Pr(\omega = 0|s_3, a_1, a_2)} = \frac{\bar{\sigma}(a_1|\omega = 1)\bar{\sigma}(a_2|\omega = 1)}{\bar{\sigma}(a_1|\omega = 0)\bar{\sigma}(a_2|\omega = 0)}m(s_3).$$

The probability of observing action $a_2 = \alpha(s_1, s_2)$ in state ω is just the probability that in state ω the signals in dates 1 and 2 are s_1 and s_2 , as for any other date and/or signal realization the action would be different (given our conjecture). It follows that

$$\bar{\sigma}(a_2|\omega) = \frac{\mu^\sigma(h_2 = \{a_1\}, s_1 = 0|\omega)}{T} = \frac{\Pr(s_1, s_2|\omega)}{T},$$

and thus

$$\frac{\bar{\sigma}(a_2|\omega = 1)}{\bar{\sigma}(a_2|\omega = 0)} = \frac{\Pr(s_1, s_2|\omega = 1)}{\Pr(s_1, s_2|\omega = 0)} = m(s_1)m(s_2).$$

Combining this with (2), the result is that agent 3 chooses action $\alpha(s_1, s_2, s_3)$ such that

$$\frac{\alpha(s_1, s_2, s_3)}{1 - \alpha(s_1, s_2, s_3)} = m(s_1)^2 m(s_2) m(s_3).$$

Note that signal s_1 is “counted” twice in this expression, first for the inference from action $\alpha(s_1)$ and then for the inference from action $\alpha(s_1, s_2)$.¹⁴ The inference from $\alpha(s_1)$ is standard. The double counting is the direct consequence of the agent bundling states into analogy classes and only considering the aggregate behavior in the two classes.

A similar reasoning applies to any time t . After observing $a_1 = \alpha(s_1), \dots, a_{t-1} = \alpha(s_1, \dots, s_{t-1})$ and signal s_t , agent t forms a (subjective) likelihood ratio

$$\begin{aligned} \frac{\Pr(\omega = 1 | s_t, a_1, a_2, \dots, a_{t-1})}{\Pr(\omega = 0 | s_t, a_1, a_2, \dots, a_{t-1})} &= \frac{\prod_{i=1}^{t-1} \bar{\sigma}(a_i | \omega = 1)}{\prod_{i=1}^{t-1} \bar{\sigma}(a_i | \omega = 0)} m(s_t) = \\ &= \frac{\Pr(s_1 | \omega = 1) \Pr(s_1, s_2 | \omega = 1) \cdots \Pr(s_1, s_2, \dots, s_{t-1} | \omega = 1)}{\Pr(s_1 | \omega = 0) \Pr(s_1, s_2 | \omega = 0) \cdots \Pr(s_1, s_2, \dots, s_{t-1} | \omega = 0)} m(s_t), \end{aligned}$$

where the second equality follows from our conjecture that all α 's are different.

In this updating, signal s_1 is counted $t - 1$ times, signal s_2 is counted $t - 2$ times and so on. Thus, agent t chooses action $\alpha(s_1, s_2, \dots, s_t)$ such that¹⁵

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

To conclude, note that, given our genericity assumption (A1), our conjecture that all α 's are different is satisfied, which implies that we have just found an ABEE.¹⁶

We summarize this result in the following proposition:

Proposition 1 *There exists an ABEE in which, after a sequence of signals s_1, s_2, \dots, s_t , agent t chooses action $a_t^* = \alpha(s_1, s_2, \dots, s_t)$ such that $\frac{a_t^*}{1 - a_t^*} = \frac{\alpha(s_1, s_2, \dots, s_t)}{1 - \alpha(s_1, s_2, \dots, s_t)} = \prod_{i=1}^{t-1} \left(\frac{q_i}{1 - q_i} \right)^{(2s_i - 1)(t - i)} \left(\frac{q_t}{1 - q_t} \right)^{2s_t - 1}$.*

In the above proposition, the strategy of agent t is not explicitly constructed as a function of the history and of agent t 's private signal. Yet, such a strategy is easily determined for the histories $\{a_1, a_2, \dots, a_{t-1}\}$ such that $a_1 = \alpha(s_1), a_2 = \alpha(s_1, s_2), \dots, a_{t-1} = \alpha(s_1, s_2, \dots, s_{t-1})$ and by the signal s_t ,

¹⁴The signal is “counted” twice in the sense that the choice of action is the same as the one that would have been made by a rational agent observing two independent draws of that signal realization.

¹⁵Note that, throughout the paper, we use the convention that $\prod_{i=1}^0 x^i$ is equal to 1.

¹⁶From this construction it should be clear that agents need not know the precisions of other agents' signals (neither they need to know their realizations, of course). Indeed, the equilibrium is constructed considering the aggregate distributions given the states of the world, which does not require knowledge of others' precisions.

by identifying the strategy $\sigma_t(a|a_1, a_2, \dots, a_{t-1}, s_t)$ with the choice of the action $a = \alpha(s_1, s_2, \dots, s_t)$ with probability 1. For other histories, the strategy is not specified, but this is irrelevant for the analysis of the equilibrium path.¹⁷

3.1 A Benchmark: The Perfect Bayesian Equilibrium

Before we continue our analysis, it is worth contrasting the result in Proposition 1 with what happens in the case in which agents are fully rational. The analog of Proposition 1 for the case of fully rational agents writes (see Lee, 1993, for a detailed analysis):

Proposition 2 *There exists a unique Perfect Bayesian Equilibrium (PBE). In the PBE, after a sequence of signals s_1, s_2, \dots, s_t , agent t chooses action $a^{PBE} = \alpha^{PBE}(s_1, s_2, \dots, s_t)$ such that*

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

In the PBE, agents can perfectly infer the signals observed by their predecessors from their choices of actions. As a result agents pick the action that corresponds to the expected value of the state of the world conditional on the signals received by themselves and all their predecessors.¹⁸

The comparison between Propositions 1 and 2 makes it very easy to appreciate the difference between the two approaches. Essentially, while in the ABEE earlier signals receive a higher weight, in a PBE they all have the same weight. To see what difference this can make, let us present a simple example.

3.2 An Example

For illustration, let us consider an example with $T = 5$. The five agents receive a private signal of precision $q_1 = 0.95$, $q_2 = 0.9$, $q_3 = 0.85$, $q_4 = 0.8$, and $q_5 = 0.75$. Moreover, let us assume that the first agent receives a signal equal to 0 and the remaining four agents a signal equal to 1.

The first agent does not have a problem of learning from others. He receives the signal $s_1 = 0$ and computes his likelihood ratio as

$$\frac{\Pr(\omega = 1 | s_1 = 0)}{\Pr(\omega = 0 | s_1 = 0)} = \frac{\Pr(s_1 = 0 | \omega = 1)}{\Pr(s_1 = 0 | \omega = 0)} = \frac{1 - 0.95}{0.95}.$$

¹⁷Obviously, that presented above is only our way (i.e., the modelers' way) of constructing the equilibrium strategies. Boundedly rational agents do not go through our steps of reasoning to choose their actions. On the contrary, our interpretation is that they learn $\bar{\sigma}$ (and so act according to our formula in Proposition 1) simply through repeated play of the game, as discussed above.

¹⁸Also in the case of the PBE agents do not need to know the other agents' signal precisions. The reason is quite different from the case of the ABEE, though. Here the knowledge is not required since rational agents can infer the precision of the signals (as well as their realizations) from the observation of the sequence of actions.

As a result, he chooses $a_1^* = E[\omega | s_1 = 0] = 0.05$.

The second agent observes a_1^* and forms his likelihood ratio

$$\frac{\Pr(\omega = 1 | s_2 = 1, a_1 = 0.05)}{\Pr(\omega = 0 | s_2 = 1, a_1 = 0.05)} = \frac{\bar{\sigma}(0.05 | \omega = 1)}{\bar{\sigma}(0.05 | \omega = 0)} m(s_2 = 1) \\ = \frac{(0.05)(0.9)}{(1 - 0.05)(1 - 0.9)} = 0.47368,$$

which implies that he chooses $a_2^* = E^{SUBJ}[\omega | h_1 = \{0.05\}, s_2 = 1] = 0.32143$.

Now, let us consider agent 3. After observing $(h_3, s_3) = (\{0.05, 0.32143\}, 1)$, his a likelihood ratio is

$$\frac{\Pr(\omega = 1 | a_1^* = 0.05, a_2^* = 0.32143, s_3 = 1)}{\Pr(\omega = 0 | a_1^* = 0.05, a_2^* = 0.32143, s_3 = 1)} = \\ \frac{\bar{\sigma}(0.05 | \omega = 1) \bar{\sigma}(0.32143 | \omega = 1)}{\bar{\sigma}(0.05 | \omega = 0) \bar{\sigma}(0.32143 | \omega = 0)} m(s_3 = 1) = \\ \frac{(0.05)^2 (0.9) (0.85)}{(1 - 0.05)^2 (1 - 0.9) (1 - 0.85)} = 0.14127.$$

This implies that his action is $a_3^* = 0.12378$. Similarly, agent 4 observes $(h_4, s_4) = (\{0.05, 0.32143, 0.12378\}, 1)$, has a likelihood ratio of 0.26768, and his optimal action is $a_4^* = 0.21116$. Finally, agent 5, after observing $(h_5, s_5) = (\{0.05, 0.32143, 0.12378, 0.21116\}, 1)$, has a likelihood equal to 2.1555, and chooses $a_5^* = 0.68309$.

For comparison, note that in the PBE, while agents 1 and 2 would make exactly the same decisions as in our model ($a_1^{PBE} = a_1^* = 0.05$, $a_2^{PBE} = a_2^* = 0.32143$), agents 3, 4 and 5 would not, since they would update their beliefs differently. In particular, agent 3's likelihood ratios would be

$$\frac{(0.05)(0.9)(0.85)}{(1 - 0.05)(1 - 0.9)(1 - 0.85)} = 2.6842,$$

and, similarly, agent 4 and agent 5's likelihood ratios would be 10.737 and 32.211, which implies that they would choose $a_3^{PBE} = 0.72857$, $a_4^{PBE} = 0.91480$, and $a_5^{PBE} = 0.96989$.

The difference between the predictions of the two models is striking in this example. In the PBE, at time 5, after four good signals and one bad signal the optimal action is close to 1. In our ABEE, in contrast, due to the high weight that agent 5 gives to the first negative realization, it is lower than 0.7. Moreover, in the PBE, since the signal realizations following the first are all good, the optimal action is monotonically increasing: agents become more and more convinced that the true state of the world is $\omega = 1$. This contrasts with the ABEE, where agent 3 chooses an action

lower than agent 2. The result is due to agent 3 double counting the bad realization of the signal that led agent 1 to choose 0.05, while agent 2 updates his belief counting it only once.

Clearly, in this example, the difference between the two equilibria is particularly strong due to our choice of a history in which the first signal (which always carries more weight) is different from the others, and, furthermore, is the one with the highest precision. For other histories and other parameter values, the differences may be less striking, yet they never disappear.

4 Behavioral Bias and Long Run Convergence

The behavioral bias we have identified implies that if early in the history of play agents receive the incorrect signals, this will have a more severe effect on future actions in an ABEE than in the PBE, since these signals receive more weight. A natural question is whether this effect will persist over time, so that beliefs may converge to the wrong value, or whether, eventually, despite the behavioral bias, convergence of beliefs to the truth obtains. And if convergence obtains, obviously we are also interested in whether boundedly rational agents learn as fast as rational agents do.

