Social Learning with Coarse Inference

Antonio Guarino and Philippe Jehiel

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 make perfect inferences from their predecessors' decisions: they only understand the relation between the aggregate distribution of actions and the state of nature and make their inferences accordingly. We show that, in a discrete action space, even if agents receive signals of unbounded precision, convergence to the truth does not occur. In a continuous action space, compared to the rational case, agents overweight early signals. Despite this behavioral bias, convergence to the truth eventually obtains.

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. 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. The process of social learning is, instead, eventually efficient in the discrete action space case whenever signals can have unbounded precision (and agents are fully rational) (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.

In another important modification of the basic herd model, agents choose an action (in a continuous action space) that matches their expectation about the state of the economy (Lee, 1993). When agents are fully rational, the process of social learning is eventually efficient because agents can perfectly infer the history of 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.

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 agent 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 address this: either modify some aspects of the game while maintaining the rationality assumptions, 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 Specifically, we develop an equilibrium

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. We will discuss this again after we have presented our main results.

²For example, if along with Smith and Sørensen (2008), we assume the order of moves is not observed, the rational inferences in such a model are even more complex than in the standard model.

approach in which agents make their inferences 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 make inferences based on 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 agents' understanding, although partial, is assumed to be correct: the state-dependent distribution of actions assumed by the agents matches the aggregate long run frequencies of actions in each state. In a learning dynamics, this approach requires that agents only keep track of (or pay attention to) how the frequencies of actions depend on the state. Agents do not need any further knowledge about other agents' preferences, information of ways of reasoning. In fact, an important motivation for our approach is that, in a number of real life social interactions, individuals do not really know other agents' exact structure of information, preferences and modes of reasoning.³

Formally, our analysis relies on the Analogy Based Expectation Equilibrium, developed by Jehiel (2005), in which agents use payoff-relevant analogy partitions (as defined in Jehiel and Koessler, 2008) to form their expectations about others' behaviors and make inferences from observed actions as to the likelihood of the state of the economy. Agents use the payoff-relevant analogy partition in that they relate others' behaviors to what they only care about in terms of payoffs, that is, the state of the economy (as opposed to what may drive the motives of others' decisions, i.e., their preferences, their information and possibly their mode of reasoning).

We apply the Analogy Based Expectation Equilibrium both to the discrete and to the continuous action space set ups. In both set ups, our analysis provides new insights. A key finding in the case of discrete action space is that the overturning principle does not hold when signals are of unbounded precision. Longer "herds" are less fragile and a single deviation is not enough to destroy it. The long run consequence is that actions do not settle on the correct one with probability close to 1, a result in stark contrast with the existing literature. In the case of a continuous action space, we find an interesting behavioral bias. Early actions have an overwhelming importance for subsequent decisions. In that sense, in contrast with the standard approach, history does matter even in a

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. Relatedly, one may argue that even if such information is available, looking at previous social learning experiences it is probably easier to remember how actions are distributed as a function of the state of the world than to know how actions depend on the private history of decisions.

continuous action space. Observing the entire history of actions leads the decision maker to make (biased) choices that he would not make, were he only able to observe his immediate predecessor(s). Despite this bias, though, in this set up, convergence to the truth eventually obtains.

1.1 Related literature

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. As a result, both in a discrete or in a continuous action space, 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 considered in Eyster and Rabin, 2005), early signals could not have a stronger effect 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.⁴

Elaborating on their previous work, in a recent paper, Eyster and Rabin (2010) consider a 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. A key observation in their model 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. It should be noted that there is no analog of our treatment of the discrete action space with varying precision in Eyster and Rabin (2010).

⁴Note 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). We believe the payoff relevant analogy partition is more suited to the analysis of social learning given that it is more salient to remember from past interactions how actions are distributed as a function of the underlying state than it is to remember how actions are distributed as a function of the own private information. The payoff-relevant analogy partition also allows for some form of non-trivial inference unlike the private information analogy partition.

A common feature of Eyster and Rabin (2010) and our paper is that agents do not fully understand other agents' strategies. But, there are important qualitative differences in the two approaches. Eyster and Rabin's approach implicitly requires that agents have a good knowledge and understanding of others' preferences and of the distribution of private signals. Our approach requires instead some form of learning, but is less demanding in terms of what agents must know about other agents' private information and motivation for their choices (preferences).⁵

Another paper that applies subjective priors to social learning is Bohren (2009). In her model, some agents only use private signals to make their choices, while others use all the available (public and private) information. There is, however, uncertainty on the proportions of the two types of agents. Moreover, agents can have incorrect beliefs about these proportions. In a set up with discrete action space, when agents overestimate the proportion of the first type of individuals, incorrect herds can persist forever. In contrast, when they overestimate the proportion of the second type of individuals, correct herds may break and beliefs fluctuate forever.

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 whereas our approach maintains the basic ingredients of Bayesian updating even if applied to a misspecified model (as motivated by incomplete learning considerations). 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 discrete action space, define the solution concept and analyze both the short and the long run properties. In Section 3 we study the case of a continuous action space. Section 4 concludes. An Appendix contains the proofs.

⁵Approaches based on subjective priors are difficult to justify from a learning perspective (Dekel *et al.*, 2004), in contrast with the equilibrium approach pursued here.

