

Bayesian Social Learning with Local Interactions

1 2 3
, ,

Antonio Guarino
University College London

Antonella Ianni
University of Southampton

September 8, 2010

¹The authors wish to thank the Editor and two anonymous referees for constructive comments and suggestions, and S. Camussi and C. Zedan for valuable research assistance.

²**Keywords:** social learning; Bayesian learning; local informational externalities; path dependence; consensus; clustering; convergence rates.

³ADDRESS FOR CORRESPONDENCE: A. Ianni, Economic Division, School of Social Sciences, University of Southampton, Southampton SO17 1BJ, UK

Abstract

We study social learning in a large population of agents who only observe the actions taken by their neighbours. Agents have to choose one, out of two, reversible actions, each optimal in one, out of two, unknown states of the world. Each agent chooses rationally, on the basis of private information and of the observation of his neighbours' actions. Agents can repeatedly update their choices at revision opportunities that they receive in a random sequential order. We show that if agents receive equally informative signals and observe both neighbours, then actions converge exponentially fast to a configuration where some agents are permanently wrong. In contrast, if agents are unequally informed (in that some agents receive a perfectly informative signal and others are uninformed) and observe one neighbour only, then everyone will eventually choose the correct action. Convergence, however, obtains very slowly, at rate \sqrt{t} .

1 Introduction

In many economic and social situations, we make our decisions after observing the choices of others. We can learn from such choices, since they can reveal information that others hold. Observing others can help us to form a more precise evaluation and make a better choice. Of course, observing others can also lead to conformism, in so far as we decide to just follow the crowd.

The process of learning from others has first been studied by Bikhchandani *et al.* [6], and Banerjee [4]. These papers focused on a simple case in which agents, endowed with some private information, act sequentially and make one irreversible choice, after observing the entire history of actions taken by their predecessors. This set up was, of course, very convenient for the analysis, but, at the same time, quite restrictive.

One of the features that restrict the applicability considerably is that every agent can observe the whole set of choices made by others. In most cases, we do not observe everyone's decisions but only those made by agents that we know, like our friends or neighbours. We know, for example, which restaurant our friends go to, which bank they use, which car they bought. We may try to infer information from their actions. But this inference is, of course, very complicated. While we observe their actions, and can take these into account to make our decision, they themselves could have gone through a similar process when it was their turn to make a choice. They could have observed our previous decisions and those of their friends and neighbours, that maybe we do not observe. And their neighbours, in turn, could have observed the decisions of others, and so on. Clearly, when we take all this into account, we realize that the process of social learning is quite an intricate and complex phenomenon.

The purpose of our paper is to shed some light on social learning when agents can only observe the actions taken by their neighbours. The seminal papers by Bikhchandani *et al.* [6] and Banerjee [4] illustrated a striking phenomenon: in their set up, eventually agents decide to disregard their private information and just conform to the prevailing action chosen by their predecessors. Conformism prevails, and it may well be that the entire population settles on the wrong action. Beliefs may never converge to the truth, despite the population is only formed of rational agents. In situations in which agents can only observe their neighbours' decisions, should we still expect uniformity of behavior? Can we still expect agents to neglect their private information? Will the population as a whole eventually learn the right decision? And in the positive case, will this convergence process be slow or fast?

When the agents' observation is limited to their neighbours' choices, the amount of information that they receive is limited and the possibility of learning seems reduced. On the other hand, the way in which this information is disseminated may be more efficient, since agents may rely more on their private information, and feed this into the social learning process by their choice of actions. Therefore, the social learning process under local interactions can differ in many ways from the one studied in the

canonical models.

To address these issues, we build a simple model that departs from the canonical one in various ways. First, as we said, we assume that agents can only observe the behavior of their neighbours, that is, a subset of other individuals who live close by, on an appropriately characterized spatial structure. Second, we let each agent revise his original decision repeatedly and postulate that updating opportunities occur in a random sequential order. Third, we assume that agents do not know the time of their decision and have no recall of past (own and neighbours') experiences. This last assumption makes the inference problem easier but far from trivial. We believe that not only it helps to make the model tractable (assuming perfect memory would make the inference problem overwhelmingly complicated), but it is also plausible in the analysis of the dynamics of choices taken by a large population.

In this set up, we address two issues. The first concerns the social learning process in terms of its asymptotic properties. Starting from an initial random configuration of beliefs in the population, we ask whether social learning is complete, in the sense that beliefs converge to the truth, or at least adequate, in the sense that all agents choose the action that is optimal given the state of the world. The second issue concerns the speed of learning. In our view, this issue is particularly relevant as the distinction between slow convergence to the truth and fast convergence to the false is not an obvious one, neither in practical terms, nor in terms of efficiency. We are able to provide an analytical solution for the speed of convergence of actions, and, by pursuing a space-time analysis (i.e., by relating the two dimensions, time and space, over which our process is defined), for the process of cluster formation.

Our main results are the following. When agents are equally informed (i.e., endowed with a signal of equal precision) learning is not complete and the process of actions converges exponentially fast to a configuration where somebody is permanently wrong. When agents are unequally informed, in the sense that some receive a fully informative signal and some receive a completely uninformative signal, social learning is adequate, that is, there is convergence to a state in which everybody chooses the correct action. Convergence, however, obtains very slowly.

The paper is organized as follows. Section 2 describes the model. Section 3 contains the main results. Section 4 relates our work to the existing literature. Finally, Section 5 concludes and Appendices A and B contain the proofs.

2 The Economy

We consider a set $\Omega = \{0, 1\}$ of possible states of nature. The two states are equally likely. In the economy there is a set X of countably many agents. Each agent $x \in X$ has to choose an action in the action space $A = \{0, 1\}$. Time t runs continuously ($t \geq 0$). Each agent x makes a decision at time $t = 0$ and then may be called to revise it more than once. In particular, each agent may have to choose a new action at a random exponential time, with mean 1. In any small time interval, at most

one agent can reassess his decision, and every agent is equally likely to receive an updating opportunity. We denote the action of agent x at time t by $\eta_t(x) \in A$. Let $\{\tau_{x_l}\}$ be the sequence of times when x receives an updating opportunity, with $\tau_{x_0} \equiv 0$ and $l = 0, 1, 2, \dots$. Agent x 's choice at t is equal to the decision he made at his last updating opportunity, that is, $\eta_t(x) = \eta_{\tau_{x_l}}(x)$ for $\tau_{x_l} \leq t < \tau_{x_{l+1}}$.