The answer to these questions is that in our ABEE beliefs converge almost surely to the true state of the world and, eventually, actions settle on the correct one. Furthermore, convergence occurs exponentially fast in our ABEE as it does in the PBE. These results are reported in the next proposition, where, for tractability, we restrict attention to the limit case in which the precisions of signals tend to be all identical and equal to q . To introduce the proposition, we first denote the public belief at time t by $B_t := \Pr(\omega = 1|h_t)$, and the subjective public belief (i.e., the belief of a boundedly rational agent after observing the history, but before receiving his signal) by $B_t^{SUBJ} := \Pr^{SUBJ}(\omega = 1|h_t)$. In other words, B_t and B_t^{SUBJ} are the public belief in a PBE and an ABEE, respectively. Observe that these beliefs are independent of the total duration T , as can be inferred from Propositions 1 and 2. In the next proposition we analyze how B_t and B_t^{SUBJ} vary with t (accordingly we assume that $T \geq t$). To this end, it is convenient to define $X_t := \frac{\sum_{i=1}^t s_i}{t}$ and $Z_t := \frac{\sum_{i=1}^t u_i}{t(t+1)/2}$, where the random variables u_i are distributed as follows: $\Pr(u_i = i|\omega = 1) = \Pr(u_i = 0|\omega = 0) = q$, $\Pr(u_i = 0|\omega = 1) = \Pr(u_i = i|\omega = 0) = 1 - q$. The random variable X_t is the average number of signals 1 realized until time t (included). Since in the PBE the observation of each action is equivalent to the observation of each signal, X_t summarizes the public belief at time $t + 1$ in the PBE. In an ABEE, at time $t + 1$, the subjective public belief is obtained by counting the first signal t times, the second $t - 1$ times, etc. These are the numbers taken by the random variable u_i (for $i = t, t - 1, \dots$) when the first signal takes value $\frac{u_t}{t}$, the second signal takes value $\frac{u_{t-1}}{t-1}$, etc. Therefore, Z_t summarizes the subjective public belief at time $t + 1$ in our ABEE. We have the following proposition:

Proposition 3 *Consider the limit case in which the precisions of all signals tend to the same value $q \in (0.5, 1)$. In both the ABEE and the PBE the public belief converges almost surely to the true state of the world, that is, when $\omega = 1$, $B_t^{SUBJ} \xrightarrow{a.s.} 1$, and $B_t \xrightarrow{a.s.} 1$ (and, similarly, when $\omega = 0$, $B_t^{SUBJ} \xrightarrow{a.s.} 0$, and $B_t \xrightarrow{a.s.} 0$). Moreover, in both cases, convergence obtains exponentially fast, that is, $\Pr(|X_t - q| > \varepsilon) \leq 2e^{-2t\varepsilon^2}$ and $\Pr(|Z_t - q| > \varepsilon) \leq 2e^{-tK(\varepsilon)\varepsilon^2}$, where $K(\varepsilon)$ is the value of K that solves the equation $(q + \varepsilon)K = \log((1 - q) + qe^K)$.*

Despite the bias in our ABEE, the signals are taken into account in the choice of each action. Since the distribution of signals is markedly different in the two states, eventually the true state of the world is discovered almost surely. Given that early signals receive a higher weight, clearly, for histories in which early signal realizations happen not to be representative of the true state, convergence will be slowed down. If the early signal realizations are representative of the state, instead, convergence will be faster. Our proposition shows that, for almost all sequences of signal realizations, the bias does not affect the form of long run convergence, since it occurs exponentially in the ABEE as it does in the PBE. Our boundedly rational agents learn in the long run, as fully rational agents do. To gain further intuition on this result, it is useful to note the following. Consider the first n consecutive signals. We know from the analysis of the standard case that when n is large enough, the probability that the difference between q and the frequency of signal realizations 1 in state $\omega = 1$ (0 in state $\omega = 0$) is higher than ε is exponentially small in n . Now, in our ABEE, at time $t > n$ the first signal is counted $t - 1$ times and the n -th signal is counted $t - n$ times. When t grows large, the different weight between the first and the n -th signal becomes, however, negligible since $(t - n)/t$ approximates 1. In other words, since the over-counting determined by the bias takes a polynomial form, it vanishes in the limit and convergence takes an exponential form in the ABEE as it does in the PBE.

Figure 1 here

We have simulated the model for various parameter values. Figure 1 shows the average distance of the public belief from the true state of the world (i.e., $|\Pr(\omega|h_t) - \omega|$ and $|\Pr^{SUBJ}(\omega|h_t) - \omega|$) in the case in which all signals are drawn randomly from the interval $(0.69, 0.71)$. The average distance at any time t is taken over 100,000 replications.¹⁹ Essentially this figure is a graphical representation of our proposition. The two graphs in the figure are almost overlapping, indicating that the long run properties of our ABEE are not dissimilar from those of the PBE.

¹⁹We have used this same number of repetitions also for the other simulations presented below.

Figure 2 here

Figure 2 shows the average distance between belief and fundamental in the ABEE and the PBE, conditional on the first five signals being incorrect. In the PBE, starting at time 6, when signal realizations are randomly drawn (from the same interval as before, and, therefore, are correct approximately 70% of the repetitions), the belief starts approaching the true state of the world: the distance between belief and state of the world decreases quickly and monotonically. The graph for the ABEE looks rather different. First, after the first wrong signals, the agents become almost certain of the wrong state of the world (the distance is close to 1). After time 5, despite signals are now correct with frequency close to 70%, the difference between subjective public belief and the true state of the world remains close to one. This is because agents put more and more weight on early signals. It takes 25 periods before the impact of the first signals is offset by the later signals and the belief starts converging to the truth.

While Figures 1 and 2 were obtained for the case in which the q 's are all close to 0.7, we have run further simulations for the case in which q would be drawn according to a uniform distribution between 0.5 and 1, and we have obtained very similar graphs, thereby suggesting that the results on convergence summarized in our proposition extend beyond the limit case of identical precisions.

To conclude this section, let us observe that the inability of agents to understand the fine details of the inference problem from the predecessors' actions could lead, in principle, to different predictions. For instance, suppose agents essentially ignore the information content of past history. Then, we would observe no convergence of actions (even though a Bayesian external observer would, of course, learn the true state of the world). On the other hand, suppose agents put a lot of weight on previous actions. Then they may be prone to a sort of herding and, although actions would settle, there would be no information aggregation. We view our results as somehow in between these two extremes. In our model, coarse inference does determine a behavioral bias. This, however, does not preclude the aggregation of information, which actually occurs as in a world of rational agents.

This is in sharp contrast with what happens if subjects use the heuristic reasoning that upon observing action $a \in (\frac{1}{2}, 1)$ of one of his predecessors, the agent believes this corresponds to an independent signal $s = 1$ having precision $q(a)$ such that $\frac{q(a)}{1-q(a)} = a$. In this case (which corresponds to the case studied in Eyster and Rabin, 2008), the weight of the first signal is approximately equal to the sum of the weights of all other signals and, given the overwhelming weight of the early signals (which may be wrong), there must be asymptotic inefficiencies.

5 On Existence and Uniqueness of ABEE

In Section 3 we have constructed an ABEE under the genericity assumption (A1). In this section we want to understand, more generally, whether there are other ABEE and what happens if assumption (A1) is violated. The main message will be that, generically, the ABEE constructed in Section 3 is unique; moreover, when (A1) is violated, there is no guarantee that an ABEE exists.

We will proceed in the following way. We will first discuss, in Section 5.1, how to construct an ABEE generally, and, then, we will move to the issue of the uniqueness. In Section 5.2, we will discuss the issue of the non-existence along with the case of equal precisions of signals.

5.1 Construction of an ABEE and uniqueness

To start our construction of an ABEE, recall that $a_t^* = E^{SUBJ}[\omega|h_t, s_t]$, which means that for any (h_t, s_t) the best response reduces to the choice of a single action. Although agents do not observe others' private signals, each action a_t^* is taken after a sequence of private signals (s_1, s_2, \dots, s_t) . This means that also for any (s_1, s_2, \dots, s_t) the best response reduces to the choice of a single action. We denote this choice in a candidate ABEE by $\alpha(s_1, s_2, \dots, s_t)$.

Let us now define the set $\Psi = \{(s_1, \dots, s_t) | t = 1, 2, \dots, T\}$ of all $(2 + \dots + 2^{T-1} + 2^T)$ possible sequences of signal realizations. Given this set, let us define the partition $\mathcal{P}(\Psi)$ whose elements P are such that if $(s_1, s_2, \dots, s_t) \in P$ and $(s'_1, s'_2, \dots, s'_t) \in P$ then $\alpha(s_1, s_2, \dots, s_t) = \alpha(s'_1, s'_2, \dots, s'_t)$. In particular, we define $P(a) = \{(s_1, s_2, \dots, s_t) | \alpha(s_1, s_2, \dots, s_t) = a, t = 1, 2, \dots, T\}$ as the set of all sequences of signal realizations after which an agent chooses the same action a .

We now show that all ABEE can be constructed in the following way: 1) conjecture a specific partition $\mathcal{P}(\Psi)$; 2) compute the actions chosen after each sequence of signal realizations; 3) finally, verify that the conjecture is satisfied. In our construction in Section 3 we used the finest partition in which each $P(a) \neq \emptyset$ is a singleton. We now consider the general case in which the $P(a) \neq \emptyset$ are not necessarily singletons.

Take any partition $\mathcal{P}(\Psi)$. After observing action a , an agent forms the likelihood ratio

$$\frac{\Pr(\omega = 1|a)}{\Pr(\omega = 0|a)} = \frac{\bar{\sigma}(a|\omega = 1)}{\bar{\sigma}(a|\omega = 0)},$$

where the probability of observing action a in state ω should satisfy

$$\begin{aligned} \bar{\sigma}(a|\omega) &= \frac{\sum_{t=1}^T \sum_{h_t, s_t} \sigma_t(a|h_t, s_t) \mu^\sigma(h_t, s_t|\omega)}{T} = \\ &= \frac{\sum_{t=1}^T \sum_{(s_1, s_2, \dots, s_t) \in P(a)} \Pr(s_1, s_2, \dots, s_t|\omega)}{T}, \end{aligned}$$

since $\sigma_t(a|h_t, s_t)$ is equal to 1 if $(s_1, s_2, \dots, s_t) \in P(a)$ and to 0 otherwise. Thus,

$$\begin{aligned} \frac{\bar{\sigma}(a|\omega = 1)}{\bar{\sigma}(a|\omega = 0)} &= \\ \frac{\sum_{t=1}^T \sum_{(s_1, s_2, \dots, s_t) \in P(a)} \Pr(s_1, s_2, \dots, s_t | \omega = 1)}{\sum_{t=1}^T \sum_{(s_1, s_2, \dots, s_t) \in P(a)} \Pr(s_1, s_2, \dots, s_t | \omega = 0)} &= \\ \frac{\sum_{t=1}^T \sum_{(s_1, s_2, \dots, s_t) \in P(a)} q_1(s_1)q_2(s_2)\dots q_t(s_t)}{\sum_{t=1}^T \sum_{(s_1, s_2, \dots, s_t) \in P(a)} (1 - q_1(s_1))(1 - q_2(s_2))\dots(1 - q_t(s_t))}, \end{aligned}$$

where $q_i(s_i)$ is equal to q_i if $s_i = 1$ and to $1 - q_i$ if $s_i = 0$.

Given these likelihood ratios, one can obtain the actions $\alpha(s_1)$, $\alpha(s_1, s_2), \dots$, $\alpha(s_1, s_2, \dots, s_t)$. Then, for this to be an equilibrium, one has to verify that these actions are measurable with respect to the partitions, that is, $\alpha(s_1, s_2, \dots, s_t) = \alpha(s'_1, s'_2, \dots, s'_t) = a^*$ if and only if (s_1, s_2, \dots, s_t) and $(s'_1, s'_2, \dots, s'_t)$ belong to the same partition $P(a^*)$.