2 Social Learning with Discrete Action Space

Imagine we run the following experiment. A group of subjects have to choose in sequence either action "a" or action "b" (with action a (b) giving the highest payoff when the state is A (B)). Before making his decision, a subject observes previous decisions and receives a signal with varying precision on whether the true state of the world is A or B. The experiment is repeated many times. After each repetition, each subject is informed of the true state of the world and of his payoff. Suppose now at the end of the experiment, we consider all the repetitions in which 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 a0, instead, action a1 was chosen 35% of the time and action a2 Suppose you were told these empirical frequencies. How would you play in this experiment if you had the possibility of participating in it?

Our approach assumes that agents expect that the probability of action a conditional on the state A is 67% and that the probability of action b conditional on the state B is 65% irrespective of the time when the action is taken and the specific sequence of actions that the agent has observed until then. Moreover, agents best respond to these frequencies, that is, they choose the best action assuming others' actions in state A or B are distributed according to these aggregate frequencies, no matter what specific sequence of actions was observed.

Putting the example of the laboratory experiment aside, such a mode of reasoning is — we believe — plausible when, in social interactions, agents have doubts about what drives the behavior of others, or simply do not know the payoff or the information that previous agents had when they made their decisions. Instead of making complicated and speculative inferences, these agents simply connect the state of the world (they care about) to the empirical frequencies of actions that they observe. In our formal analysis, we consider a steady state of such a dynamic process, in which the behaviors induced by the past empirical frequencies of actions by state give rise to the same empirical frequencies. As we said in the introduction, this idea is formalized in a solution concept called the Analogy Based Expectation Equilibrium. We define it in subsection 2.1 and analyze it in subsection 2.2.

2.1 The Model

In our economy there are T agents who make a decision in sequence. Time is discrete and indexed by t = 1, 2, ..., 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 words 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 $\omega \in \{0,1\}$. The two states of the world are equally likely. If $\omega = 1$, an agent receives a payoff of 1 if he chooses action 1, and a payoff of 0 otherwise; vice versa if $\omega = 0$.

We denote the history of actions until time t-1 by h_t , that is, $h_t = \{a_1, a_2, ..., a_{t-1}\}$ (and $h_1 = \emptyset$). We denote the set of such histories by H_t . We assume that agent t observes the entire history of actions h_t . In addition to observing the sequence of actions taken by the predecessors, each agent observes a symmetric binary private signal s_t , distributed as follows:

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

where q_t represents the precision of the signal (and is private information to agent t). We also assume that each signal's precision q_t is distributed on the support [0.5, 1] according to a density function $f(q_t)$ with cumulative distribution $F(q_t)$, where f(.) is continuous and f(1) > 0. The random variables q_t and $q_{t'}$ for $t' \neq t$ are independent of each other; given q_t and $q_{t'}$, and conditional on the realized state ω , the private signals s_t and $s_{t'}$ are independently distributed. Agent t's information set is represented by the triple (h_t, s_t, q_t) . Given the information (h_t, s_t, q_t) , the agent chooses a_t to maximize his expected payoff.

The main innovation of our work is in modelling how agents make inferences from past actions as to the likelihood of ω . We adopt the Analogy Based Expectation Equilibrium concept first introduced by Jehiel (2005) using, more specifically, the payoff-relevant analogy partition (Jehiel and Koessler, 2008). We assume that agents are unable to understand 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. 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

⁶Observe that we are considering a time-independent distribution of precisions. Furthermore, the distribution of the likelihood ratio $\frac{q_t}{1-q_t}$ has support $[0,\infty)$, which means that the distribution of beliefs is unbounded.

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, $\sigma = (\sigma_1, \sigma_2, ... \sigma_T)$, where an agent's strategy σ_t maps (h_t, s_t, q_t) into a distribution of possible actions, that is, $\sigma_t : H_t \times \{0, 1\} \times [0.5, 1] \to \Delta\{0, 1\}$. For consistency with later sections, and with a slight abuse of notation, we denote by $\sigma_t(a|h_t, s_t, q_t)$ the probability that agent t picks action $a \in \{0, 1\}$ when the history is h_t , the signal is s_t , and the precision is q_t .

Let us consider a vector of realized precisions $q \equiv (q_1, q_2, ..., q_T)$, and let $\mu^{\sigma}(h_t, s_t, q|\omega)$ denote the probability that history h_t is realized, s_t is the signal at t, and q is the vector of precisions when ω is the state of the world. Observe that, for any t, there are finitely many (h_t, s_t) , so that, for a given vector of precisions q, $\mu^{\sigma}(h_t, s_t, q|\omega) > 0$ only for finitely many (h_t, s_t) . Given the strategy σ , the aggregate distribution of action i = 1, 2 as a function of the state of the world ω can be expressed as

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

where E_q denotes the expectation over the possible realizations of q. The right-hand side of (1) is the empirical frequency of action a = i that would result in the long run if agents kept behaving according to the strategy σ .⁷ Note 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.⁸

In an Analogy Based Expectation Equilibrium with payoff-relevant analogy partitions, every agent t assumes that when the state is ω other agents choose action a=i with probability $\overline{\sigma}(a=i|\omega)$, that is, the probability that matches the aggregate distribution of the action. Agents also assume that these behaviors are randomized independently across periods, and that the signal s_t (and its precision q_t) is independent of the actions of his predecessors (conditional on the state of the world). He chooses a best-response accordingly:

⁷Observe the endogenous weight of the behavior at (h_t, s_t, q) in this expression, which reflects the fact that according to σ various histories have different frequencies of visit.