Each agent gathers information on the state of the world in two ways. At time $t = 0$ he observes a private symmetric binary signal on the realized state of nature. We denote the signal observed by agent x by $\theta^\sigma(x) : \Omega \rightarrow \{0, 1\}$, where the index $\sigma \in \{h, l\}$ refers to the precision of the signal. An agent can receive a signal of high precision $q^h \equiv \Pr[\theta^h(x) = 1 \mid \omega = 1] = \Pr[\theta^h(x) = 0 \mid \omega = 0] \in (0.5, 1]$ or a signal of low precision $q^l \equiv \Pr[\theta^l(x) = 1 \mid \omega = 1] = \Pr[\theta^l(x) = 0 \mid \omega = 0] \in [0.5, q^h]$. Note that, conditional on a state of nature, the signals that agents receive are independently distributed. Note also that q^l is weakly lower than q^h . We will analyze a case in which the precisions are identical and one in which one precision is strictly lower than the other.

Having observed the signal at time 0, each agent x makes his first choice, $\eta_0(x)$. When the agent receives another opportunity to take a decision (i.e., to revise the choice previously made), he observes the decisions taken by a subset of other agents in the population. This is the second way in which he gathers information on the state of nature. We provide each agent with a spatial location on a 1-dimensional lattice Z^1 (an *address*), and assume that he can only interact with the set of agents who live in his vicinity. Formally, we take $X \subseteq Z^1$ and define the set of x 's nearest neighbours as $N(x) = \{y : \|y - x\| = 1\}$, that is, the set of two agents who live at Euclidean distance 1 from agent x . We denote these two agents by $x - 1$ and $x + 1$, and the information set upon which agent x takes a decision at time t by $I_t(x)$. In order for agents to be able to fully draw inference upon the observation of their neighbours' action, one should endow each agent with a very rich information set, including the history of actions chosen in their neighbourhood and the exact order in which updating opportunities have been assigned until that point. Although we assume that agents are able to perform Bayesian updating, we take the view that these requirements are unreasonable in the set up of a large population of agents. On these grounds, we assume that agents have limited memory, in the sense that, if at time t agent x has to choose an action, his information set is¹:

$$I_t(x) \subseteq \begin{cases} \{\theta^\sigma(x), \sigma\} & \text{for } t = 0, \\ \{\eta_t(x), \eta_t(y), y \in \{x \pm 1\}, \sigma\} & \text{for } t > 0. \end{cases}$$

where $\sigma \in \{h, l\}$. For the entire paper, we keep the assumption that at time $t = 0$ agents observe only their private signal. For $t > 0$, instead, we will consider two cases, one in which the agent observes both neighbours, and one in which he only observes one neighbour (randomly drawn). Note that, because of limited memory,

¹An entirely equivalent specification could be obtained by imposing that agents only use stationary strategies, that is, strategies that do not depend on time, nor on location.

agents revise their decisions every time completely ignoring past history. Note also that the time t does not belong to the information set, that is, the agent does not know the time of his decision, neither he knows how many times he already had an opportunity to revise his choice. On the other hand σ does belong to agent x 's information set, meaning that at each time a decision is to be made, x knows the precision of the signal he received.

Agents' payoff only depends on the realized state of nature ω and on the chosen action. While action 0 always gives a payoff of zero, action 1 gives a payoff of one if the state is 1 and a payoff of minus one if the state is 0. On the basis of the available information, at time t agent x chooses $\eta_t(x)$ to maximize $E[U(\eta_t(x), \omega) | I_t(x)]$ and sticks to this decision until a new updating opportunity arises.

If we denote the belief at time t that $\omega = 1$ by $\pi_t(x) = \Pr[\omega = 1 | I_t(x)]$, then x 's optimal action is

$$\eta_t^*(x) = \begin{cases} 1 & \text{if } \pi_t(x) > 0.5, \\ \{0, 1\} & \text{if } \pi_t(x) = 0.5, \\ 0 & \text{if } \pi_t(x) < 0.5. \end{cases} \quad (1)$$

or, equivalently,

$$\eta_t^*(x) = \begin{cases} 1 & \lambda_t(x) > 0, \\ \{0, 1\} & \lambda_t(x) = 0, \\ 0 & \lambda_t(x) < 0. \end{cases} \quad (2)$$

where $\lambda_t(x) \equiv \log \frac{\Pr[\omega=1|I_t(x)]}{\Pr[\omega=0|I_t(x)]}$ denotes the log-likelihood ratio (LLR).

Essentially, agents choose their best action whenever they have the opportunity to do so. A simple interpretation is that the actual payoff is realized only when at a random time the game is over. Since agents do not know the end of the game, maximizing the payoff every time is in fact their optimal behavior.

We now define the equilibrium in our game.

Definition 1 (Equilibrium) *An equilibrium is a profile of strategies $\{\eta_t^*(x)\}_{x \in X}$ such that, for all $x \in X$, $\eta_t^*(x) \in \arg \max E[U(\eta(x), \omega) | I_t(x)]$.*

Clearly, an equilibrium is absorbing for our learning processes if there exists a time T such that for any $x \in X$ and for any $t > T$, $\eta_t^*(x) = \eta^*(x)$. The dynamics are as follows. Agent x makes a decision at time $t = 0$ and then, whenever he has an updating opportunity, chooses an action according to (1), that is,

$$\eta_t(x) = \eta_{\tau_{x_l}}^*(x) \quad \text{for } \tau_{x_l} \leq t < \tau_{x_{l+1}} \quad \text{and } l = 0, 1, 2, \dots$$

At the beginning of time, ω is realized and each agent x receives a signal $\theta^\sigma(x)$ which determines $\eta_0(x)$. The process then evolves stochastically in continuous time. We refer to the process of social learning as to the dynamic process generated by the collection of all individual actions and we are interested in analyzing its properties. We shall denote the state of the process at time t by $\eta_t \in \{0, 1\}^X$ and we are interested in characterizing its evolution over time and over space.

3 Social Learning

Before proceeding to the analysis, we find it useful to discuss the relation between a canonical model of social learning and a model of social learning with local interactions as ours. Consider the standard model of sequential social learning of Birkhchandani *et al.* [6] and suppose that each agent can directly observe the signals received by his predecessors before making a decision. Suppose, for simplicity, that $q^h = q^l \equiv q \in (0.5, 1)$ and that agents are indexed by $1, 2, 3, \dots, n$. Since the precisions of the signals are identical, we omit the superscript for θ . Then, the LLR for agent n on the basis of his information set $I(n) = \{\{\theta(y)\}_{y \leq n}\}$ is

$$\lambda(n) \equiv 2 \left(\log \frac{q}{1-q} \right) \left(\theta(n) - \frac{1}{2} \right) + 2 \left(\log \frac{q}{1-q} \right) \sum_{y < n} \left(\theta(y) - \frac{1}{2} \right).$$

Suppose, for instance, that the true state of nature is $\omega = 0$. Then, the random variables $\theta(\cdot)$ have mean $1 - q$. This implies that $\lambda(n)$ tends to $-\infty$, as n tends to infinity. In other words, the assessment of the probability that state $\omega = 1$ is true, $\pi(n) \equiv \exp[\lambda(n)] / [1 + \exp[\lambda(n)]]$, will tend to zero exponentially fast, as the number of observations increases. Therefore, in a canonical model of social learning in which agents can observe all signals, learning is *complete* in the sense that beliefs converge to the truth (and, hence, actions to the correct decisions), that is, $\Pr[\lim_{n \rightarrow \infty} \pi(n) = \eta(n) = \omega] = 1$ for all n . In particular, convergence obtains exponentially fast.