Now that we have illustrated the general construction of an ABEE, let us turn to the uniqueness issue. From this construction, we can show that the ABEE is generically unique. To see this, let us define the set of all partitions of Ψ by $C(\Psi)$ and let $C_j(\Psi) \in C(\Psi)$, for $j = 1, 2, \dots, |C(\Psi)|$.²⁰ Consider the set

$$Q_j = \left\{ \frac{\sum_{t=1}^T \sum_{(s_1, s_2, \dots, s_t) \in C_j(a_i)} \prod_{i=1}^t q_i(s_i)}{\sum_{t=1}^T \sum_{(s_1, s_2, \dots, s_t) \in C_j(a_i)} \prod_{i=1}^t (1 - q_i(s_i))}, \text{ for all } C_j(a_i) \in C_j(\Psi) \right\},$$

and let $Q = \bigcup_j Q_j$, and let z_k 's be elements of Q . Note that, for any T , the sets Q_j 's are finite since so is any $C_j(\Psi)$. The (subjective) likelihood ratio that any agent t forms results from the product of elements of the set Q times the multiplier $\left(\frac{q_t}{1-q_t}\right)^{2s_t-1}$ (for the private signal). In particular, given a conjectured partition $\mathcal{P}(\Psi)$, they come from the product of elements of the set Q_i , where $C_i(\Psi) \equiv \mathcal{P}(\Psi)$.

Clearly, if we impose the condition that whenever, for some $z_i \in Q_i$ and $z'_j \in Q_j$, $z_1 z_2 \dots z_l \left(\frac{q_t}{1-q_t}\right)^{2s_t-1} = z'_1 z'_2 \dots z'_{l'} \left(\frac{q_t}{1-q_t}\right)^{2s_t-1}$, then $l = l'$ and for any z_k ($k = 1, \dots, l$) there exists a z'_h ($h = 1, \dots, l'$) such that $z_k = z'_h$, we have that the ABEE is unique: it is the equilibrium in which $\mathcal{P}(\Psi)$ is the finest partition (i.e., whose elements are all singletons). Indeed, under this condition, if we start from a partition in which each action is chosen only for one sequence of signal realizations, we then verify that this is indeed the case. No other conjecture on the partition would be verified, since the condition guarantees that each sequence of realizations leads to a different action. Finally, note

²⁰We use the symbol $|C(\Psi)|$ to denote the cardinality of the set.

that the condition under which the ABEE is unique is generically satisfied: therefore, our ABEE is unique for almost all (q_1, \dots, q_T) .²¹

5.2 The case of signals of equal precision

We now want to consider the case in which the precisions of the signals are all identical, so that assumption (A1) is violated. We will show that in this case the existence of an ABEE is not guaranteed, which justifies the choice of the genericity assumption (A1) in the first place.²²

To see this, we consider a simple example with $T = 3$ and in which the precision is equal to 0.7 for all signals. As a first step, we try to construct an equilibrium based on our usual conjecture that, for any two different sequences of signal realizations, an agent chooses two distinct actions. We will see that this conjecture is actually incorrect.

Following the reasoning of the previous sections, time 1 agent would choose 0.3 or 0.7 upon receiving $s_1 = 0$ or $s_2 = 1$. Under the maintained conjecture, agent 2 would then choose 0.15517 or 0.5 or 0.84483. Consider now agent 3's decision after observing $a_1 = 0.3$, $a_2 = 0.5$ and a private signal $s_3 = 0$. Under our conjecture, he would double count the signal leading to the first action, $s_1 = 0$, and count once the signal $s_2 = 1$ thus obtaining a likelihood ratio equal to $\frac{(0.3)^2(0.7)(0.3)}{(0.7)^2(0.3)(0.7)} = \frac{(0.3)^2}{(0.7)^2}$. Note, however, that this is the same likelihood ratio of an agent at time 2 after $a_1 = 0.3$ and $s_2 = 0$. Therefore, the conjecture would not be verified, since the same action 0.15517 would be chosen after $\{s_1 = 0, s_2 = 0\}$ as well as after $\{s_1 = 0, s_2 = 1, s_3 = 0\}$. Note also that, if, instead, we conjectured that action 0.15517 is chosen after $\{s_1 = 0, s_2 = 0\}$ as well as after $\{s_1 = 0, s_2 = 1, s_3 = 0\}$, agent 3, using our equilibrium concept, would then compute $\bar{\sigma}(a = 0.15517|\omega)$ as equal to $\frac{1}{3} \Pr(s_1 = 0, s_2 = 0|\omega) + \frac{1}{3} \Pr(s_1 = 0, s_2 = 1, s_3 = 0|\omega)$. Therefore, he would not anymore have a likelihood ratio of $\frac{(0.3)^2(0.7)(0.3)}{(0.7)^2(0.3)(0.7)}$ and would not anymore choose 0.15517, making the conjecture wrong again.

But even more is true. We have verified that whatever partition of the set $\Psi = \{(s_1, \dots, s_t) | t = 1, 2, 3\}$ one conjectures, the consistency requirement is never satisfied.²³ In other words, this is an

²¹This condition is slightly stronger than the condition used in the previous analysis. Indeed, the genericity assumption of Section 2 can be derived from this by considering only the finest partition, since in that case the elements of \mathcal{Q} have all the form $\prod_{i=1}^t \frac{q_i(s_i)}{1-q_i(s_i)}$.

²²Obviously, we can choose precisions arbitrarily close to each other and still satisfy (A1).

²³The set Ψ contains 14 elements, and the number of partitions $\mathcal{P}(\Psi)$ is obviously huge. The analysis is, however, simplified by two considerations. First, after any sequence belonging to the set which included $s_1 = 0$ ($s_1 = 1$) the consistency requires that the chosen action is 0.3 (0.7). Second, since the aim is to prove an inconsistency, finding it for two elements of a conjectured partition is already enough. Since these two elements can belong to many partitions, this considerably reduced the number of cases one has to study. The complete analysis is tedious but simple, and is

example in which no ABEE exists.

We conjecture that restricting attention to the equal precision case, for generic values of q an ABEE would fail to exist. This does not contradict the previous existence results by Jehiel (2005) and Jehiel and Koessler (2008), since they consider cases in which the space of actions is finite. What causes the inexistence of an ABEE in our framework is that the action space is continuous.

6 Beyond the Canonical Model

So far we have presented the model of social learning with continuous action space in its simplest form. Now we discuss other variants of the model, thereby shedding more light on the robustness of our results. Throughout this Section, we maintain our assumption (A1).

6.1 Stochastic Signal Precision

One aspect of our analysis is that agents are asymmetrically privately informed, in the sense that the signals they receive have all different precisions. We now move to a set up in which there is symmetry at least in an ex-ante sense. While we maintain the assumption that signal precisions are different, we assume that they are randomly allocated to agents. In other words, given a set of signal precisions $\{q_{(1)}, q_{(2)}, q_{(3)}, \dots, q_{(T)}\}$, the actual sequence of signal precisions may be given by any permutation of this set.²⁴ Each agent t will be assigned a signal of a particular precision, which will be his own private information. If we also assume that each signal precision is assigned to each agent with the same probability, we have agents whose private information is, ex-ante, of the same precision.²⁵

The social learning problem is now more complex than in the original set up given that histories are now parameterized by the sequence of signal precisions together with the signal realizations.²⁶ Despite this extra level of complexity, we now show that there exists an ABEE in which, if the sequence of signal precisions and signal realizations up to t is $(q_{(1)}, s_1, \dots, q_{(t)}, s_t)$ then agent t chooses

available on request from the authors.

²⁴Of course, the same assumptions illustrated in Section 2 apply to the set of signals $\{q_{(1)}, q_{(2)}, q_{(3)}, \dots, q_{(T)}\}$.

²⁵We can assume that the stochastic process according to which the signals are allocated to subjects is common knowledge, but this is not necessary for our result, as will become clear.

²⁶In particular, note that while in our original analysis we defined

$$\bar{\sigma}(a|\omega) = \frac{\sum_{t=1}^T \sum_{h_t, s_t} \sigma_t(a|h_t, s_t) \mu^\sigma(h_t, s_t|\omega)}{T},$$

now, since signal precisions are random, we must re-define it as $\bar{\sigma}(a|\omega) = \frac{\sum_{t=1}^T \sum_{j=1}^{T!} \sum_{h_t, s_t} \sigma_t(a|h_t, s_t, P_j, \omega) \mu^\sigma(h_t, s_t, P_j|\omega)}{T!T}$, where the second summation considers all the $T!$ possible permutations P_j of the signal precisions $\{q_{(1)}, q_{(2)}, \dots, q_{(T)}\}$.

the action $\alpha(q_{(1)}, s_1, \dots, q_{(t)}, s_t)$ such that

$$\frac{\alpha(q_{(1)}, s_1, \dots, q_{(t)}, s_t)}{1 - \alpha(q_{(1)}, s_1, \dots, q_{(t)}, s_t)} = \prod_{i=1}^{t-1} \left(\frac{q_{(i)}}{1 - q_{(i)}} \right)^{(2s_i-1)(t-i)} \left(\frac{q_{(t)}}{1 - q_{(t)}} \right)^{(2s_t-1)}.$$

The reason is that, under our genericity assumption, action $\alpha(q_{(1)}, s_1, \dots, q_{(t)}, s_t)$ can only arise at date t if the sequence of signal precisions and signal realizations is $(q_{(1)}, s_1, \dots, q_{(t)}, s_t)$. Hence, we obtain similar likelihood ratios as those of Proposition 1, when the sequence of signal precisions is known to be $(q_{(1)}, \dots, q_{(t)})$.

The analysis of the PBE is also similar to that in the original set up. Indeed, the agents are able to infer the precision of signals by observing the sequence of actions, since our genericity assumption guarantees that action 1 can only be taken by an agent receiving a signal of a particular precision and realization, action 2 can only be taken after observing action 1 and a signal of a particular precision and realization at time 2, and so on. Therefore, rational agents choose the same actions as in the original set up in which (q_1, \dots, q_T) coincide with the realizations of the precisions at the various dates in this stochastic environment.

We summarize these findings in the next proposition:

Proposition 4 *When the signal precisions are stochastic, there exists an ABEE in which, after a sequence of signal precisions and signal realizations $(q_{(1)}, s_1, \dots, q_{(t)}, s_t)$, agent t chooses action $a_t^* = \alpha(q_{(1)}, s_1, \dots, q_{(t)}, s_t)$ such that*

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

There exists a unique PBE. In the PBE, after the sequence $(q_{(1)}, s_1, \dots, q_{(t)}, s_t)$, agent t chooses action $a_t^{PBE} = \alpha^{PBE}(q_{(1)}, s_1, \dots, q_{(t)}, s_t)$ such that

$$\frac{a_t^{PBE}}{1 - a_t^{PBE}} = \prod_{i=1}^t \left(\frac{q_i}{1 - q_i} \right)^{2s_i-1}.$$

6.2 Unknown Order of Play

The social learning literature in the tradition of Bikhchandani *et al.* (1992) has sometimes been criticized on the ground that it assumes agents know in which order all previous actions have been chosen (see, e.g., Smith and Sørensen, 2008). We now remove this assumption and discuss the case of an unknown order of play: agents know the previous actions, but not the order in which they were

taken.²⁷ Therefore, the history h_t that agent t observes is now the set $\{a_{(1)}, a_{(2)}, a_{(3)}, \dots, a_{(t-1)}\}$, where $a_{(i)}$ is the action taken by one of t 's predecessors. In other words, the true sequence of actions can be any permutation of this set. Our basic construction of an ABEE is robust to this change of observational assumptions.