⁸This is the reason we find it technically convenient to present the model first with a finite number T of agents.

⁹Compared 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.

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

We think of 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, that is, that the conjecture about $\overline{\sigma}(a=i|\omega)$ matches the empirical frequency as defined in the right-hand side of (1), 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 (as we have already informally discussed above).

2.2 Equibrium Analysis

An ABEE in our environment is fully characterized by the aggregate distributions of actions in each state of the world ω , that is, $\beta_T(\omega) := \overline{\sigma}(a = \omega | \omega)$, where we make it explicit that the distribution depends on T. Given the symmetry of our game, we further restrict attention to ABEE in which the aggregate distribution of action $a = \omega$ in state ω is the same for $\omega = 0$ and $\omega = 1$, and we denote it by $\beta_T^* \equiv \beta_T^*(0) = \beta_T^*(1)$.

Given β_T^* , it is readily verified that agent t in an ABEE chooses his best response on the basis of the following likelihood ratio:

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

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). More precisely, if this likelihood ratio is greater than 1, then he finds it optimal to choose action 1, if it is lower than 1, he chooses action 0 (otherwise he is indifferent between the two actions).

As we will show, in an ABEE, it is always the case that $\beta_T^* > \frac{1}{2}$. Agents' strategies are then derived as follows. Agent 1 obviously follows his signal, whatever its precision, given that he has no predecessors and thus no other information on which to base his decision. The decision of agent t

to follow his signal or not depends on the following trade-off. Given that $\beta_T^* > \frac{1}{2}$, in the absence of any private signal, this agent would follow the action chosen by the majority of predecessors. With a private signal, this agent would, of course, continue to follow the majority if the private signal agreed with the choice of the majority; and he may consider following his own signal against the majority if the private signal were sufficiently precise. The exact cut-off precision is determined by the value of q_t for which the likelihood ratio as defined in (2) is equal to 1.

In the next proposition we show that an equilibrium exists, and that any symmetric equilibrium is of the form just described. We also derive important asymptotic properties of the equilibrium as the number T of periods grows large.

Proposition 1 a) For any T, there exists a symmetric ABEE with $\beta_T^*(0) = \beta_T^*(1) \equiv \beta_T^* \in (1/2, 1)$. There exists no symmetric ABEE with $\beta_T^* \leq 1/2$.

b) There exists no sequence $\beta_T^*(\omega)$ of corresponding ABEE such that as $T \to \infty$, $\beta_T^*(\omega) \longrightarrow 1$ for $\omega = 0$ and $\omega = 1$.

The argument for the existence of a symmetric ABEE is standard in this finite environment (Jehiel, 2005; Jehiel and Koessler, 2008). In the proof we simply invoke the intermediate value theorem to show the existence of a fixed point. The fixed point cannot be $\beta_T^* < 1/2$, as it would lead agents either to go against the majority (thereby leading to a long run frequency of $\frac{1}{2}$) or to follow their own signal, implying that the overall frequency of action $a = \omega$ in state ω would be no less than 1/2, which contradicts the premise that $\beta_T^* < 1/2$.

As for the asymptotic properties of our economy, as stated in the second part of the proposition, when T grows large, the actions cannot settle on the correct one with probability close to 1. In the Appendix we show that assuming that $\beta_T^*(\omega)$ converges to 1 as T goes to ∞ implies that the probability of taking the incorrect action is bounded away from 0, thereby leading to a contradiction. Intuitively, if the aggregate distribution of the correct action converged to 1, it would mean that in an ABEE all agents except the first one would disregard their own signal and follow the decision that was most popular among the predecessors: that is, the first agent would follow his own signal and all subsequent agents would imitate him. 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 assumed

convergence to efficiency. 10

In order to understand more concretely the asymptotic properties of our social learning problem, we have simulated the equilibrium value of β_T for various T, assuming the q_t 's are uniformly distributed on $[\frac{1}{2}, 1]$.¹¹ 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.

2.3 Discussion

It is instructive to compare our result to that of a model with 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 long-lasting effect, impeding complete learning. One may wonder what causes these differences. 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 updates his belief knowing that the previous agent had received a very high precision signal. Thus, this agent may decide to go against the majority even upon receiving a low precision, as long as the signal contradicts the majority. In other words, 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 (the "overturning principle," as defined by

¹⁰Such an intuition is somehow reminiscent of the Grossman-Stiglitz paradox (Grossman and Stiglitz, 1980): our economy cannot be informationally efficient since otherwise all agents but the first one would neglect their signals, contradicting the efficiency.

¹¹To 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 .

Smith and Sørensen, 2000). In our model, when an agent receives a very precise signal contradicting the consensus, he goes against the majority as in the standard case, but the following agent does not do so unless he himself receives a very precise signal. The reason is that, by considering the aggregate distributions only, he misses the inference that if someone went against a strong majority it must be that he received a very precise signal. In other words, in an ABEE, even though private signals can have unbounded precision, the observed actions always have a bounded informational content, thus not being able to guarantee complete learning.