Essentially, in the canonical model of social learning, the observability of private information guarantees that the outcome of the learning process is informationally efficient. Actions are clearly not as informative as signals in this setting: potential inefficiencies may arise when agents only observe the actions taken by their predecessors and not their signals. Indeed, given a discrete action space, these inefficiencies may take the extreme form of an informational cascade, in which beliefs do not converge to the truth and the entire population settles on the wrong action.

Things are different if interaction is local, as in our framework. To see this, assume that agents do observe signals (and not actions), but only those of their nearest neighbours. Since signals are noisy and each agent observes three signals only, agents' beliefs, as measured by their LLR, are bounded. As a result, we cannot expect any convergence in beliefs to the truth in this case, and agents will continue choosing the same action (perhaps the incorrect one) indefinitely. Observing actions may in fact make agents better off, as actions taken by one's neighbour may convey information on signals received by that neighbour's neighbours and so on.

We are interested in characterizing the properties of social learning processes with local interactions in terms of the degree of informational efficiency. For the reasons just mentioned, requiring complete learning would be too demanding in our model. A weaker requirement is, instead, that all agents eventually choose the correct action, as in the following definition:

Definition 2 (Adequate Learning) *The social learning process shows adequate learning if*

$$\lim_{t \rightarrow \infty} \Pr[\eta_t(x) = \omega \text{ for all } x \text{ in } X] = 1.$$

In fact, $\lim_{t \rightarrow \infty} \Pr[\eta_t(x) = \omega]$ is the limit measure of agents who choose the action appropriate for the true state of the world and is our measure of how informationally efficient the social learning process is.

We now move to address these issues with reference to our specific models.

3.1 Equally Informed Agents

We start by studying the case in which each agent receives a signal of equal precision² $q^h = q^l \equiv q \in (0.5, 1)$ and observes both his neighbours. At any revising opportunity, agent x 's information set is:

$$I_t(x) = \begin{cases} \{\theta(x), \sigma\} & \text{at } t = 0, \\ \{\eta_t(x), \eta_t(y), y \in \{x \pm 1\}, \sigma\} & \text{at each } t = \tau_x. \end{cases}$$

where σ is dropped for notational convenience. We now show that in this set up in equilibrium extreme inefficiencies arise.

Theorem 3 *If each agent x receives a symmetric binary signal with precision $q \in (0.5, 1)$ and can observe both neighbours, then the process of social learning is not adequate, and $\lim_{t \rightarrow \infty} \Pr[\eta_t(x) = \omega \text{ for all } x \text{ in } X] = 0$. The process converges exponentially fast to a configuration where some agents are permanently wrong.*

The proof of the Theorem is contained in Appendix A. It relies on the explicit characterization of the process of inference underlying agents' optimal choices. Let $\{\tau^i\}$ be the sequence of random times at which agents in the population have an opportunity to revise their choice. In other words, τ^1 is the first time an agent (randomly chosen) has the opportunity to change his choice, τ^2 is the second time, etc. Consider the following strategy **S1**:

$$\begin{aligned} \text{At any time } \tau_x: & \quad \text{if } \eta_{\tau_x}(x-1) = \eta_{\tau_x}(x+1), \text{ then choose } \eta_{\tau_x}(x \pm 1); \\ & \quad \text{if } \eta_{\tau_x}(x-1) \neq \eta_{\tau_x}(x+1), \text{ then stick to } \eta_{\tau_x}(x). \end{aligned}$$

Recall that by the assumptions of the model, $\Pr[\tau_x = \tau_y = \tau^i] = 0$ (i.e., no two agents act at the same time) and $\Pr[\tau_x = \tau^i] = \Pr[\tau_y = \tau^i]$ for all x, y (i.e., within any time period agents are equally likely to receive an updating opportunity). Hence, to show that the above strategy is optimal, we need to show that it is so at any time, that is, for any $\tau_x = \tau^i$, and at any stage of the revision process, that is, for any τ_{x_l} , $l = 1, 2, \dots$

²Note that this is the case contemplated by the canonical model of social learning.

To start the analysis, it is useful to notice that this strategy is clearly optimal if x is the first agent to receive an updating opportunity. Indeed, at time $t = 0$, upon receiving the signal $\theta(x)$, any agent x has an LLR equal to

$$\lambda_0(x) = 2 \log \left[\frac{q}{1-q} \right] \left(\theta(x) - \frac{1}{2} \right)$$

and, given the incentive structure, $\eta_0(x) = \theta(x)$ is the optimal choice at time $t = 0$. Since at time $t = 0$ all agents are playing their signals, if agent x is the first to act, his information set will consist of $\{\theta(x) = \eta_0(x), \theta(x \pm 1) = \eta_0(x \pm 1)\}$ and

$$\lambda_{\tau^1}(x) = 2 \log \left[\frac{q}{1-q} \right] \left\{ \left(\theta(x) - \frac{1}{2} \right) + \left(\theta(x-1) - \frac{1}{2} \right) + \left(\theta(x+1) - \frac{1}{2} \right) \right\}.$$

As a result, the above strategy is the only optimal strategy if x is the first agent to receive an updating opportunity. As such, it is a candidate to be an optimal strategy at any time $\tau_x > 0$. By induction, in the proof, Remark 5 shows that this strategy is optimal as long as it is followed by all other agents. On the basis of this result, Remark 6 proceeds to characterize the equilibrium of the social learning process, as well as to compute the rate of convergence of the process of action choices. Finally, Remark 7 emphasizes that this social learning process with local interactions gives rise to an extreme form of informational inefficiency, as the probability that the whole population learns to behave optimally, given the state of the world, is in fact zero. The result shows that the process of social learning may get absorbed in one of an infinite number of states where someone chooses the correct action and someone does not. In essence, the reason for this endemic multiplicity of stable limit configurations is that agents are extremely inward looking, in the sense that their choices are entirely determined by what happens inside their small neighbourhood, and although neighbourhoods are overlapping, information fails to be transmitted. To gain further intuition, consider the border between a cluster (of at least two agents) choosing action 0 and a cluster (of at least two agents) choosing action 1. As each of the two bordering agents has at least one neighbour choosing the same action as he does, neither of them will ever flip and information transmission will come to a halt.

It is interesting to consider what would happen if such bordering agents did not rely so much on their private (possibly wrong) information and allowed for the possibility of changing action in any situation where the actions chosen by their neighbours were in conflict. In what follows, we build on this intuition.