Proposition 5 *When the order of play is unknown, there exists an ABEE in which, after a sequence of signals s_1, s_2, \dots, s_t , agent t chooses action $a_t^* = \alpha(s_1, s_2, \dots, s_t)$ such that $\frac{a_t^*}{1-a_t^*} = \prod_{i=1}^{t-1} \left(\frac{q_i}{1-q_i}\right)^{(2s_i-1)(t-i)} \left(\frac{q_t}{1-q_t}\right)^{2s_t-1}$.*

This result is a simple consequence of the way inferences are formed in an ABEE. Given that agents conjecture that the distributions of actions are independently distributed across periods (with a different distribution for each state), knowing in which order actions have been chosen has no effect on the updated belief, which in turn implies that an ABEE when the order of actions is unknown remains an ABEE when the order is known.

We note that the same robustness would also hold for the Perfect Bayesian Equilibrium (PBE), although for a quite different reason. In the perfect rationality case and under our genericity assumption, agents can infer the order of actions, as different order of moves would result in different actions. This in turn implies that the PBE found when the order of moves is observed remains a PBE when this order is not observed.

6.3 Observing Only Some Predecessors

Assume now that each agent can only observe the action taken by *some* predecessors. We are not the first ones to discuss this case. Çelen and Kariv (2004) have studied social learning when agents can only observe their immediate predecessor's action.²⁸ Banerjee and Fudenberg (2004) and Smith

²⁷In the literature there are other papers that study similar set ups. Collander and Horner (2009) analyze the case in which agents have to make a binary decision after receiving signals of different precision and observing only the total number of previous agents who have opted for either alternative (and not the order in which decisions were taken). They show that in some cases, in equilibrium, agents follow the minority action. Guarino *et al.* (2007) study the case in which only one of two possible actions is observable. For instance agents can observe the number of previous investments, but not the number of predecessors who had the opportunity to invest and passed on it. They show that only a cascade on the observable action can occur. Larson (2008) analyzes a situation in which agents observe the average action of a population of their predecessors (before making a choice in a continuous action space). The focus of the study is on the speed of learning.

²⁸In this study, agents can only choose between two actions. The result is that behavior does not settle on a single action. Long periods of herding (i.e., conformity of actions) are observed, but switches to the other action occur. As time passes, the periods of herding become longer and longer, and the switches increasingly rare.

and Sørensen (2008) have studied, in different contexts, the case in which agents observe random samples of past actions.²⁹ We analyze both cases.

Observing Only the Immediate Predecessor

In the case in which agents can only observe the action of their immediate predecessor, we get that the ABEE coincides with the PBE.

Proposition 6 *When agents can only observe their immediate predecessor's action, there exists an ABEE (and a unique PBE) in which, after a sequence of signals s_1, s_2, \dots, s_t , agent t chooses action $a_t^* = \alpha(s_1, s_2, \dots, s_t)$ ($a_t^{PBE} = \alpha^{PBE}(s_1, s_2, \dots, s_t)$) such that $\frac{a_t^*}{1-a_t^*} = \frac{a_t^{PBE}}{1-a_t^{PBE}} = \prod_{i=1}^t \left(\frac{q_i}{1-q_i} \right)^{(2s_i-1)}$.*

The intuition for this result is very simple. In the original analysis of Section 3, signal 1 has weight $t-1$ since it enters the Bayesian updating through the observation of agent 1's action, agent 2's action...up to agent $t-1$'s action. Similarly for the other signals. Here, instead, a signal in the sequence (s_1, \dots, s_{t-1}) is counted only once, since it enters the Bayesian updating through the only observation available to agent t , which is the action taken by agent $t-1$.³⁰

Observing Only n Immediate Predecessors From the previous discussion it should be clear that the equivalence between the two concepts (i.e., the lack of bias in an ABEE) is a special result of the case in which only one predecessor is observed. As soon as more than one predecessor is observed, some form of bias reappears. The reason is that when more than one predecessor's action is observed, the same signal may be counted more than once.

Proposition 7 *Suppose agents can only observe the actions of their $n > 1$ immediate predecessors. Then, there exists an ABEE in which, after a sequence of signals (s_1, s_2, \dots, s_t) , agent t chooses action*

²⁹Smith and Sørensen (2008) study a sequential decision model in which agents can only observe unordered random samples from predecessors' actions (e.g., because of word of mouth communication) before making a binary decision. With unbounded private signals complete learning eventually obtains in their model. Banerjee and Fudenberg (2004) present a model in which, at every time, a continuum of agents choose a binary action after observing a sample of previous decisions (and, possibly, of signals on the outcomes). This can be interpreted as a model of word of mouth communication in large populations. The authors find sufficient conditions (on the sampling rule, etc.) for herding to arise, and conditions for all agents to settle on the correct choice.

³⁰A similar observation is made in Eyster and Rabin (2008) in the context of their solution concept, from which they interestingly infer that letting agents observe less may be beneficial to agents when these are boundedly rational. The same conclusion applies to our form of bounded rationality too.

$a_t^* = \alpha(s_1, s_2, \dots, s_t)$ such that

$$\frac{a_t^*}{1 - a_t^*} = \prod_{i=1}^{t-n} \left(\frac{q_i}{1 - q_i} \right)^{n(2s_i-1)} \prod_{i=1}^{n-1} \left(\frac{q_{t-n+i}}{1 - q_{t-n+i}} \right)^{(n-i)(2s_{t-n+i}-1)} \left(\frac{q_t}{1 - q_t} \right)^{2s_t-1}.$$

There exists a unique PBE, in which the optimal action a_t^{PBE} is such that $\frac{a_t^{PBE}}{1 - a_t^{PBE}} = \prod_{i=1}^t \left(\frac{q_i}{1 - q_i} \right)^{2s_i-1}$.

Essentially, when only the n immediate predecessors' actions are observable, in an ABEE agent t counts signals from 1 to $t - n$, a number n of times, signal $t - (n - 1)$ a number $n - 1$ of times, signal $t - (n - 2)$ a number $n - 2$ of times, and so on. The first n signals are counted n times since they enter the likelihood ratio through the observation of all the n predecessors' actions. The signal $t - n + 1$ is counted $n - 1$ time since it enters the likelihood ratio through the observation of the $n - 1$ immediate predecessor, etc.

The equality defining an ABEE in this case is, obviously, identical to that of our main analysis for $n = t - 1$, that is, when the entire history is observable. On the other hand, when only the immediate predecessor is observable, the ABEE does not produce any bias with respect to the PBE. The more predecessors are observable, the higher the bias, in the sense that early signals have a higher and higher weight.

As for the PBE, the analysis is identical to that of the case in which each agent can only observe his immediate predecessor. We know from that analysis that observing the immediate predecessor's action is sufficient to recover the entire history of actions. Of course observing also other previous actions does not modify the inference and the equilibrium outcome.

Observing Only a Sample of Predecessors' Actions We have already observed that in our ABEE the order in which actions are taken does not really matter. It is, therefore, intuitive that the previous analysis extends, with some simple qualifications, to the case in which agents can only observe a sample of n predecessors' actions (knowing or not the order of each action in the sequence).

If the sample consists only of one action, again there will be coincidence between our ABEE (constructed under the usual conjecture) and the PBE, even though ABEE and PBE will now take a different form due to the new sampling assumption. For $n > 1$, our ABEE will, instead, give more weight to earlier actions and a form of bias will appear. Moreover, the overweight of previous signals will be higher, the higher the sample size.

In the interest of space, we only present the result here, and postpone the formal analysis to the Appendix. Let us define the set of indices identifying agent t , his observed predecessors, the

predecessors that they observed and so on by $O_n(t) = \{t, t_{11}, t_{12}, \dots, t_{1n}, t_{111}, t_{112}, \dots, t_{11n\dots}\}$. We have the following proposition:

Proposition 8 *When agents can only observe a sample n of their predecessors' actions, there exists an ABEE in which, after a sequence of signals s_1, s_2, \dots, s_t , agent t chooses action $a_t^* = \alpha(s_1, s_2, \dots, s_t)$ such that $\frac{a_t^*}{1-a_t^*} = \prod_{i \in O_n(t)} \left(\frac{q_i}{1-q_i} \right)^{(2s_i-1)(n-I_i)}$, where the indicator function I_i takes value 0 if agent i belongs to all sets $O_n(t_1), O_n(t_2), \dots, O_n(t_n)$, takes value 1 if agent i belongs to $n-1$ of these sets, and so on.*

6.4 More Actions at the Same Time

In all the previous analysis, we have adopted the traditional sequential decision model, and assumed that at each time t one agent only makes a decision. Our set up, however, is flexible enough to accommodate the case in which several agents make their decisions simultaneously. Nothing substantial changes in our ABEE, although, obviously, we have now to redefine the history of actions until time t , taking into account that at each time there is a number n_t of actions: $h_t = \{a_{1,1}, a_{1,2}, \dots, a_{1,n_1}, \dots, a_{t-1,1}, a_{t-1,2}, \dots, a_{t-1,n_{t-1}}\}$, where $a_{i,j}$ denotes one of the n_i actions at time i . It is easy to show that one can construct an ABEE exactly as in the main analysis, conjecturing that, for any two different sequences of signal realizations, an agent chooses two distinct actions and then observing that our genericity assumption on the signal precisions verifies it. Similarly one can show that there exists a unique PBE in which each action reveals the signal to the subsequent agents. For space constraints we do not present a formal analysis. It is, however, clear that in this set up the bias will be even higher than in the standard set up, since more signals will be overweighted, and, furthermore, the overweight will be higher. For instance, at time 3 any signal $s_{1,j}$ will be overcounted by all agents acting at that time. Furthermore, now it will be counted not twice (as in the main analysis) but $n_2 + 1$ times: once through the observation of action $a_{1,j}$, and n_2 times through the observation of all actions at time 2.³¹ At time 4 it will be counted $n_3 n_2 + n_2 + 1$ by agents acting at that time, with the first term reflecting the count of the signal through the observation of the n_3 actions at time 3.³² The reasoning generalizes to any time t .

7 Social Learning with Discrete Action Space

The previous analysis has shown how coarse inference biases agents' decisions, and affects the process of social learning when the action space is continuous. An obvious question is what happens

³¹Indeed, each action $a_{i,j}$ at time 2 is only taken after a specific sequence $\{s_{1,1}, s_{1,2}, \dots, s_{1,n_1}, s_{2,i}\}$.

³²Each action at time 3 is only taken after a specific sequence $\{s_{1,1}, s_{1,2}, \dots, s_{1,n_1}, s_{2,1}, s_{2,2}, \dots, s_{2,n_2}, s_{2,i}\}$.

in the case in which agents can only choose in a discrete action space. This is, indeed, the set up first analyzed in the social learning literature in the canonical model of Bikhchandani *et al.* (1992). In Bikhchandani *et al.* (1992) model, each agent receives a signal of precision $q \in (0.5, 1)$ and chooses an action in $\{0, 1\}$. We first consider this simple case (mainly for completeness's sake) and then we will consider the case of signals of unbounded precision, which is our main interest. All the other aspects of the model are identical to those described in Section 2. An ABEE can be defined in this set up analogously to the previous, continuous action, case: agents make their inferences simply based on the frequencies of actions $a = 0, 1$ in the two states $\omega = 0, 1$. This means that we can define

$$\bar{\sigma}(a = i|\omega) = \frac{\sum_{t=1}^T \sum_{h_t, s_t} \sigma_t(a = i|h_t, s_t) \mu^\sigma(h_t, s_t|\omega)}{T},$$

for $i = 0, 1$ as the aggregate distribution of actions given the state of the world ω .

7.1 The Case of Bounded Beliefs

In the model of Bikhchandani *et al.* (1992), in equilibrium, agents follow their private signals, unless their predecessors having chosen one option outnumber those in favor of the other by at least two.³³ This is the case in which an informational cascade (i.e., a situation in which agents neglect their private information) arises.