Ours is a model of bounded rationality. A different approach would be to maintain the rationality assumption and study whether similar insights could be derived by changing some other assumptions, e.g., on what agents observe. Smith 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. Callander and Hörner (2009) propose a model in which agents do not know the order in which previous actions have been made. With imperfect observability, the inference process is typically much harder than with perfect observability. This is why, for example, Callander and Hörner (2009) restrict attention to the case in which signals only have two possible precisions. One interesting result they obtain is that it may be optimal for an agent to neglect his own private information and follow the minority (rather than the majority). In their model, they observe that learning is asymptotically efficient, since the herd on the minority action is eventually broken.¹² Smith and Sørensen also obtain asymptotic efficiency results under their finite sampling assumptions.

These results are in sharp contrast with ours, which are obtained under fairly general specification of the distribution of precisions. In our set-up, agents never follow the minority (remember that $\beta_T^* > 1/2$), and asymptotic inefficiencies necessarily arise. Besides, our analysis does not depend on whether agents observe the order of the previous actions or not, since their behavior in an ABEE depends only on the number of various actions previously observed (not their order). Even though we did not consider the possibility of finite sampling, it is not hard to see that we would a fortiori obtain asymptotic inefficiencies with finite sampling.

So far we have analyzed the case in which agents receive signals with varying precision. One may wonder what happens in the canonical model with signals of deterministic and equal precision (i.e., $q_t = \overline{q} \in (0.5, 1)$ for all t). In the standard Perfect Bayesian Equilibrium, first considered

 $^{^{12}}$ Callander and Hörner (2009) consider the case in which one signal precision is 1/2 and the other is 1.

by Bikhchandani et al. (1992), agents follow their private signals, unless their predecessors having chosen one option outnumber those in favor of the other by at least two. ¹³ This is the case in which an informational cascade (i.e., a situation in which agents neglect their private information) arises. In our set-up, with agents making inferences according to the payoff-relevant analogy reasoning, instead, 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. 4 While the strategies in the standard equilibrium and in the ABEE are similar, still there is an interesting difference. In the standard equilibrium, when the cascade starts, no information is aggregated. As a result, after a cascade of one million people, agents have the same beliefs as after a cascade of two people (i.e., the public belief never exceeds the belief arising from having two signals in favor of the chosen state), a quite unsettling conclusion. In the ABEE, as more and more agents choose the same action, agents' beliefs are updated every time in favour of the chosen action. The likelihood in favor of the chosen action, similar to that in 2, keeps increasing. Eventually, agents put weight 1 on the state corresponding to the action chosen by the first agent — not necessarily the correct one. In the standard set-up, since it is built on very limited information, a cascade is fragile. If after a cascade of one million people there is a small shock (e.g., public information), in the canonical model, the cascade is broken. In our set-up, instead, the herd becomes less and less fragile, and bigger shocks would be required to make successors switch to the alternative action. This is consistent with experimental findings, according to which longer cascades are less fragile (Kübler and Weizsäcker, 2005).

3 Social Learning with Continuous Action Space

We now turn to illustrate which insights our bounded rationality approach offers when agents can choose their action in a continuous space, in a model similar to Lee's (1993). Specifically, we assume that agent t takes an action a_t in the space [0,1]. The agent's payoff function is quadratic and equal to $-(\omega - a_t)^2$. This ensures that agents want to take an action as close as possible to what they believe the state is (given what they observe and what they infer from it). As in Section 2, we assume that agent t observes previous actions $h_t = (a_1, a_{t-1})$ together with a binary signal

¹³This 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.

 $^{^{14}}$ We refer the reader to the working paper version of the article for a complete analysis.

 s_t distributed as follows:

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

where q_t is the precision of signal s_t . As before, conditional on the state of the world, the signals are independent over time. In contrast with Section 2, however, we assume that each signal has a given precision, lower than 1 (i.e., the distribution of beliefs is bounded), that may deterministically vary from one period to another.

For convenience, we denote the ratio $\frac{\Pr(s_i|\omega=1)}{\Pr(s_i|\omega=0)} = \left(\frac{q_i}{1-q_i}\right)^{2s_i-1}$ by $m(s_i)$. As will prove useful in the following analysis, we make the following (genericity) assumption:¹⁵

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

Such an assumption means that receiving k_i signals s_i of precision q_i for i = 1, ...T must result in a different belief than that resulting from receiving k'_i signals s'_i of precision q_i for i = 1, ...T unless for all i, $k_i = k'_i$ and $s_i = s'_i$.

Given the information (h_t, s_t) , agent t chooses a_t to maximize the expected payoff $E^{SUBJ}[-(\omega - a_t)^2|h_t, s_t]$. That is, he chooses $a_t^* = E^{SUBJ}[\omega|h_t, s_t]$. (We use the superscript "SUBJ" in the expectation to emphasize that each agent forms a subjective expectation.) We can restrict attention to pure strategies, given that $E^{SUBJ}[\omega|h_t, s_t]$ reduces to the choice of a single action.