3.2 Unequally Informed Agents

We now move to a different scenario, in which agents receive signals of different precisions. We study, in particular, the case in which some agents in the population are perfectly informed, while others receive an uninformative signal. In terms of our notation, this means that $q^h = 1$, $q^l = 0.5$ and that the probability that each agent

receives a strongly-informative signal is denoted by $r \in (0, 1)$. We assume that at each time x is to take a decision, he observes the action currently chosen by one of his two neighbours, drawn with equal probabilities in $\{x \pm 1\}$.

The main difference with respect to the model previously analyzed is that now information is not homogeneous among agents: while agents who receive a fully informative signal will always choose the correct action independently of their neighbours, agents who receive a (fully) uninformative signal will draw Bayesian inference on the basis of their observation, that now consists of the action currently chosen by a single neighbour. The next result shows that the properties of the entailed social learning process with local interactions are very different in this set-up.

Theorem 4 *If each agent is perfectly informed with probability $r \in (0, 1)$ and perfectly uninformed otherwise, and can observe a randomly drawn neighbour, then the process of social learning with local interactions is adequate, as $\Pr[\lim_{t \rightarrow \infty} \eta_t(x) = \omega \text{ for all } x \text{ in } X] = 1$. The process converges slowly (at rate \sqrt{t}) to a configuration where all agents choose the correct action.*

The proof is contained in Appendix B and its logic parallels that of Theorem 3. Since agents who are perfectly informed always choose the correct action, the focus is on the characterization of the behaviour of the remaining uninformed agents. For convenience, we denote agents who are perfectly informed as \mathbf{x} and agents who are perfectly uninformed as x . Under the assumptions of this model,

$$I_t(x) = \begin{cases} \{\theta^\sigma(x), \sigma\} & \text{at } t = 0, \\ \{\eta_t(x), \eta_t(y), \Pr[y = x - 1] = \Pr[y = x + 1] = 0.5, \sigma\} & \text{at each } t = \tau_x. \end{cases}$$

Consider the following strategy **S2** for agent x :

$$\text{At any time } \tau_x: \quad \text{choose } \eta_{\tau_x}(x) = \eta_{\tau_x}(y).$$

To understand why this strategy is optimal, suppose agent x is the first to act and observes $\eta_{\tau_1}(y) = 1$. As x is uninformed, $\lambda_0(x) = 0$ and

$$\lambda_{\tau_1}(x) \equiv \log \frac{\Pr[\eta(y) = 1 \mid \omega = 1] \Pr[\omega = 1]}{\Pr[\eta(y) = 1 \mid \omega = 0] \Pr[\omega = 0]} = \log \frac{r + (1 - r)\frac{1}{2}}{(1 - r)\frac{1}{2}} = \log \frac{1 + r}{1 - r} > 0$$

where, we recall, r is the probability that agent y is perfectly informed. By the same token, observing $\eta_{\tau_1}(y) = 0$, leads agent x to revise his LLR to $\log \frac{1-r}{1+r} < 0$. Since each neighbour is equally likely to be observed, choosing $\eta_{\tau_1}(x) = \eta_{\tau_1}(y)$ is optimal for agent x at his first updating opportunity, τ^1 . *De facto*, this strategy posits that when the evidence provided by the observation of the neighbours is strong (i.e., when both neighbours choose the same action), agent x optimally chooses to agree with them (as he did in the previous model), but whenever the actions observed by x

provide only weak evidence on the unknown state (i.e., when neighbours disagree), agent x may choose any of the two actions with equal probability.

Remark 8 shows that strategy S2 is optimal for any agent x at any stage of the revision process, as long as it is followed by all other uninformed agents. Remark 9 characterizes the limit behaviour of this social learning process and Remark 10 shows that this social learning process is adequate, in the sense that, albeit very slowly, all agents will eventually choose the correct action.

The intuition behind this result is that, in this model, those agents who are aware of being uninformed will disregard their private information and be more prone to changing actions. On the other hand, those agents who are aware of being perfectly informed do unerringly play the correct action. As a result, this modeled heterogeneity in available private information significantly improves the efficiency of the mechanism of information transmission.

4 Related Literature

The literature on social learning has been growing very fast over the last two decades. Early models of social learning (Bikhchandani *et al.* [6], Banerjee [4]), typically referred to as models of herding and informational cascades, show that individuals who take choices sequentially and observe the choices made by all predecessors may neglect their private information and base their decisions entirely on what is publicly observed (an informational cascade). Herds may occur because the informational content of the public history of choices may overwhelm the information contained in the agents' private signals. Since in this case agents' private information is not revealed through their actions, learning may come to a halt and beliefs never converge to the truth.

After this seminal work, many papers have extended the analysis in various directions and applied it to study different issues (for comprehensive surveys see, among others, Gale [19], Devenow *et al.* [14], Hirshleifer and Theo [23], Chamley [10], Bikhchandani *et al.* [7], Vives [28]). The papers most related to ours study social learning when agents can only observe other individuals to which they are connected in a network. This is the case in Gale and Kariv [20]. In their model, agents act simultaneously, have perfect recall, and can revise their previous decisions. Their results show that, under some conditions, despite the fact that agents cannot observe the entire population, eventually, uniformity of actions occurs. Acemoglu *et al.* [1] analyze a situation in which agents observe the past actions of a stochastically-generated neighbourhood of individuals. In their set up, when beliefs are unbounded, there is asymptotic learning (defined as convergence of the actions to the correct one) as long as there is some minimal amount of "expansion in observations". For many common deterministic and stochastic networks, bounded private beliefs are, instead, incompatible with asymptotic learning, as in the canonical model of social learning.

Nevertheless, the authors find conditions under which asymptotic learning obtains even with bounded private beliefs for a large class of stochastic network topologies.

Other models of social learning in networks include Bala and Goyal [3], De Marzo *et al.* [13], Acemoglu *et al.* [2] and Ellison and Fudenberg [17]. What these papers have in common is that they assume some form of bounded rationality. In Bala and Goyal [3], agents in a network choose after observing their neighbours' actions and payoffs. It should be noticed that this is a model of social experimentation rather than social learning: agents learn by observing the outcome (payoff) of an experiment (choice of action) rather than by inferring another agent's private information from his action. There is private information in their model, but agents are assumed to ignore it to some extent. By assumption, each agent learns from his neighbour's actions, but does not ask what information might have led the neighbour to choose those actions. They show that, in a connected network, in the long run, everyone adopts the same action and that the action chosen can be suboptimal. De Marzo *et al.* [13] and Acemoglu *et al.* [2] also focus on networks, but learning in these models is non-Bayesian. In Ellison and Fudenberg [17], finally, agents consider the experiences of their neighbours 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.