We now show that, when agents make inferences according to the payoff-relevant analogy reasoning, there exists a unique ABEE in which an informational cascade starts already at time 2, with the first agent choosing the action dictated by his signal and all the following ones imitating him.³⁴

Proposition 9 *Suppose agents receive a signal of the same precision $q \in (0.5, 1)$. There exists a unique ABEE in which the first agent in the sequence chooses the action according to his signal ($a_1^* = s_1$) and all the following agents imitate him ($a_t^* = a_1^*$ for all t).*

The proof that such an ABEE exists is very simple. The first agent obviously follows his own signal. Consider now the following agents. If all agents play according to the strategy indicated in

³³This is true under any tie-breaking rule such that, if indifferent, an agent plays the action in agreement with his signal with some positive probability.

³⁴In existing experiments on Bikhchandani et al. (1992), such a conformity of actions is typically not observed. Given the documented tendency of subjects to put relatively more weight on their own signal, this is not surprising. Still, in future work, it would be interesting to study how a modified version of our ABEE which includes the possibility of errors in the laboratory (in the spirit of the Quantal Response Equilibrium) could fit the data of the existing experimental studies.

the proposition, the aggregate distribution of actions as a function of the state of the world ω is given by $\bar{\sigma}(a = 1|\omega = 1) = \bar{\sigma}(a = 0|\omega = 0) = q$. Indeed, given the strategy, all actions will be identical. Whenever the first agent chooses the correct action ($a = \omega$) (which occurs with probability q), so will everyone else; whenever he chooses the incorrect one (which occurs with probability $1 - q$), so will all the following decision makers. Since agents only consider this aggregate distribution, they will form the following, subjective, likelihood ratios:

$$\frac{\Pr(\omega = 1|s_2, a_1)}{\Pr(\omega = 0|s_2, a_1)} = \frac{\bar{\sigma}(a_1|\omega = 1)}{\bar{\sigma}(a_1|\omega = 0)} \left(\frac{q}{1-q}\right)^{2s_2-1} = \left(\frac{q}{1-q}\right)^{2s_1-1} \left(\frac{q}{1-q}\right)^{2s_2-1}$$

for agent 2, and

$$\frac{\Pr(\omega = 1|s_t, a_1, \dots, a_{t-1})}{\Pr(\omega = 0|s_t, a_1, \dots, a_{t-1})} = \frac{\bar{\sigma}(a_1|\omega = 1)}{\bar{\sigma}(a_1|\omega = 0)} \dots \frac{\bar{\sigma}(a_{t-1}|\omega = 1)}{\bar{\sigma}(a_{t-1}|\omega = 0)} \left(\frac{q}{1-q}\right)^{2s_t-1} = \left(\frac{q}{1-q}\right)^{(t-1)(2s_1-1)} \left(\frac{q}{1-q}\right)^{2s_t-1}$$

for any agent t . Given these subjective beliefs, ignoring his own signal and following agent 1 is weakly optimal for agent 2 and strictly optimal for all the others. These best responses clearly make the beliefs on the aggregate distributions correct.

The proof that this is the unique ABEE is only slightly more complicated. Intuitively, in any other ABEE, agents would consider the others' action i either as strictly more informative than their own signal, or as strictly less informative. In the Appendix, we show that in either case a contradiction arises. For instance, if the action $a = i$ is considered strictly more informative (i.e., $\bar{\sigma}(a = i|\omega = i) > q$), then each agent would simply imitate the first agent. Following the argument above, though, this would mean that $\bar{\sigma}(a = i|\omega = i)$ would be equal to q , contradicting that $a = i$ is more informative than the own signal. On the other hand, if the action $a = i$ is considered strictly less informative (i.e., $\bar{\sigma}(a = i|\omega = i) < q$), then each agent would rely more on his own private information and the resulting frequency of actions $a = i$ in state $\omega = i$ could not be strictly lower than q .

We note that in the unique ABEE identified in Proposition 9 agents' beliefs put eventually weight 1 on the state corresponding to the action chosen by the first agent. This is of course very different from the beliefs in the full rationality case (which never exceed the belief arising from having two signals in favor of the chosen state).

7.2 The Case of Unbounded Beliefs

We now modify the simple model, to allow for a more general structure of the private information agents receive. In particular, we now assume that each agent t receives a signal of precision q_t , where q_t is distributed on the support $[0.5, 1]$ according to a density function $f(q_t)$, and a distribution function $F(q)$ where we assume that $f(\cdot)$ is a continuous function with $f(1) > 0$. Note that we are assuming the same density for all times t . Furthermore, the likelihood ratio $\frac{q_t}{1-q_t}$ has support $[0, \infty]$, which means that the distribution of beliefs is unbounded. Since we want to take the length of the sequence T explicitly into account, we find it convenient to redefine the aggregate distributions as $\beta_T(\omega) := \bar{\sigma}(a = \omega|\omega)$.

As a first step in our analysis, we now establish an existence result:

Proposition 10 *Suppose agents receive signals of precision q_t distributed on the support $[0.5, 1]$ according to a density function $f(q_t)$. Then, for any T , there exists a symmetric ABEE with $\beta_T^*(0) = \beta_T^*(1) \equiv \beta_T^* \in (1/2, 1)$.*

The argument for the existence of an ABEE is standard in this finite environment (Jehiel, 2005; Jehiel and Koessler, 2008). Here it is worth noting that in a symmetric equilibrium, given the aggregate distribution β_T^* , agent t chooses his best response on the basis of the following likelihood ratio:

$$\frac{\Pr(\omega = 1|h_t, s_t)}{\Pr(\omega = 0|h_t, s_t)} = \left(\frac{\beta_T^*}{1 - \beta_T^*} \right)^{|a=1| - |a=0|} \left(\frac{q_t}{1 - q_t} \right)^{(2s_t - 1)}, \quad (3)$$

where $|a = i|$ denotes the number of times action i was chosen from period 1 to period $t - 1$ (so that the sum $|a = 1| + |a = 0|$ is obviously equal to $t - 1$). If this likelihood ratio is greater than 1, he chooses action 1, and if it is lower he chooses action 0. The decision of the agent to follow the majority or not depends on the precision of his own signal. A very precise signal induces an agent to follow private information independently of the consensus, even if the majority (i.e., the absolute value of the difference $|a = 1| - |a = 0|$) is large.

Let us now move to the asymptotic properties of our economy. In particular, we want to prove that, when T grows large, the actions cannot settle on the correct one with probability close to 1. Note that if the actions settled on the correct one, an external, Bayesian, observer would learn the true state of the world. In contrast, in our model, since there is no convergence of actions to the correct one, there is no convergence of beliefs to the truth.

Proposition 11 *Suppose agents receive signals of precision q_t distributed on the support $[0.5, 1]$ according to a density function $f(q_t)$. In this economy, there exists no ABEE such that, for $T \rightarrow \infty$, $\beta_T^*(\omega) \rightarrow 1$ for $\omega = 0$ and $\omega = 1$.*

In the Appendix we show that assuming that $\beta_T^*(\omega) \rightarrow 1$ implies that the probability of taking the incorrect action is bounded away from 0, thereby leading to a contradiction. Intuitively, if the aggregate probability of the correct action converged to 1, it would mean that in an ABEE all agents would disregard their own signal and follow the decision of the predecessors (unless they receive an extremely precise signal against previous decisions). Since there is a strictly positive probability (equal to $1 - Eq_1$) that the first agent makes the incorrect choice, this implies that the probability that everyone chooses the incorrect action would be bounded away from zero, contradicting the convergence.

In order to understand better the asymptotic properties of our social learning problem, we have simulated the equilibrium value of β_T for various T assuming q is uniformly distributed on $(\frac{1}{2}, 1)$.³⁵ The equilibrium value β_T^* is approximately equal to 0.79 when $T = 3$, and changes only slightly when we increase T . For large T , β_T^* tends towards 0.82. This means that there is a rather significant inefficiency, since approximately 18% of agents are making the wrong decision. Moreover, the probability that the agent makes a mistake is approximately the same whether an agent acts at time 10 or at time 200, since the impact of early actions becomes soon overwhelming.

It is instructive to compare our result to that in a model of fully rational agents. We know from Smith and Sørensen (2000) that in a model like ours (with unbounded beliefs) but with fully rational agents, beliefs converge to the truth and the actions settle on the correct one (see their Theorems 1 and 3). The bias introduced by the coarse inference has, therefore, a strong long-lasting effect, impeding complete learning. One may wonder what lays at the root of the different results. Intuitively, note that in the case of fully rational agents, unbounded precisions of signals have a very powerful effect for the decision of an agent and of his successors. Even after a large majority of agents have chosen one alternative, an agent with a high precision signal that contradicts the previous history chooses an action against the majority. The probability of this event is never zero. Moreover, after observing the deviation of this agent, the following one will update his belief knowing that the previous agent had received a very high precision signal. Also this agent could, therefore, go against the majority even if he himself does not have a very precise signal. In our model, in contrast, even if an agent receives a very precise signal and goes against the majority, the following agent would not go against the majority unless he himself receives a very precise signal. This is because by considering the aggregate distributions only, agents miss the inference that if someone went against a strong majority it must be that he received a very precise signal.

³⁵To find the fixed point, for a given β_T we have set $\omega = 1$ and simulated the agents' choices according to the best response described above. We have repeated the simulation 100,000 times. We have then computed the empirical frequency of $a = 1$. We have repeated the procedure until the empirical frequency was indeed equal to β_T .

8 Conclusion

Social learning in real economies is a fascinating and complex phenomenon. The models of rational social learning have helped us in understanding many mechanisms through which people learn from others. They have shed light on phenomena such as fads, fashion and cultural change. While we find these models very insightful, in this paper we have offered a new perspective, with less demanding requirements on the rationality of people. Specifically, we have assumed that agents make their inferences based only on how the aggregate distributions of individual actions differ across states rather than on more elaborate statistics that would allow the agents to differentiate these distributions according to the various histories.

Our aim was not to attack all the aspects of rational social learning that one can find problematic, or all conclusions that one can find implausible. We simply studied two environments (the case of the continuous action space and that of unbounded signal precision) where the existing results seem to require more investigation. In the continuous action space, where according to the standard paradigm, history does not matter, we have found that early signals are assigned a higher weight, a bias that echoes the intuitive idea that early choices may have lasting effects in interactions with social learning. In the case of unbounded precision of signals, we have found that the effect of deviators from the crowd is not as strong (and implausible) as implied by the overturning principle, and the process of learning does not necessarily lead agents to settle on the correct action.

While understanding social learning in real economies is the final aim, probably the most sensible environment where to study it is the laboratory, given that our models are highly stylized. A vast experimental literature has shown that the classical model of social learning of Bikhchandani *et al.* (1992) is only partially successful in capturing the behavior human subjects exhibit in the laboratory, and that important deviations from standard equilibrium predictions occur (see, e.g., the recent, comprehensive study by Weizsäcker, 2008). For instance, experimental subjects seem to take into account that the predecessors' actions are based on private information, but fail to internalize completely that the predecessors also went through a similar process of inference from others' actions (see, e.g., Kübler and Weizsäcker, 2004). The view resulting from these experiments is consistent with our approach in which agents make their inferences based on a simple statistical model (which relates the distribution of actions to the state) rather than by putting themselves in the shoes of other agents. Clearly, further experiments (both for continuous action spaces and varying precisions of signals) should be made to test the predictions of our model. This is left for future research.