3.1 Equilibrium Analysis

An agent's strategy σ_t maps (h_t, s_t) into an action, that is, $\sigma_t : H_t \times \{0, 1\} \to [0, 1]$. Similarly to the previous section, $\sigma_t(a|h_t, s_t)$ 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, we will consider equilibria in pure strategies, $\sigma_t(a|h_t, s_t)$ will either be equal to 0 or to 1. Given a particular strategy profile $\sigma = (\sigma_1, \sigma_2, ... \sigma_T)$,

¹⁵Such an assumption is met for almost every $(q_1, ... q_T)$, whatever the Lebesgue measure on $(\frac{1}{2}, 1)^T$. (Note that \mathbb{Z} denotes the set of integers.) Essentially, the assumption means that, whatever the weight given to two or more signals, they never cancel out, giving a posterior likelihood ratio of 1. One obvious case in which the assumption is not satisfied is when all signals have the same precision. In that case, for instance, a signal 1 and a signal 0 taken with equal weight (i.e., $n_i = 1$) would offset each other. Of course, one can choose precisions arbitrarily close to each other and still satisfy A1.

¹⁶Obviously agent t also knows the (deterministic) precision q_t of his signal. We do no explicitly indicate this in the information set. Whether the agent also knows the other signals' precisions or not is instead immaterial, as will become clear in the analysis.

the probability $\mu^{\sigma}(h_t, s_t|\omega)$ that history h_t is realized and s_t is the signal at t when ω is the state of the world is strictly positive only for finitely many (h_t, s_t) (since s_t takes values in $\{0, 1\}$ and the strategies σ_t 's are pure). As a result, analogously to the previous analysis, we can define the aggregate distribution of actions as a function of the state of the world ω as

$$\overline{\sigma}(a|\omega) = \frac{\sum_{t=1}^{T} \sum_{h_t, s_t} \sigma_t(a|h_t, s_t) \mu^{\sigma}(h_t, s_t|\omega)}{T}.$$
(3)

Given this aggregate distribution of actions, an ABEE is defined analogously to the previous section. Based on the history of actions $h_t = (a_1, a_{t-1})$ and the signal s_t at date t, the agent forms the subjective likelihood ratio

$$\overline{LR}(h_t, s_t) = \prod_{i=1}^{t-1} \frac{\overline{\sigma}(a_i \mid \omega = 1)}{\overline{\sigma}(a_i \mid \omega = 0)} m(s_t), \tag{4}$$

and chooses action a_t so that

$$\frac{a_t}{1 - a_t} = \overline{LR}(h_t, s_t). \tag{5}$$

We now start the construction of an ABEE. Let $\alpha(s_1, s_2, ..., s_k)$ denote the equilibrium action taken after the sequence of signal realizations $\{s_1, s_2, ..., s_k\}$. We conjecture that the agents' strategies are such that, in equilibrium, for any two different sequences of signal realizations, $\{s_1, s_2, ..., s_k\} \neq \{s'_1, s'_2, ..., s'_l\}$, an agent chooses two distinct actions in [0, 1], that is, $\alpha(s_1, s_2, ..., s_k) \neq \alpha(s'_1, s'_2, ..., s'_l)$. 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, ..., k, and the case in which $k \neq l$. Given this conjecture, we construct an equilibrium and then we verify that, because of A1, it satisfies the conjecture.

Specifically, consider a sequence of signals $\{s_1,...s_{t-1}\}$ and the corresponding actions $a_1 = \alpha(s_1),..., a_{t-1} = \alpha(s_1,...s_{t-1})$. Given our conjecture that all α 's are different, it is readily verified that a given action is observed only at one time and after one sequence of signals. Thus, $\frac{\overline{\sigma}(a_i|\omega=1)}{\overline{\sigma}(a_i|\omega=0)} = \frac{\Pr(s_1,...,s_i|\omega=1)}{\Pr(s_1,...,s_i|\omega=0)}$, which (given the independence of s_k 's conditional on ω) simplifies into

$$\frac{\overline{\sigma}(a_i \mid \omega = 1)}{\overline{\sigma}(a_i \mid \omega = 0)} = \prod_{k=1}^{i} m(s_k).$$
(6)

A simple rewriting of $\overline{LR}(h_t, s_t)$ yields¹⁷

$$\overline{LR}(h_t, s_t) = \prod_{i=1}^{t-1} m(s_i)^{t-i} m(s_t).$$
 (7)

¹⁷We use the convention that $\Pi_{i=1}^0 x^i$ is equal to 1.

Note that 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 $\alpha(s_1, s_2, ..., s_t)$ such that

$$\frac{\alpha(s_1, s_2, ..., s_t)}{1 - \alpha(s_1, s_2, ..., s_t)} = \prod_{i=1}^{t-1} m(s_i)^{t-i} m(s_t).$$
(8)

In other words, agent t chooses his action as if he had observed signal s_1 t-1 times, signal s_2 t-2 times,..., signal s_i t-i times. The reason is that, in this construction, a signal s_i affects the aggregate distribution of actions at every time $t \ge i$. This overweight of early signals is the essence of the bias induced by the (boundedly rational) reasoning the agents use.