Beyond the papers that focus on social learning in networks, our paper is also related in the motivation to studies that analyze social learning when agents can only observe a subset of other agents' actions. Banerjee and Fudenberg [5] and Smith and Sørensen [27] study social learning when agents observe a sample of predecessors' actions. Banerjee and Fudenberg [5] 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 for herding to arise, and conditions for all agents to settle on the correct choice. Smith and Sørensen [27] study a sequential decision model in which agents can only observe unordered random samples of predecessors' actions. They characterize different conditions on the sampling procedure and on the beliefs to have complete or incomplete learning. When the past is not over-sampled, that is, not affected forever by any one individual, and when beliefs are unbounded, complete learning eventually obtains.

The theme of imperfect observability of other agents' actions is common to a series of other papers in the social learning literature. Çelen and Kariv [11] extend the standard model of sequential social learning by allowing each agent to observe the decision of his immediate predecessor only. The prediction of these authors is that behavior does not settle on a single action. Long periods of herding can be observed, but switches to the other action may occur. As time passes, the periods of herding become longer and longer, and the switches increasingly rare. Larson [24] analyzes a situation in which agents observe a weighted average of past actions before making a choice in a continuous action space. Similarly to our work, the focus

is on the speed of learning (since the continuous action space guarantees that complete learning eventually occurs). An interesting observation of this study is that the speed of learning depends on how effectively the noise coming from early actions is purged. Guarino *et al.* [21] introduce a model of aggregate information cascades where only one of two possible actions is observable to others. Agents make a binary decision in a random order and agents are not aware of their own position in the sequence. When called upon, they are only informed about the total number of others who have chosen the observable action before them. The result of this study is that only one type of cascade arises in equilibrium, the cascade on the observable action. Collander and Hörner [12] present a model in which an agent observes only the total number of choices of each type (in a binary action setting), rather than the full sequence of actions. They characterize conditions under which later agents optimally imitate the minority, rather than the majority action.

Finally the issue of imperfect observability is also discussed in recent papers by Bohren [8], Eyster and Rabin [18] and Guarino and Jehiel [22] in contexts in which agents are not fully rational. The imperfect observability can actually alleviate some biases that bounded rationality produces in a classical model of learning with continuous action space similar to that of Lee [25].

5 Conclusions

In our economy, a large population of individuals have to choose one out of two available actions. Each action is optimal in one of two unknown states of the world. Agents repeatedly and reversibly choose an action, the payoff to which will materialize when the state of the world realizes. Agents derive a posterior probability on the basis of a symmetric binary signal that they receive and by observing a sample of other agents, called their neighbours. Observed choices can be informative, since signals are, and this raises an issue of informational externality. While signals are generated by a probability distribution that is exogenously given to each agent, observed choices are endogenous to the model and, given the postulated spatial structure of the process, show a potentially high degree of spatial correlation.

We have studied two social learning processes, one in which agents are homogenous in the quality of the private information they receive (as measured by the precision of their signals) and one in which the quality of the information differs (in that some agents are perfectly informed, while others are completely uninformed). We have compared the two social learning processes in terms of the probability with which they may prove to be adequate, that is reach a configuration where every agent adopts the action that is optimal given the true state of the world. As we pointed out, since beliefs are bounded by the local nature of the social interaction, complete learning is out of reach within this class of models. We have shown that the specific kind of heterogeneity embedded in the second model guarantees that, albeit very slowly, the social learning process is adequate, since it converges to a configuration

where all agents adopt the correct action. This cannot be so in the first model, as the social learning process gets absorbed exponentially quickly in a configuration where some agents permanently adopt the incorrect action. The explicit characterization of the rates of convergence proves to be relevant if one wants to compare the two models in terms of informational efficiency: while in the first model we observe a quick and complete blockage of information transmission, in the second information does get disseminated, but this occurs very slowly.

We conclude with a few remarks and conjectures.

In the model of Section 3.1 social learning is not adequate. One may wonder whether learning could be adequate in the presence of some agents who are perfectly informed about the state of the world. To address this issue suppose that $1 = q^h > q^l > 0.5$ and $0 < r < 1$, that is, suppose that all agents receive an informative signal and a positive measure of them are perfectly informed. We conjecture that Theorem 3 would carry on in this case as well. Any perfectly informed agent, say agent 0, knows the true state of the world, ω , and sticks to $\eta(0) = \omega$ independently of the actions adopted in his neighbourhood. This relevant information is not necessarily transmitted to others. Since agents ± 1 are themselves informed and observe both neighbours, if the other neighbour, ± 2 respectively, chooses the wrong action, $1 - \omega$, they would stick to the wrong action. Also, the speed at which the social learning process converges would still be driven by the use of strategy S1 on the part of the less informed agents, and hence would still be exponentially fast. The result would hold, *a fortiori*, if $q^h < 1$, since now agents know that no neighbour is perfectly informed about the state of the world. This suggests that heterogeneity in the quality of information *per se* is not sufficient to guarantee that the social learning process is adequate.

In the model of Section 3.2 an extreme form of heterogeneity in the quality of information leads to adequate social learning. One may wonder whether the result extends to the case in which in the population there are uninformed agents and other agents that receive a noisy signal: $1 > q^h > q^l = 0.5$ and $0 < r < 1$. We conjecture that this heterogeneity is not sufficient to guarantee adequate learning. Strategy S2 would still be optimal for the uninformed agents in this case. Also, for a non empty set of values of q^h a better informed agent, say 0, would play his signal independently of his neighbours. As a result, with positive probability, agents ± 1 (if uninformed) will learn agent 0's action, and transmit it to agents ± 2 (if uninformed), etc. The difference with Section 3.2 is that with positive probability some better informed agents receive the wrong signal and never correct their initial decision, which precludes adequate learning.

In essence, sufficient conditions that guarantee that the social learning process is adequate, are the existence of some agents who know the truth and unerringly choose the correct action, together with the existence of some agents who, being uninformed, decide to change their original choice. Perhaps contrary to intuition, in this set-up of limited memory a process of slow clustering on the correct decision ensues, by thus

guaranteeing that the social learning process reaches an informationally efficient state. Our explicit characterization of the convergence rates shows that the speed at which this cluster grows does not depend on the proportion of perfectly informed agents in the population: our estimate is in fact unaltered for any value of the parameter $0 \leq r < 1$.³

Finally, one may wonder if endowing agents with the ability to convey more information to their neighbours, for example by fully revealing their posterior belief to each other, may enhance informational efficiency. Duffie *et al.* [15] show that this is the case in a model that does not account for a spatial dimension and are able to characterize explicitly the convergence of the cross-sectional distribution of beliefs to a common posterior. Their results do not apply to our setting,⁴ but the question grants future research.

Appendix A

Proof of Theorem 3.

The proof of the Theorem 3 is split into a few Remarks: Remark 5 shows that the model admits an equilibrium; Remark 6 characterizes limit behaviour and convergence rates of this process of social learning with local interactions and finally Remark 7 evaluates the degree of informational efficiency of the process.