References

- [1] Acemoglu, D., Ozdaglar, A., and ParandehGreibi, A. (2009) Spread of (Mis)Information in Social Networks. *MIT WP* 09-15.
- [2] Bala V. and Goyal S. (1998) Learning form Neighbours, *Review of Economic Studies*, 65, 595-621.
- [3] Banerjee, A.V. (1992) A Simple Model of Herd Behavior. *Quarterly Journal of Economics* 107, 797-817.
- [4] Banerjee, A.V., and D. Fudenberg (2004) Word-of-Mouth Learning. *Games and Economic Behavior* 46, 1-23.
- [5] Bergemann, D. and Morris, S., 2005, Robust Mechanism Design, *Econometrica*, 73(6), 1771–1813.
- [6] Bikhchandani, S., D. Hirshleifer, and I. Welch (1992) A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades. *Journal of Political Economy* 110, 992-1026.
- [7] Çelen, B., and S. Kariv (2004) Observational Learning under Imperfect Information. *Games and Economic Behavior* 47, 72-86.
- [8] Collander S. and Hörner J. (2009) The Wisdom of the Minority. *Journal of Economic Theory* 144, 1421-1439.
- [9] De Marzo, P., Vayanos, D., and Zwiebel, J. (2003) Persuasion Bias, Social Influence, and Unidimensional Opinions, *Quarterly Journal of Economics*, 118, 909-968.
- [10] Ellison G. and Fudenberg D. (1993) Rules of Thumb for Social Learning. *Journal of Political Economy* 111, 612-643.
- [11] Eyster, E. and Rabin, M. (2005) Cursed Equilibrium. *Econometrica* 73(5), 1623-1672.
- [12] Eyster, E. and Rabin, M. (2008) Naive Herding. *Mimeo*, Berkeley and LSE.
- [13] Gale, D. (1996) What Have We Learned from Social Learning, *European Economic Review*, 40, 617-628.
- [14] Guarino, A., Harmgart, H. and S. Huck (2008) When Half the Truth is Better than the Truth: A Theory of Aggregate Information Cascades, *mimeo*, UCL
- [15] Hirshleifer, D. and S. Teoh (2003) Herd Behaviour and Cascading in Capital Markets: a Review and Synthesis, *European Financial Management*, 9, 25-66.
- [16] Huck, S. and J. Oechssler (1998) Informational Cascades with Continuous Action Spaces, *Economics Letters*, 60, 163-166.
- [17] Informational Cascades in the Laboratory: Do They Occur for the Right Reasons? *Journal of Economic Psychology*, 21, 661-671.

- [18] Jehiel, P., 2005, Analogy-Based Expectation Equilibrium, *Journal of Economic Theory*, 123, 81-104.
- [19] Jehiel, P. and Koessler, F. (2008) Revisiting Games of Incomplete Information with Analogy-Based Expectations, *Games and Economic Behavior*, 62, 533-557.
- [20] Kübler, D. and Weizsäcker, G. (2004) Limited Depth of Reasoning and Failure of Cascade Formation in the Laboratory. *Review of Economic Studies*, 71, pp. 425-441.
- [21] Larson, N. (2008) In with the New Or Out with the Old? Bottlenecks in Social Learning. *mimeo*
- [22] Lee, I. H., 1993, On the Convergence of Informational Cascades, *Journal of Economic Theory*, 61, 395-411.
- [23] Shiryaev, A. (1996) *Probability*, Springer-Verlag, New York.
- [24] Smith, L., and P. Sørensen (2000) Pathological Outcomes of Observational Learning. *Econometrica*, March 2000, 68, pp. 371-398.
- [25] Smith, L., and P. Sørensen (2008) Rational Social Learning with Random Sampling. *Mimeo*.
- [26] Weizsäcker, G. (2008) Do We Follow Others when We Should? A simple test of rational expectations. *Mimeo*, LSE.

9 Appendix

9.1 Proof of Proposition 3

The proof of convergence in the case of the PBE is standard (see, e.g., Shiryaev, 1996), and we omit it. For the proof of convergence in the case of the ABEE, we study the properties of

$$Z_t = \frac{\sum_{i=1}^t u_i}{t(t+1)/2}.$$

We prove the proposition in four steps.

Step 1. Consider a number $a \in (p, 1)$. By applying Chebychev's inequality, we obtain

$$\Pr(Z_t \geq a) = \Pr\left(\lambda \sum_{i=1}^t u_i \geq \lambda a \frac{t(t+1)}{2}\right) = \Pr\left(e^{\lambda \sum_{i=1}^t u_i} \geq e^{\lambda a \frac{t(t+1)}{2}}\right) \leq \frac{E e^{\lambda \sum_{i=1}^t u_i}}{e^{\lambda a \frac{t(t+1)}{2}}},$$

where $\lambda > 0$. Since λ is arbitrary, it is also true that

$$\Pr\left(\lambda \sum_{i=1}^t u_i \geq \lambda a \frac{t(t+1)}{2}\right) \leq \inf_{\lambda > 0} E \left(e^{\lambda \left(\sum_{i=1}^t u_i - a \frac{t(t+1)}{2}\right)} \right) = \inf_{\lambda > 0} E e^{\lambda \sum_{i=1}^t u_i} e^{-\lambda a \frac{t(t+1)}{2}}.$$

Now, note that

$$Ee^{\lambda \sum_{i=1}^t u_i} = \prod_{i=1}^n Ee^{\lambda u_i} = \prod_{i=1}^n ((1-q) + qe^{\lambda i}) = \exp \left\{ \sum_{i=1}^t \log((1-q) + qe^{\lambda i}) \right\},$$

where we use the fact that the random variables u_i are independently distributed. Therefore, we can conclude that

$$\Pr(Z_t \geq a) = \Pr \left(\lambda \sum_{i=1}^t u_i \geq \lambda a \frac{t(t+1)}{2} \right) \leq \inf_{\lambda > 0} \exp \left\{ \sum_{i=1}^t \log((1-q) + qe^{\lambda i}) - \lambda a \frac{t(t+1)}{2} \right\}.$$

Step 2. Now we rewrite this upper bound in a more convenient form. First, let us rewrite the sum in the exponent as follows:

$$\sum_{i=1}^t \log((1-q) + qe^{\lambda i}) - \lambda a \frac{t(t+1)}{2} = \sum_{i=1}^t \left(\log((1-q) + qe^{\lambda i}) - \lambda a i \right).$$

Now let us replace λ with $m = n\lambda$ and obtain

$$\sum_{i=1}^t \left(\log((1-q) + qe^{\lambda i}) - \lambda a i \right) = n \sum_{i=1}^t \frac{1}{n} \left(\log((1-q) + qe^{m \frac{i}{n}}) - m a \frac{i}{n} \right).$$

When n gets large, the sum on the right hand side approaches the Riemann integral of the function $\log((1-p) + pe^{mx}) - amx$, for x that goes from 0 to 1. For large n the following approximation is, therefore, exact:

$$\sum_{i=1}^t \frac{1}{n} \left(\log((1-q) + qe^{\lambda i}) - \lambda a i \right) \approx t \int_0^1 (\log((1-q) + qe^{mx}) - amx) dx.$$

Hence, we can write

$$\Pr(Z_t \geq a) \leq \inf_{m > 0} \exp \left\{ -t \int_0^1 (amx - \log((1-q) + qe^{mx})) dx \right\}.$$

Step 3. Now we look for the m that makes the integral as large as possible, since we want to make our upper bound tight. Note that the integrand $f(y) = (ay - \log((1-q) + qe^y))$ is strictly concave, it takes value 0 when $y = 0$, it is then positive and eventually becomes negative. In

particular, for $x = 1$, it becomes zero in the point $m^* > 0$ that solves $am^* = \log(1 - q + qe^{m^*})$. To maximize the integral, we want to integrate the function under all its positive area. Therefore, we have

$$\Pr(Z_t \geq a) \leq \exp \left\{ -t \int_0^1 \left(am^*x - \log((1 - q) + qe^{m^*x}) \right) dx \right\}.$$

Step 4. Finally, we find an approximation for the integral. The integrand $f(y)$ is maximized at $y = \log \frac{a(1 - q)}{q(1 - a)} = \log \frac{a(1 - q)}{q(1 - a)} > 0$ (since $q < a$). Moreover, at the maximum, its value is $a \log \frac{a}{q} + (1 - a) \log \frac{(1 - a)}{(1 - q)} := H(a) \geq 0$. Now, since $(ay - \log((1 - q) + qe^y))$ is positive in the interval $(0, m^*)$, is concave and has a maximum value $H(a)$, we can draw a triangle underneath it with a base $[0, m^*]$ and a height $H(a)$, and the area of this triangle is a lower bound on our integral. That is,

$$\int_0^1 \left(am^*x - \log((1 - q) + qe^{m^*x}) \right) dx \geq \frac{m^*}{2} H(a) \geq m^*(a - q)^2,$$

where the last inequality comes from the fact that $H(p + x) \geq 2x^2$.

Therefore, we can conclude that

$$\Pr(Z_t - q \geq \varepsilon) \leq e^{-n \frac{m^*}{2} H(q + \varepsilon)} \leq e^{-nm^* \varepsilon^2}.$$

Analogous arguments show that

$$\Pr(Z_t - q \leq -\varepsilon) \leq e^{-n \frac{m^*}{2} H(q - \varepsilon)} \leq e^{-nm^* \varepsilon^2}.$$

Finally, we can conclude that

$$\mathbf{P}(|Z_t - q| \geq \varepsilon) \leq 2e^{-nm^* \varepsilon^2}.$$

9.2 Proof of Proposition 4

We provide the proof for the ABEE. We omit the the proof for the PBE since it is standard.

For the ABEE, consider the decision problem of agent 2. After observing action a_1 and the private signal s_2 the agent has the following likelihood ratio:

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

As before, since the agent computes expectations only considering the average behavior given a state of the world, this ratio will be equal to

$$\frac{\bar{\sigma}(a_1 | \omega = 1)}{\bar{\sigma}(a_1 | \omega = 0)} m(s_2).$$

Recall that in our original analysis we defined

$$\bar{\sigma}(a|\omega) = \frac{\sum_{t=1}^T \sum_{h_t, s_t} \sigma_t(a|h_t, s_t) \mu^\sigma(h_t, s_t|\omega)}{T}.$$

Now, since signal precisions are random, we must re-define it as $\bar{\sigma}(a|\omega) = \frac{\sum_{t=1}^T \sum_{j=1}^{T!} \sum_{h_t, s_t} \sigma_t(a|h_t, s_t, P_j, \omega) \mu^\sigma(h_t, s_t|\omega)}{T!T}$ where the second summation considers all the $T!$ possible permutations P_j of the signal precisions $\{q_{(1)}, q_{(2)}, \dots, q_{(T)}\}$. By conjecture, there is only one sequence of signals associated with each action a . Therefore, the probability of action $a = q_{(i)}$ ($a = 1 - q_{(i)}$) is equal to one for the sequence $\{s_1 = 1\}$ ($\{s_1 = 0\}$) and $q_1 = q_{(i)}$ and zero otherwise. In other words, the action $q_{(i)}$ ($1 - q_{(i)}$) is only taken if the first agent in the sequence receives a signal of precision $q_{(i)}$. This implies that

$$\begin{aligned} \bar{\sigma}(a = q_{(i)}|\omega) &= \\ \frac{\mu^\sigma(h_1, s_1 = 1, q_1 = q_{(i)}|\omega)}{T!T} &= \frac{\Pr(s_1 = 1, q_1 = q_{(i)}|\omega)}{T!T}, \text{ and} \\ \bar{\mu}^\sigma(a = 1 - q_{(i)}|\omega) &= \\ \frac{\mu^\sigma(h_1, s_1 = 0, q_1 = q_{(i)}|\omega)}{T!T} &= \frac{\Pr(s_1 = 0, q_1 = q_{(i)}|\omega)}{T!T}. \end{aligned}$$

Finally, the likelihood ratio will simply be

$$\begin{aligned} \frac{\bar{\mu}^\sigma(a = q_{(i)}|\omega = 1)}{\bar{\mu}^\sigma(a = q_{(i)}|\omega = 0)} m(s_2) &= \frac{q_{(i)}}{1 - q_{(i)}} m(s_2) \\ \text{or} \\ \frac{\bar{\mu}^\sigma(a = 1 - q_{(i)}|\omega = 1)}{\bar{\mu}^\sigma(a = 1 - q_{(i)}|\omega = 0)} m(s_2) &= \frac{1 - q_{(i)}}{q_{(i)}} m(s_2), \end{aligned}$$

where we use the fact that the probability that the precision of the first signal, $q_{(i)}$, is independent of the signal realization. An identical analysis applies to all the following agents.