To ensure that the construction of this equilibrum is correct, we still need to check our conjecture that the actions $\alpha(s_1, s_2...s_t)$'s, as defined by (8), are all different. This is clearly implied by our genericity assumption (A1), thereby allowing us to conclude:¹⁸

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

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 $\{a_1, a_2, ..., a_{t-1}\}$ such that $a_1 = \alpha(s_1)$, $a_2 = \alpha(s_1, s_2)$, ..., $a_{t-1} = \alpha(s_1, s_2...s_{t-1})$ and by the signal s_t , by identifying the strategy $\sigma_t(a|\{a_1, a_2, ..., a_{t-1}\}, s_t)$ with the choice of the action $a = \alpha(s_1, s_2, ..., s_t)$ with probability 1. For other histories, the strategy is not specified, but this is irrelevant for the analysis of the equilibrium path.^{19'} ²⁰

¹⁹Obviously, the construction 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. The premise is that they have learned $\overline{\sigma}(a|\omega)$ for the various a and ω through past observations, and choose their actions according to the likelihood ratio \overline{LR} as defined in (4).

²⁰¿From 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 of actions given the state of the world, which does not require knowledge of others' precisions.

¹⁸We have constructed an ABEE under the genericity assumption A1. A natural question is whether there are other ABEE and what happens if assumption A1 is violated. Since these are mainly technical issues, we refer the reader to the working paper version of the article. Here we only mention our main findings: generically, the ABEE constructed in the previous section is unique. Moreover, when A1 is violated, like in the case of signals of equal precision, there is no guarantee that an ABEE exists. 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. The reason why an ABEE may not exist when A1 is not satisfied is that the action space is continuous.

It is worth contrasting the result in Proposition 2 with what happens in the case in which agents are fully rational. As we know from Lee (1993), in the case of fully rational agents, in the Perfect Bayesian Equilibrium (PBE), after a sequence of signals $\{s_1, s_2, ..., s_t\}$, agent t chooses action $a^{PBE} = \alpha^{PBE}(s_1, s_2, ..., s_t)$ such that $\frac{a_t^{PBE}}{1-a_t^{PBE}} = \frac{\alpha^{PBE}(s_1, s_2, ..., s_t)}{1-\alpha^{PBE}(s_1, s_2, ..., s_t)} = \prod_{i=1}^t \left(\frac{q_i}{1-q_i}\right)^{2s_i-1}$. In the PBE, agents perfectly infer the signals observed by their predecessors from their 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.²¹ The comparison between the expressions obtained for the ABEE and for the PBE 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. In the ABEE, history matters, in that a signal (and thus an action) early in the game has more impact than later signals (actions).

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 the 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 the 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 Proposition 2 and from the discussion on the PBE. 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

²¹Also 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 this 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.

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,...) 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 $\omega = 1$, $B_t^{SUBJ} \stackrel{a.s.}{\to} 1$, and $B_t \stackrel{a.s.}{\to} 1$ (and, similarly, when $\omega = 0$, $B_t^{SUBJ} \stackrel{a.s.}{\to} 0$, and $B_t \stackrel{a.s.}{\to} 0$). Moreover, in both cases, convergence obtains exponentially fast, that is, $\Pr(|X_t - q| > \varepsilon) \le 2e^{-2t\varepsilon^2}$ and $\Pr(|Z_t - q| > \varepsilon) \le 2e^{-tK(\varepsilon)\varepsilon^2}$, 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.²² 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

²²The result would hold a fortiori if we considered unbounded beliefs.

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 (the same number of repetitions is also used for the other simulations presented below). 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.

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.

3.2 Discussion

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.

It is also worth noting that our result on convergence, similar to that of the PBE, is instead in sharp contrast with what happens if agents 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, 2010), the weight of the first signal is approximately equal to the sum of the weights of all other signals. Given the overwhelming weight of the (possibly wrong) early signals — growing at an exponential rate with t, as opposed to a linear rate in our model — there must be asymptotic inefficiencies.

The analysis we have presented can be extended in several directions. One can consider the case in which agents are not aware of the order of moves; the case in which agents can only observe some (e.g., a sample of, or the immediate) predecessor(s); the case in which more than one action is taken at the same time. Our finding that early signals have more weight continues to hold in these extensions. One notable exception is when agents only observe their immediate predecessor, in which case ABEE and PBE offer the same result. In fact, according to the PBE, different assumptions on the observability of past actions do not change the agents' behavior: observing the entire sequence or a number $n \geq 1$ of immediate predecessors leads anyway to efficient learning. In the ABEE, instead, the longer the sequence of immediate predecessors an agent can observe, the higher the bias in favor of early signals. We refer the reader to the working paper version of the article for a comprehensive analysis of these cases.

4 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 unbounded signal precision and that of the continuous action space) where the previously available results required more investigation. 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. 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.

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, 2010). 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 the continuous action space and

for the discrete action space with varying precisions of signals) are needed to test the predictions of our model. This is left for future research.