Remark 5 *Suppose all agents $y \neq x$ choose stationary strategy S1. Then this strategy is also optimal for x at any time τ_{x_1} .*

Proof. First, suppose that $\tau_{x_1} = \tau^1$. Then, the statement is true, since $\eta_0(y) = \theta(y)$ for all y . Suppose now that $\tau_{x_1} > \tau^1$. Let us describe the process of inference undertaken by agent x in such a case (i.e., if he knew that at least one other agent had received an updating opportunity before). We drop the time subscript for notational convenience. Due to the symmetry of the model, WLOG we consider $\theta(x) = 0$.

Let us consider first the case in which $\eta(x-1) = 1$. Agent x needs to infer $\theta(x-1)$ on the basis of $I(x) = \{\theta(x) = 0 = \eta(x), \eta(x-1) = 1, \eta(x+1)\}$. By Bayesian updating,

$$\begin{aligned} & \Pr[\theta(x-1) = 1 \mid I(x)] \equiv \\ & (\Pr[\eta(x-1) = 1 \mid \theta(x-1) = 1, \theta(x) = 0, \eta(x+1)] \Pr[\theta(x-1) = 1 \mid \theta(x) = 0, \eta(x+1)]) \\ & + \Pr[\eta(x-1) = 1 \mid \theta(x-1) = 0, \theta(x) = 0, \eta(x+1)] \Pr[\theta(x-1) = 0 \mid \theta(x) = 0, \eta(x+1)]^{-1} \end{aligned}$$

³Clearly, for $r = 0$ the social learning process could fail to be adequate, as it would only show consensus on a particular action, not necessarily the correct one.

⁴Due to the amount of spatial correlation, the information sets upon which our agents base their decisions are not disjoint, which contradicts a basic assumption of their model.

If $\tau_{(x-1)_1} > \tau_{x_1}$, this probability is one, as $x - 1$ is playing his signal, by construction. If $\tau_{(x-1)_1} < \tau_{x_1}$ and agent $x - 1$ has followed strategy S1, this probability is also equal to one, since $\Pr[\eta(x - 1) = 1 \mid \theta(x - 1) = 0, \theta(x) = 0, \eta(x + 1)] = 0$. Hence an agent who observes a neighbour choosing an action different from the action he himself is choosing, infers that neighbour is playing his signal:

$$\Pr[\theta(x - 1) = 1 \mid I(x)] = 1.$$

Let us consider now the case of $\eta(x - 1) = 0$. Agent x needs to infer $\theta(x - 1)$ on the basis of $I(x) = \{\theta(x) = 0 = \eta(x), \eta(x - 1) = 0, \eta(x + 1)\}$. If $\tau_{(x-1)_1} > \tau_{x_1}$, clearly $\theta(x - 1) = \eta(x - 1) = 0$. If $\tau_{(x-1)_1} < \tau_{x_1}$ and agent $x - 1$ has followed strategy S1, then by a logic analog to that followed in the previous paragraph,

$$\Pr[\eta(x - 1) = 0 \mid \theta(x - 1) = 1, \theta(x) = 0, \eta(x + 1)] = \Pr[\eta(x - 2) = 0 \mid I(x)]$$

and

$$\Pr[\eta(x - 1) = 0 \mid \theta(x - 1) = 0, \theta(x) = 0, \eta(x + 1)] = 1$$

Let $\Pr[\eta(x - 2) = 0 \mid I(x)] \equiv 1 - \alpha$ and $\Pr[\theta(x - 1) = 1 \mid \theta(x) = 0, \eta(x + 1)] \equiv 1 - \beta$. Then,

$$\Pr[\theta(x - 1) = 1 \mid I(x), \tau_{(x-1)_1} < \tau_{x_1}] \equiv \frac{(1 - \alpha)(1 - \beta)}{(1 - \alpha)(1 - \beta) + \beta}$$

Note that β is the belief that agent x has on the signal of $x - 1$ being equal to zero. As such, β depends on the value of $\eta(x + 1)$, i.e., either

$$\Pr[\theta(x - 1) = 1 \mid \theta(x) = 0, \eta(x + 1) = 1] = 0.5$$

or

$$\Pr[\theta(x - 1) = 1 \mid \theta(x) = 0, \eta(x + 1) = 0] < \Pr[\theta(x - 1) = 1 \mid \theta(x) = 0] < 0.5$$

Hence $\beta \geq 0.5$ and, for all $0 \leq \alpha \leq 1$,

$$\Pr[\theta(x - 1) = 1 \mid I(x), \tau_{(x-1)_1} < \tau_{x_1}] \equiv \frac{(1 - \alpha)}{(1 - \alpha) + \frac{\beta}{1 - \beta}} \equiv \gamma < \frac{1}{2},$$

As a result,

$$\Pr[\theta(x - 1) = 1 \mid I(x)] \leq \Pr[\tau_{(x-1)_1} < \tau_{x_1}] \gamma < \frac{1}{2}.$$

Hence an agent who observes a neighbour choosing the same action as he himself is choosing, infers that the neighbour is more likely to be playing his signal (than to have used an updating opportunity).

As a result of the above considerations, and for $y \in \{x \pm 1\}$ the conditional expectations of $\theta(y)$ are

$$\begin{aligned} E[\theta(y) \mid \theta(x) = 0, \eta(y) = 1] &= 1, \text{ and} \\ E[\theta(y) \mid \theta(x) = 0, \eta(y) = 0] &< \frac{1}{2}. \end{aligned}$$

We now proceed to show that, given these conditional expectations, the strategy is optimal for x at time τ_{x_1} for any possible $I(x) = \{\theta(x) = 0 = \eta(x), \eta(x \pm 1)\}$.

Let us first prove the "if" part. Suppose that $I(x) = \{\theta(x) = 0, \eta(x - 1) = \eta(x + 1) = 1\}$. By the above considerations $\theta(x - 1) = \theta(x + 1) = 1$. Hence:

$$\begin{aligned}\lambda_1(x) &\equiv 2 \log \left[\frac{q}{1-q} \right] \left(E[\theta(x-1) | I(x)] + \theta(x) + E[\theta(x+1) | I(x)] - \frac{3}{2} \right) = \\ &= 2 \log \left[\frac{q}{1-q} \right] \left(2 - \frac{3}{2} \right) = \log \left[\frac{q}{1-q} \right] > 0\end{aligned}$$

Now let us prove the "only if" part. We have to consider different cases.

Case a): Suppose that $I(x) = \{\theta(x) = 0, \eta(x - 1) = \eta(x + 1) = 0\}$. By the above considerations $E[\theta(x - 1) | I(x)] < 0.5$. Hence:

$$\begin{aligned}\lambda_1(x) &\equiv 2 \log \left[\frac{q}{1-q} \right] \left(E[\theta(x-1) | I(x)] + \theta(x) + E[\theta(x+1) | I(x)] - \frac{3}{2} \right) = \\ &< 2 \log \left[\frac{q}{1-q} \right] \left(1 - \frac{3}{2} \right) = -\log \left[\frac{q}{1-q} \right] < 0\end{aligned}$$

since $q > 0.5$, as assumed.