9.3 Proof of Proposition 5

To prove the result for the ABEE, let us consider the decision problem of agent 3 (the decision problem of agents 1 and 2 is of course identical to that already discussed above, since there is no ambiguity on the order). Agent 3 observes the history $\{a_{(1)}, a_{(2)}\}$, not knowing which action has been taken first. Recall that, by our conjecture, the optimal action a_t^* is only taken after one sequence of signals, that is, $a_t^* = \alpha(s_1, s_2, \dots, s_t)$. Therefore, the first action is actually only taken after a specific realization of s_1 and the second after a specific sequence of realizations s_1 and s_2 . Since the agent considers the aggregate distribution given a particular state of the world, he will again compute his posteriors in the following way:

$$\bar{\sigma}(a_{(1)}|\omega) = \frac{\Pr(s_1|\omega)}{T}, \text{ or } \bar{\sigma}(a_{(1)}|\omega) = \frac{\Pr(s_1, s_2|\omega)}{T},$$

depending on whether $a_{(1)} = a_1$ or $a_{(1)} = a_2$. (A similar expression obviously holds for $a_{(2)}$.) It should be clear that the agent does not really know whether $a_{(1)}$ is the first or the second action, but he knows the probability with which that particular action is taken conditional on the state of the world. And this is the information he uses to update his belief. Therefore agent 3's likelihood ratio can be expressed as

$$\frac{\Pr(\omega = 1|s_3, a_1, a_2)}{\Pr(\omega = 0|s_3, a_1, a_2)} = \frac{\bar{\mu}^\sigma(a_1|\omega = 1)\bar{\mu}^\sigma(a_2|\omega = 1)}{\bar{\mu}^\sigma(a_1|\omega = 0)\bar{\mu}^\sigma(a_2|\omega = 0)}m(s_3) = \frac{\Pr(s_1|\omega = 1)\Pr(s_1, s_2|\omega = 1)}{\Pr(s_1|\omega = 0)\Pr(s_1, s_2|\omega = 0)}m(s_3),$$

which is exactly the same expression obtained in the original analysis. An identical argument shows that the likelihood ratio is the same as that obtained in the original analysis for any agent t .

Let us now prove the result for the PBE. Let us consider the decision problem of agent 3. Although he does not observe the order in which actions $\{a_{(1)}, a_{(2)}\}$ have been taken, he can infer it from the actions themselves. Recall that

$$\begin{aligned} \frac{a_1^{PBE}}{1 - a_1^{PBE}} &= \left(\frac{q_1}{1 - q_1} \right)^{2s_1 - 1} \\ &\text{and} \\ \frac{a_2^{PBE}}{1 - a_2^{PBE}} &= \prod_{i=1}^2 \left(\frac{q_i}{1 - q_i} \right)^{2s_i - 1}. \end{aligned}$$

Only one of the two observed actions can be the first in the sequence, since only one will be consistent with one of the two possible likelihood ratios at time 1: $\left(\frac{q_1}{1 - q_1} \right)^{2s_1 - 1}$. Indeed, by our genericity assumption, it can never be the case that $\left(\frac{q_1}{1 - q_1} \right)^{2s_1 - 1} \left(\frac{q_2}{1 - q_2} \right)^{2s_2 - 1} = \left(\frac{q_1}{1 - q_1} \right)^{-(2s_1 - 1)}$. Therefore, agent 3 will be able to infer which action was taken first, and his updating problem will be identical to the one discussed in the main analysis. He will make his decision exactly as if he observed the order of the sequence. A similar logic applies to the decision problem of agent 4. Again, our genericity assumption guarantees that only one action actually taken by the predecessors will be consistent with one of the two possible likelihood ratios given by the first signal with precision q_1 . This will allow agent 4 to infer which action is a_1^{PBE} . Then, only one action taken by the two remaining predecessors will be consistent with one of the two possible likelihood ratios at time 2: $\frac{a_1^{PBE}}{1 - a_1^{PBE}} \left(\frac{q_2}{1 - q_2} \right)^{2s_2 - 1}$. This will allow the agent to infer which action was a_2^{PBE} and, by exclusion, which one was a_3^{PBE} . Therefore, agent 4 will perfectly infer the sequence of actions and make his decision in the same way as if he observed the entire sequence. Similarly for any agent t .

9.4 Proof of Proposition 6

For the first two agents in the sequence, of course, nothing changes. Let us discuss how agent 3 makes his decision in an ABEE. After observing action $a_2 = \alpha(s_1, s_2)$ and the private signal s_3 the agent has the following likelihood ratio:

$$\frac{\Pr(\omega = 1|s_3, a_2)}{\Pr(\omega = 0|s_3, a_2)} = \frac{\bar{\sigma}(a_2|\omega = 1)}{\bar{\sigma}(a_2|\omega = 0)}m(s_3).$$

Given our conjectures that all $\alpha(s)$'s are different, we obtain that

$$\frac{\bar{\sigma}(a_2|\omega = 1)}{\bar{\sigma}(a_2|\omega = 0)} = \frac{\Pr(s_1, s_2|\omega = 1)}{\Pr(s_1, s_2|\omega = 0)} = m(s_1)m(s_2),$$

and thus agent 3 chooses action $\alpha(s_1, s_2, s_3)$ such that

$$\frac{\alpha(s_1, s_2, s_3)}{1 - \alpha(s_1, s_2, s_3)} = m(s_1)m(s_2)m(s_3).$$

A similar reasoning applies to any time t . The likelihood ratio for agent t can, therefore, be written as

$$\frac{\Pr(\omega = 1|s_t, a_{t-1})}{\Pr(\omega = 0|s_t, a_{t-1})} = \frac{\Pr(s_1, s_2, \dots, s_{t-1}|\omega = 1)}{\Pr(s_1, s_2, \dots, s_{t-1}|\omega = 0)}m(s_t),$$

and the agent chooses action $\alpha(s_1, s_2, \dots, s_t)$ such that

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

for the PBE Since the argument is similar to those presented above, and fairly standard, we refrain from providing a formal proof.

9.5 Proof of Proposition 7

To prove that the ABEE is unique we proceed by contradiction. Suppose that $\bar{\sigma}(a = i|\omega = i) > q$, for $i = 0$ or 1 (or both). Then agents would perceive others' action i as more informative than their own signal. This implies that if the first agent chooses $a = i$, all agents would follow the first agent's action, independently of their signal, yielding in turn $\bar{\sigma}(a = i|\omega = i) = q$, which is a contradiction. Suppose now that $\bar{\sigma}(a = 0|\omega = 0) = \bar{\sigma}(a = 1|\omega = 1) = \bar{q} < q$. Then an agent would follow his own signal, unless the predecessors who have chosen one action outnumber the others by at least $k(\bar{q}) > 1$, where the number $k(\bar{q})$ is the smallest integer such that $\left(\frac{\bar{q}}{1 - \bar{q}} \right)^{k(\bar{q})} > \frac{q}{1 - q}$. In this case, a

cascade arises, and an agent just follows the majority action, independently of his signal. Now, if $k(\bar{q}) \geq T$, a cascade cannot occur and the agents always follow their private signals. Therefore, the probability of action $a = i$ in state $\omega = i$ is equal to q for any time and any history, which implies that $\bar{\sigma}(a = i|\omega = i) = q$, a contradiction. If, instead, $k(\bar{q}) < T$, then $\bar{\sigma}(a = i|\omega = i) > q$, yielding again a contradiction.³⁶ The remaining case in which $\bar{\sigma}(a = 0|\omega = 0) = \bar{q} < q$ and $\bar{\sigma}(a = 1|\omega = 1) = q$ (or vice versa) can also be immediately ruled out based on the same considerations.

9.6 Proof of Proposition 8

As usual we construct an ABEE based on our conjecture and then verifying it. Let us consider the case in which $n = 2$. In this case, the decision process for the first three agents is obviously identical to the main analysis. The difference arises at time 4. Agent 4 observes $a_2 = \alpha(s_1, s_2)$ and $a_3 = \alpha(s_1, s_2, s_3)$ and, given our conjecture, computes the following probabilities:

$$\begin{aligned}\bar{\sigma}(a_2|\omega) &= \frac{\Pr(s_1, s_2|\omega)}{T}, \\ \bar{\sigma}(a_3|\omega) &= \frac{\Pr(s_1, s_2, s_3|\omega)}{T}.\end{aligned}$$

Therefore, agent 4 will compute his likelihood ratio as follows:

$$\begin{aligned}\frac{\Pr(\omega = 1|s_4, a_2, a_3)}{\Pr(\omega = 0|s_4, a_2, a_3)} &= \\ \frac{\bar{\sigma}(a_2|\omega = 1)\bar{\sigma}(a_3|\omega = 1)}{\bar{\sigma}(a_2|\omega = 0)\bar{\sigma}(a_3|\omega = 0)}m(s_4) &= \\ \frac{\Pr(s_1, s_2|\omega = 1)\Pr(s_1, s_2, s_3|\omega = 1)}{\Pr(s_1, s_2|\omega = 0)\Pr(s_1, s_2, s_3|\omega = 0)}m(s_4),\end{aligned}$$

which, in turn, implies that

$$\frac{a_4^*}{1 - a_4^*} = \prod_{i=1}^2 \left(\frac{q_i}{1 - q_i} \right)^{2(2s_i - 1)} \left(\frac{q_3}{1 - q_3} \right)^{2s_3 - 1} \left(\frac{q_4}{1 - q_4} \right)^{2s_4 - 1}.$$

³⁶To show that $\bar{\sigma}(a = i|\omega = i) > q$ it is enough to observe that a correct cascade is more likely than an incorrect one. Formally, at any time $t \leq k(\bar{q})$, the probability that $a = \omega$ is equal to q for any history. Consider now $t = k(\bar{q}) + 1$. There are three possibilities: either a correct cascade (i.e., a cascade in which an agent chooses $a = \omega$ with probability 1) arises; or an incorrect cascade arises; or there is no cascade (in which case the probability that $a = \omega$ is q). (An agent could possibly be mixing when $\left(\frac{\bar{q}}{1-\bar{q}}\right)^{k(\bar{q})} = \frac{1-q}{q}$, but this does not affect the conclusion.) Now, we can show that at time $t = k(\bar{q}) + 1$ the probability of observing $a = \omega$ is greater than q . Indeed, the probability of observing $a = \omega$ is equal to $\left(1 - q^{k(\bar{q})} - (1 - q)^{k(\bar{q})}\right)q + q^{k(\bar{q})}$. This probability is higher than q when $(1 - q^{k(\bar{q})} - (1 - q)^{k(\bar{q})}) + q^{k(\bar{q})-1} > 0$, which is always satisfied since $q \in (0.5, 1)$. Similarly, one can prove that any time $t > k(\bar{q}) + 1$ the probability of $a = \omega$ is again greater than q . Therefore, $\bar{\sigma}(a = i|\omega = i) > q$.