References

- [1] Acemoglu, Daron, Asuman E. Ozdaglar, and Ali ParandehGheibi. 2009. "Spread of (Mis)Information in Social Networks." MIT WP 09-15.
- [2] Bala, Venkatesh, and Sanjeev Goyal. 1998. "Learning form Neighbours." Review of Economic Studies, 65, 595-621.
- [3] Banerjee, Abhijit V. 1992. "A Simple Model of Herd Behavior." Quarterly Journal of Economics, 107, 797-817
- [4] **Bergemann, Dirk, and Stephen Morris.** 2005. "Robust Mechanism Design." *Econometrica*, 73(6), 1771–1813.
- [5] Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch. 1992. "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades." *Journal of Political Econ*omy, 110, 992-1026.
- [6] Bohren, Aislinn J. 2009. "Information-Processing Uncertainty in Social Learning." Mimeo, University of San Diego.
- [7] Callander Steven, and Johannes Hörner. 2009. "The Wisdom of the Minority." Journal of Economic Theory, 144, 1421-1439.
- [8] **Dekel, Eddie, Drew Fudenberg, and David K. Levine.** 2004. "Learning to Play Bayesian Games." *Games and Economic Behavior*, 46, 282–303.
- [9] De Marzo, Peter M., Dimitri Vayanos, and Jeffrey Zwiebel. 2003. "Persuasion Bias, Social Influence, and Unidimensional Opinions." Quarterly Journal of Economics, 118, 909-968.
- [10] Ellison, Glenn, and Drew Fudenberg. 1993. "Rules of Thumb for Social Learning." Journal of Political Economy, 111, 612-643.
- [11] Eyster, Erik, and Matthew Rabin. 2005. "Cursed Equilibrium." *Econometrica*, 73(5), 1623-1672.
- [12] Eyster, Erik, and Matthew Rabin. 2010. "Naive Herding in Rich-Information Settings."

 American Economic Journal: Microeconomics, 2, 221-243
- [13] Grossman, Sanford J., and Joseph E. Stiglitz. 1980. "On the Impossibility of Informa-

- tionally Efficient Markets." American Economic Review, 70, 393–407.
- [14] Jehiel, Philippe. 2005. "Analogy-Based Expectation Equilibrium." Journal of Economic Theory, 123, 81-104.
- [15] **Jehiel, Philippe, and Frédéric Koessler.** 2008. "Revisiting Games of Incomplete Information with Analogy-Based Expectations." *Games and Economic Behavior*, 62, 533-557.
- [16] Kübler, Dorothea, and Georg Weizsäcker. 2004. "Limited Depth of Reasoning and Failure of Cascade Formation in the Laboratory." Review of Economic Studies, 71, pp. 425-441.
- [17] Kübler, Dorothea, and Georg Weizsäcker. 2005. "Are longer cascades more stable?" Journal of the European Economic Association, 3(2-3), 330-339.
- [18] Lee, In Ho. 1993. "On the Convergence of Informational Cascades." Journal of Economic Theory, 61, 395-411.
- [19] Shiryaev, Albert N. 1996. Probability, Springer-Verlag, New York.
- [20] Smith, Lones, and Peter Sørensen. 2000. "Pathological Outcomes of Observational Learning." Econometrica, 68, pp. 371-398.
- [21] Smith, Lones, and Peter Sørensen. 2008. "Rational Social Learning with Random Sampling." *Mimeo*, University of Michigan.
- [22] Weizsäcker, Georg. 2010. "Do We Follow Others when We Should? A simple test of rational expectations." American Economic Review, 100(5), 2340-2360.

Appendix: Proofs

Proof of Proposition 1 (Point "a")

First, we prove 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 $\phi: [0.5,1] \to [0.5,1]$ that, for a given β_T , gives the aggregate distribution of actions conditional on a state of the world (in the best response to β_T), $\phi(\beta_T)$. 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 $a=\omega$ in state ω is $\phi(1/2)=E(q_t)>1/2$. 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 $a=\omega$ in state ω is again $\phi(1)=E(q_t)<1$. Finally, notice that, at any time t, the probability that an agent receiving signal ω chooses $a=\omega$ is 1 if $|a=\omega|\geq |a=1-\omega|$ and is $1-F\left(\frac{(1-\beta_T)^{|a=\omega|-|a=1-\omega|}}{(1-\beta_T)^{|a=\omega|-|a=1-\omega|}+\beta_T^{|a=\omega|-|a=1-\omega|}}\right)$ otherwise. Since q_t is distributed according to a continuous density function, $\phi(\beta_T)$ is a continuous function. By the intermediate value theorem, there exists a $\beta_T \in (1/2,1)$ such that $\phi(\beta_T)=\beta_T$.

Now we prove that there does not exist an ABEE with $\beta_T < 1/2$.

By contradiction, suppose $\beta_T < 1/2$. Without loss of generality, let us consider the case in which $\omega = 1$. Let us define $\widetilde{\beta}_t$ as the aggregate distribution of actions at time t. Note that $\beta_T = \frac{1}{T} \sum_{t=1}^T \widetilde{\beta}_t$.

Observation: an agent who observes |a=1|-|a=0|=n>0, chooses action 1 with probability $Eq_t\left(1-F\left(\frac{(1-\beta_T)^n}{\beta_T^n+(1-\beta_T)^n}\right)\right)$ and action 0 with complementary probability $(1-Eq_t)+Eq_tF\left(\frac{(1-\beta_T)^n}{\beta_T^n+(1-\beta_T)^n}\right)$. Similarly, an agent who observes |a=1|-|a=0|=-n<0, chooses action 0 with probability $(1-Eq_t)\left(1-F\left(\frac{(1-\beta_T)^n}{\beta_T^n+(1-\beta_T)^n}\right)\right)$ and action 1 with complementary probability $Eq_t+(1-Eq_t)F\left(\frac{(1-\beta_T)^n}{\beta_T^n+(1-\beta_T)^n}\right)$. (Of course, an agent who observes |a=1|-|a=0|=0, chooses action 1 with probability Eq_t and action 0 with complementary probability $1-Eq_t$.