Case b): Suppose that $I(x) = \{\theta(x) = 0, \eta(x - 1) = 0, \eta(x + 1) = 1\}$ (or viceversa). By the above considerations ($E[\theta(x - 1) | I(x)] < 0.5$) and $E[\theta(x + 1) | I(x)] = 1$. Hence:

$$\begin{aligned}\lambda_1(x) &\equiv 2 \log \left[\frac{q}{1-q} \right] \left(E[\theta(x-1) | I(x)] + \theta(x) + E[\theta(x+1) | I(x)] - \frac{3}{2} \right) = \\ &< 2 \log \left[\frac{q}{1-q} \right] \left(\frac{1}{2} + 1 - \frac{3}{2} \right) = 0\end{aligned}$$

since $q > 0.5$, as assumed.

This concludes the proof that, under the stated assumptions, this strategy is optimal at time $\tau_{x_1} = \tau^i$ for $i = 1, 2, \dots$ (i.e., at the first updating opportunity that agent x gets, independently of when this opportunity arises). We now show that the statement holds at any time τ_{x_l} for $l = 2, 3, \dots$

Consider $\tau_{x_2} > \tau_{x_1}$ and let $I_2(x) = \{\theta(x) = 0, \eta_{\tau_{x_2}}(x \pm 1)\}$ denote x 's information set at time τ_{x_2} . If $\eta_{\tau_{x_1}}(x) = \theta(x)$, clearly the previous part of the proof holds in this case as well. Suppose instead that $\eta_{\tau_{x_1}}(x) \neq \theta(x)$. In this case x has flipped to $\eta(x) = 1$ at time τ_{x_1} , because $\eta_{\tau_{x_1}}(x \pm 1) = 1$. By the reasoning above, this means that x could perfectly infer that $\theta(x \pm 1) = 1$. Hence, strategy S1 is optimal at time τ_{x_2} as well. In other words, at the second updating opportunity and for any x :

$$\begin{aligned}\text{either } \eta_{\tau_{x_1}}(x) &= 1 \text{ and } \lambda_1(x) = \log \left[\frac{q}{1-q} \right] \\ \text{or } \eta_{\tau_{x_1}}(x) &= 0 \text{ and } \lambda_1(x) = -\log \left[\frac{q}{1-q} \right]\end{aligned}$$

An entirely analog reasoning shows that the strategy is also optimal at any time τ_{x_l} for $l > 2$, when one notices that in between any $\tau_{x_{l+1}}$ and τ_{x_l} within x 's neighbourhood the number of agents $x \pm 1$ such that $\eta(x \pm 1) = \eta(x)$ cannot decrease. To see this, consider $\tau_{x_3} > \tau_{x_2}$ and let $I_3(x) = \{\theta(x) = 0, \eta_{\tau_{x_3}}(x \pm 1)\}$. If $\eta_{\tau_{x_2}}(x) = \theta(x) = 0$, it must be that also $\eta_{\tau_{x_1}}(x) = \theta(x) = 0$ (since it must have been that $\eta_{\tau_{x_1}}(x \pm 1) = 1$) and the first part of the proof holds. If $\eta_{\tau_{x_2}}(x) = 1 \neq \theta(x) = 0$, then x must have flipped either at time τ_{x_1} or at time τ_{x_2} , and again we know from the previous part of the proof that strategy S1 was optimal in those cases. As a Corollary, the above reasoning shows that, within this model, each agent can flip at most once. ■

Remark 6 *If agents use strategy S1, the characterization of the limit behaviour of the social learning process is as follows⁵.*

Let $\{\hat{\eta}\}$ be the set of configurations such that, for each x in X , there is at least a y in $N(x) = \{\eta(x \pm 1)\}$ such that $\eta(y) = \eta(x)$. Then, starting from any given initial distribution, μ^ω , the process converges in probability to a configuration $\eta_\infty \in \{\hat{\eta}\}$:

$$P^{\mu^\omega}[\lim_{t \rightarrow \infty} \eta_t = \eta_\infty] = 1$$

Convergence obtains exponentially fast:

$$P^{\mu^\omega}[\eta_t \neq \eta_\infty] \propto \exp[-t]$$

Proof. We shall find it convenient to model transitions in terms of *flip rates*, i.e., the rates at which $\eta_t(x)$ flips to $1 - \eta_t(x)$. By *flip rate* c we mean that the probability that the transition occurs in an infinitesimal time dt is $c dt$. We shall denote flip rates by $c(x, \eta_t)$ to emphasize their dependence on the current state of action chosen in the population and assume that $\Pr[\eta_t(x) \mid 1 - \eta_t(x)] = c(x, \eta_t)t + o(t)$.

By Remark 5, the flip rates for this process are:

$$c(x, \eta) = \begin{cases} 1 & \eta(y) \neq \eta(x) \quad \forall y = \{x \pm 1\} \\ 0 & \text{otherwise} \end{cases}$$

and the characterization of $\hat{\eta}$ follows by simple inspection of these.

To show that the process of actions converges, let $\delta_{x,y}(t) = 1$ if $\eta_t(x) \neq \eta_t(y)$, and 0 otherwise. Recall that agents live on a one-dimensional lattice $X = \mathbb{Z}^1 = \{\dots, -2, -1, 0, +1, +2, \dots\}$. We follow the lines of Durrett *et al.* [16] and define the following function:

$$\Upsilon_t = \sum_{x \in X} \sum_{y \in N(x)} \exp[-|x + y|] \delta_{x,y}(t)$$

⁵In stating the results, we use the following additional notation. We denote any probability distribution over the state space by μ_t , and the initial distribution by μ^ω . Since at time $t = 0$ choices are determined by the signals and in any given state of the world ω these are stochastically independent, this initial distribution is by construction a product measure. As for any $t > 0$ choices may instead depend on the spatial configuration of action chosen within neighbourhoods, μ_t will typically display an amount of spatial correlation.

Note that, by construction, $0 \leq \Upsilon_t < \bar{\Upsilon} < \infty$. We shall show that, starting from Υ_0 , at any time in which any x flips from $\eta(x)$ to $1 - \eta(x)$, this function decreases by a strictly positive amount. To this aim, let $\Upsilon_t(x)$ be:

$$\Upsilon_t(x) = \sum_{y \in N(x)} \exp[-|x + y|] \delta_{x,y}(t)$$

and for simplicity⁶ take $x = 0$, with neighbours $y \in \{-1, +1\}$:

$$\Upsilon_t(0) = \sum_{y \in \{-1, +1\}} \exp[-|y|] \delta_{0,y}(t)$$

Note that agent $x = 0$ will flip if and only if $\sum_y \delta_{0,y}(t) = 2$, and, by the construction of the model, this can happen with positive probability. Suppose that this happens. Then the drop in Υ at site 0, after the flip occurred, is equal to $2 \exp[-1]$ which is strictly positive. As the same argument applies to any generic site, this implies that the function Υ_t is strictly decreasing at any time at which an agent flips action.