Similarly, agent 5 will observe a_3^* and a_4^* , and form expectations

$$\begin{aligned}\bar{\sigma}(a_3|\omega) &= \frac{\Pr(s_1, s_2, s_3|\omega)}{T}, \\ \bar{\sigma}(a_4|\omega) &= \frac{\Pr(s_1, s_2, s_3, s_4|\omega)}{T}.\end{aligned}$$

Therefore, his likelihood ratio will be

$$\begin{aligned}\frac{\Pr(\omega = 1|s_5, a_3, a_4)}{\Pr(\omega = 0|s_5, a_3, a_4)} &= \\ \frac{\bar{\sigma}(a_3|\omega = 1)\bar{\sigma}(a_4|\omega = 1)}{\bar{\sigma}(a_3|\omega = 0)\bar{\sigma}(a_4|\omega = 0)}m(s_5) &= \\ \frac{\Pr(s_1, s_2, s_3|\omega = 1)\Pr(s_1, s_2, s_3, s_4|\omega = 1)}{\Pr(s_1, s_2, s_3|\omega = 0)\Pr(s_1, s_2, s_3, s_4|\omega = 0)}m(s_5),\end{aligned}$$

which, in turn, implies that

$$\frac{a_5^*}{1 - a_5^*} = \prod_{i=1}^3 \left(\frac{q_i}{1 - q_i} \right)^{2(2s_i - 1)} \left(\frac{q_4}{1 - q_4} \right)^{2s_4 - 1} \left(\frac{q_5}{1 - q_5} \right)^{2s_5 - 1}.$$

The genericity assumption verifies that our conjecture is correct. The analysis immediately extends to any time t . The analysis for $n > 2$ is identical, giving the result on the ABEE in the proposition.

9.7 Proof of Proposition 9

As usual we construct an ABEE based on our conjecture and then verify it. The decision problem for the first two agents is identical to that of the main analysis. Consider the decision problem of agent 3 His posterior likelihood ratio will be:

$$\frac{\Pr(\omega = 1|s_3, h_3^O)}{\Pr(\omega = 0|s_3, h_3^O)} = \frac{\Pr(h_3^O|\omega = 1, s_3)\Pr(\omega = 1|s_3)}{\Pr(h_3^O|\omega = 0, s_3)\Pr(\omega = 0|s_3)}.$$

The agent does not know whether the observed action $h_3^O = \{a\}$ comes from the first or the second agent, that is, whether $a_1 = a$ or $a_2 = a$. This lack of knowledge is, however, irrelevant for his decision. In any case, he will obtain a likelihood ratio equal to

$$\frac{\Pr(\omega = 1|s_3, h_3^O)}{\Pr(\omega = 0|s_3, h_3^O)} = \frac{\bar{\sigma}(a|\omega = 1)}{\bar{\sigma}(a|\omega = 0)}m(s_3).$$

Since the agent computes expectations only considering the average behavior given a state of the world, and given our conjecture, this ratio will be equal to

$$\frac{\Pr(s_1|\omega = 1)}{\Pr(s_1|\omega = 0)}m(s_3), \text{ if } a_1 = a,$$

or

$$\frac{\Pr(s_1, s_2|\omega = 1)}{\Pr(s_1, s_2|\omega = 0)}m(s_3), \text{ if } a_2 = a.$$

The decision problem for agent 4 is slightly more complicated. Now, not only he does not know whether what he is observing is the action of agent 1 or 2 or 3. He does not even know which action agent 3 observed when he made his decision. But, again, this lack of knowledge turns out to be immaterial in the case of our ABEE. As we have explained above, action a_3 will either come from the likelihood ratio

$$\frac{\Pr(s_1|\omega = 1) \Pr(s_3|\omega = 1)}{\Pr(s_1|\omega = 0) \Pr(s_3|\omega = 0)}$$

or from

$$\frac{\Pr(s_1, s_2|\omega = 1) \Pr(s_3|\omega = 1)}{\Pr(s_1, s_2|\omega = 0) \Pr(s_3|\omega = 0)}.$$

Actions a_1 and a_2 will result, as usual, from the likelihood ratios based only on signal s_1 and based on both signals s_1 and s_2 , respectively. In all four possible cases (agent 4 observes agent 1, or agent 2 who has observed agent 1, or agent 3 who has observed either agent 1 or agent 2), the agent will update his belief in the following way

$$\frac{\Pr(\omega = 1|s_4, h_4^O = \{a\})}{\Pr(\omega = 0|s_4, h_4^O = \{a\})} = \frac{\bar{\sigma}(a|\omega = 1)}{\bar{\sigma}(a|\omega = 0)}m(s_4),$$

and the ratio $\frac{\bar{\sigma}(a|\omega=1)}{\bar{\sigma}(a|\omega=0)}$ will be equal to $\frac{\Pr(s_1|\omega=1)}{\Pr(s_1|\omega=0)}$ if $a = a_1$, to $\frac{\Pr(s_1, s_2|\omega=1)}{\Pr(s_1, s_2|\omega=0)}$ if $a = a_2$, to $\frac{\Pr(s_1, s_3|\omega=1)}{\Pr(s_1, s_3|\omega=0)}$ if $a = a_3$ and agent 3 has observed agent 1, and, finally, to $\frac{\Pr(s_1, s_2, s_3|\omega=1)}{\Pr(s_1, s_2, s_3|\omega=0)}$ if $a = a_3$ and agent 3 has observed agent 2.

The same reasoning can be replicated recursively for the subsequent agents. Moreover, our genericity assumption guarantees that our conjecture is correct, thus giving an ABEE.

The result for the PBE follows a standard argument. Agent 3 is perfectly able to infer who made the action a , whether agent 1 or agent 2. Moreover, if it is agent 2, he can also infer the previous action of agent 1. Therefore, he infers the signals from the actions and updates accordingly. The same reasoning applies to all subsequent agents.

9.8 Proof of Proposition 10

We want to show the existence of a symmetric equilibrium in which $\beta_T(0) = \beta_T(1) \equiv \beta_T \in (1/2, 1)$. The argument is standard. Let us define the function $\phi : [0.5, 1] \rightarrow [0.5, 1]$ that, for a given β_T , gives the aggregate distribution of actions conditional on a state of the world, $\phi(\beta_T)$. We want to prove that there exists a β_T such that $\beta_T = \phi(\beta_T)$. Recall that for a given β_T the best response of agent t consists in choosing action $a_t = 1$ (0) when the likelihood ratio

$$\frac{\Pr(\omega = 1|h_t, s_t)}{\Pr(\omega = 0|h_t, s_t)} = \left(\frac{\beta_T}{1 - \beta_T}\right)^{|a=1|-|a=0|} \left(\frac{q_t}{1 - q_t}\right)^{2s_t-1}$$

is greater (lower) than 1.

Consider, first, the case of $\beta_T = 1/2$. In this case, other agents' decisions are perceived as uninformative. Therefore, each agent just follows his private signal. As a result, the aggregate distribution of $a = \omega$ will be $\phi(\beta_T) = E(q_t) > (1/2, 1)$, which implies that $\phi(\beta_T) > \beta_T$. Consider, now, the case of $\beta_T = 1$. In this case, other agents' decisions are perceived as perfectly informative. Therefore, after the first agent chooses (by following his own signal), all the others simply imitate him. As a result, the aggregate distribution of $a = \omega$ will be again $\phi(\beta_T) = E(q) \in (1/2, 1)$, which implies that $\phi(\beta_T) < \beta_T$. Finally, notice that, at any time t , the probability that an agent receiving signal ω chooses $a = \omega$ is 1 if $\left(\frac{\beta_T}{1 - \beta_T}\right)^{|a=\omega|-|a=1-\omega|} \geq 1$ and is $1 - F(\beta_T^{|a=\omega|-|a=1-\omega|})$ otherwise. Since q_t is distributed according to a continuous density function, $\phi(\beta_T)$ is a continuous function. By Brouwer's fixed point theorem, there exists a $\beta_T \in (1/2, 1)$ such that $\phi(\beta_T) = \beta_T$.

9.9 Proof of Proposition 11

Without loss of generality, suppose the state of the world is $\omega = 0$. Suppose $\beta_T^*(0) \rightarrow \hat{\beta}$, for $\hat{\beta} = 1 - \varepsilon$, and for a small ε .

First, note that if $\beta_T^*(0) \rightarrow \hat{\beta}$, then an agent receiving a signal of precision $q_t < \hat{\beta}$ will always follow the majority. Indeed, if there is a majority in favor of one action, it means that there is at least one more predecessor in favor of such an action. Consider the least favorable case in which the majority is by only one. Clearly, the likelihood ratio will be $\left(\frac{\hat{\beta}}{1 - \hat{\beta}}\right)^I \left(\frac{q_t}{1 - q_t}\right)^{2s_t-1}$, where I takes value 1 if the majority action is 1 and -1 if it is 0. This likelihood ratio will be higher or lower than 1 (and so the action chosen by the agent will be either 1 or 0) depending only on the choice of the predecessors, since $q_t < \hat{\beta}$. The argument will hold a fortiori if the majority is by a number greater than one.

Now, we find a lower bound for the probability of the incorrect action. Consider the event that the first agent receives the signal $s_1 = 1$ (an event that occurs with probability $1 - Eq_1$). The first

agent obviously chooses $a_1 = 1$. Now, consider the second agent. If he receives the signal $s_2 = 1$, obviously he chooses $a_2 = 1$. If he observes $s_2 = 0$, he chooses $a_2 = 1$ with the same probability as the probability that $q_2 < \widehat{\beta}$, $F(\widehat{\beta})$, as we know from the previous reasoning. Using Taylor's expansion, it is easy to show that this probability is $F(\widehat{\beta}) > 1 - 2f(1)\varepsilon$, for $\varepsilon < \bar{\varepsilon}$ and some $\bar{\varepsilon} > 0$.

Similarly, if the third agent receives the signal $s_3 = 1$, obviously he chooses $a_3 = 1$. If he observes $s_3 = 0$, he chooses $a_3 = 1$ with a probability equal to $F(\widehat{\beta}^2)$. Using again Taylor's expansion, it is easy to prove that this probability is greater than $1 - 4f(1)\varepsilon^2$ for $\varepsilon < \bar{\varepsilon}$, and some $\bar{\varepsilon} > 0$. A similar analysis proves that if agent t observes $s_t = 0$, he chooses $a_t = 1$ with a probability not lower than $1 - 2^{t-1}f(1)\varepsilon^t$.

Therefore, the probability that every agent t chooses action $a_t = 1$ is higher than

$$(1 - Eq_1) \prod_{k=2}^T \left(1 - 2^{k-1}f(1)\varepsilon^{k-1}\right).$$

Taking the logarithm of $\prod_{k=1}^T (1 - 2^k f(1)\varepsilon^k)$, we obtain

$$\log \prod_{k=1}^T \left(1 - 2^k f(1)\varepsilon^k\right) = \sum_{k=1}^T \log \left(1 - 2^k f(1)\varepsilon^k\right),$$

which, using again Taylor's expansion, can be shown to be greater than

$$\sum_{k=1}^T -a2^k f(1)\varepsilon^k$$

for some $a > 1$. Therefore, the probability that every agent chooses the incorrect action is bounded below by

$$(1 - Eq_1) \exp\left\{-af(1) \sum_{k=1}^T 2^k \varepsilon^k\right\},$$

which, for $T \rightarrow \infty$, is equal to

$$(1 - Eq_1) \exp\left\{-af(1) \frac{2\varepsilon}{1 - 2\varepsilon}\right\}.$$

This expression is close to $(1 - Eq_1) > 0$ for ε close to zero, thus contradicting that $\beta_T^*(0) \rightarrow \widehat{\beta} = 1 - \varepsilon$.




Figure 1: Average distance of the public belief from the fundamental value. The solid line refers to the PBE. The dotted line refers to the ABEE.




Figure 2: Average distance of the public belief from the fundamental value after 5 incorrect signals. The solid line refers to the PBE. The dotted line refers to the ABEE.