Now, consider any time t. Observe that for odd t, the total number of actions 1 minus the total number of actions 0 until time t included (let us denote it by θ_t), can take the following values ± 1 , ± 3 ,..., $\pm (t-1)$. For even t, the values are 0, ± 2 ,..., $\pm (t-1)$. Note that for each history of actions $(a_1, a_2, a_3, ..., a_t)$ leading to $\theta_t = -m < 0$, there exists the history of actions $(1 - a_1, 1 - a_2, 1 - a_3, ..., 1 - a_t)$ leading to $\theta_t = m > 0$. By using the expressions in the above Observation, it is immediate to verify that the probability of the latter history is always higher than the probability of the former. (The only remaining histories are such that $\theta_t = 0$.) Therefore, at each time t, $\beta_t > 1/2$. It is then immediate to conclude that $\beta_T > 1/2$, a contradiction.

Proof of Proposition 1 (Point "b")

Without loss of generality, suppose the state of the world is $\omega = 0$. Suppose $\beta_T^*(0) \to \widehat{\beta}$ and

 $\beta_T^*(1) \to \widehat{\beta}$, for $\widehat{\beta} = 1 - \varepsilon$, and for a small ε .

First, note that if $\beta_T^*(0) \to \widehat{\beta}$, then an agent receiving a signal of precision $q_t < \widehat{\beta}$ always follows the majority. Consider the least favorable case in which the majority is by only one. Clearly, the likelihood ratio is $\left(\frac{\widehat{\beta}}{1-\widehat{\beta}}\right)^I \left(\frac{q_t}{1-q_t}\right)^{2s_t-1}$, where I takes value 1 if the majority action is 1 and -1 if it is 0. This likelihood ratio is higher or lower than 1 (and so the action chosen by the agent is either 1 or 0) depending only on the choice of the predecessors, since $q_t < \widehat{\beta}$. The argument holds 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 $s_2=1$, 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 $\varepsilon < \overline{\varepsilon}$ and some $\overline{\varepsilon} > 0$.

Similarly, if the third agent receives the signal $s_2=1$, he chooses $a_3=1$, and if he observes $s_3=0$, he chooses $a_3=1$ with a probability equal to $F\left(\frac{\widehat{\beta}^2}{\widehat{\beta}^2+(1-\widehat{\beta})^2}\right)$. Note that $\frac{\widehat{\beta}^2}{\widehat{\beta}^2+(1-\widehat{\beta})^2}>\widehat{\beta}^2$ and so $F\left(\frac{\widehat{\beta}^2}{\widehat{\beta}^2+(1-\widehat{\beta})^2}\right)>F(\widehat{\beta}^2)$. Using again Taylor's expansion, it is easy to prove that $F(\widehat{\beta}^2)$ is greater than $1-4f(1)\varepsilon^2$ for $\varepsilon<\overline{\varepsilon}$ and some $\overline{\varepsilon}>0$. 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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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:

$$\sum_{i=1}^{t} \frac{1}{n} \left(\log((1-q) + qe^{\lambda i}) - \lambda ai \right) \approx t \int_{0}^{1} \left(\log((1-q) + qe^{mx}) - amx \right) dx.$$

Hence, we can write

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

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

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

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

 $a \log \frac{a}{q} + (1-a) \log \frac{(1-a)}{(1-q)} := H(a) \ge 0$. 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,

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

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

Therefore, we can conclude that

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

Analogous arguments show that

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

Finally, we can conclude that

$$\Pr(|Z_t - q| \ge \varepsilon) \le 2e^{-nm^*\varepsilon^2}.$$

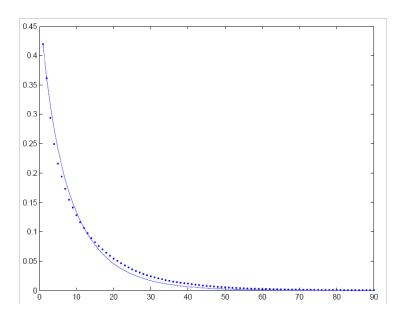


Figure 1: Average distance of the public belief from the fundamental value. The solid line refers to the PBE. The dotted line refers to the ABEE.

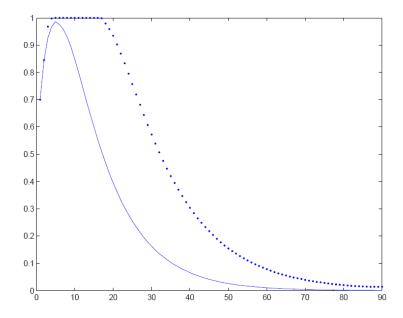


Figure 2: Average distance of the public belief from the fundamental value after 5 incorrect signals. The solid line refers to the PBE. The dotted line refers to the ABEE.