Let $\hat{\Upsilon} \equiv \hat{\Upsilon}(\hat{\eta})$ be the value of this function at any stable configuration $\hat{\eta}$ such that $\lim_{t \rightarrow \infty} \eta_t = \hat{\eta}$ and consider $\hat{\Upsilon}_t \equiv \Upsilon_t - \hat{\Upsilon}$ along the realizations of the process leading to $\hat{\eta}$. To show that convergence obtains exponentially fast we will show that there exists $k > 0$ and $\varepsilon > 0$ such that:

$$P^{\Upsilon_0}[\hat{\Upsilon}_t > 0] \leq k \Upsilon_0 \exp[-\varepsilon t]$$

To this aim, we need to make the transition from convergence along integer times (as in $\hat{\Upsilon}_t$) to convergence in real time (for η_t). Let $\delta_{x,y}[(n-1)t, nt]$ for $n \geq 1$ and $\hat{\Upsilon}_{nt} = [\hat{\Upsilon}_{(n-1)t, nt}] \leq \sum_{k=1, \dots, t} \hat{\Upsilon}_{(k-1)t, kt}$. Since, by construction, $\hat{\Upsilon}$ is finite, $E[\hat{\Upsilon}_t]$ is also finite. Since $\hat{\Upsilon} \leq \bar{\Upsilon} < \infty$, then $\hat{\hat{\Upsilon}} \equiv (\bar{\Upsilon})^{-1} \hat{\Upsilon} \leq 1$. Let

$$\tilde{\Upsilon} \equiv E[\exp[\xi \hat{\hat{\Upsilon}}_{nt}]] < 1$$

which is true for a small positive ξ . This implies that:

$$\Pr[\hat{\Upsilon}_{nt} > 0] \leq \tilde{\Upsilon}^n$$

and since $P^{\Upsilon_0}[\hat{\Upsilon}_s > 0]$ is monotonic in s , for $k \equiv \tilde{\Upsilon}^{-1}$ and for $\varepsilon \equiv t^{-1} \log \tilde{\Upsilon}^{-1}$ the assert is proved. ■

Remark 7 *If agents use strategy S1, the process of social learning with local interactions is not adequate and*

$$\Pr[\lim_{t \rightarrow \infty} \eta_t(x) = \omega \text{ for all } x \text{ in } X] = 0$$

⁶This is done WLOG, since the initial distribution that determines the initial condition is (a product measure and hence) translation invariant.

Proof. Recall that the initial choice of actions is produced by a product measure: $\Pr^{\mu_{\omega=1}}[\eta : \eta(x) = 1] = q$ or $\Pr^{\mu_{\omega=0}}[\eta : \eta(x) = 1] = 1 - q$ respectively. Since $0.5 < q < 1$ the probability that any two adjacent agents receive the same signal is strictly positive. Once this happens, as the above Remark shows, these agents will never flip. Hence, this process fails to satisfy Definition (2) and learning is not adequate. ■

Appendix B

Proof of Theorem 4.

The proof of Theorem 4 is split into three Remarks: Remark 8 shows that an equilibrium exists, Remark 9 characterizes the limit behaviour and Remark 10 evaluates the degree of informational efficiency of this model. The logic of the proof parallels that of Theorem 3.

Remark 8 *Suppose all agents $y \neq x$ choose strategy S2. Then this strategy is also optimal for x at any τ_{x_1} .*

Proof. We follow exactly the same logic as that of Remark 5 and describe the process of inference undertaken by agent x at time τ_{x_1} , (we drop the time subscript for notational convenience). Notice that, within this model, agent x cannot draw any inference from his signal. Also, agent x cares about his neighbours' signal only insofar as they are informed.

Since $q(x) = 0.5$ by construction, $\lambda_0(x) = 0$.

Suppose $\tau_{x_1} = \tau^1$ (i.e., agent x is the first to receive an updating opportunity) and $\eta(x) = 1$. Agent x needs to compute $\lambda_1(x)$ on the basis of $I(x) = \{\eta(x), \eta_{\tau_{x_1}}(y)\}$, for $y \in \{x \pm 1\}$, resulting in:

$$\begin{aligned} \lambda_1(x) &\equiv \log \frac{\Pr[\eta(y) \mid \omega = 1] \Pr[\omega = 1]}{\Pr[\eta(y) \mid \omega = 0] \Pr[\omega = 0]} \\ &= \begin{cases} \log \frac{r+(1-r)^{\frac{1}{2}}}{(1-r)^{\frac{1}{2}}} = \log \left[\frac{1+r}{1-r} \right] > 0 & \text{if } \eta(y) = \eta(x) = 1 \\ \log \frac{(1-r)^{\frac{1}{2}}}{r+(1-r)^{\frac{1}{2}}} = \log \left[\frac{1-r}{1+r} \right] < 0 & \text{if } \eta(y) \neq \eta(x) = 1 \end{cases} \end{aligned}$$

Since $\Pr[y = x - 1] = \Pr[y = x + 1] = 0.5$, this shows that S2 is optimal at time $\tau_{x_1} = \tau^1$.

Consider $\tau_x > \tau^1$ and let $I_{\tau_x}(x) = \{\eta_{\tau_x}(x), \eta_{\tau_x}(y)\}$ denote x 's information set at time τ_x . When drawing inference, agent x has now to consider the possibility that agent y may have received an updating opportunity and may have used strategy S2. As x is uninformed, $\lambda(x) = 0$. Let $s \equiv \Pr[\tau_y < \tau_x]$, i.e., the probability that agent y has received an updating opportunity before agent x .

Suppose $y = x - 1$, $\eta(x - 1) = 1$ and $\eta(x) = 1$. Then:

$$\begin{aligned} \Pr[\eta(x - 1) = 1 \mid \omega = 1] &= r + (1 - r)\left\{(1 - s)\frac{1}{2} + s\left[\eta(x - 2)\frac{1}{2} + \frac{1}{2}\right]\right\} \\ \Pr[\eta(x - 1) = 1 \mid \omega = 0] &= (1 - r)\left\{(1 - s)\frac{1}{2} + s\left[\eta(x - 2)\frac{1}{2} + \frac{1}{2}\right]\right\} \end{aligned}$$

where the term in square brackets refers to the possibility that $x - 1$ might have chosen $\eta(x - 1) = 1$ as a result of S2 (and hence observed either $\eta(x - 2) = 1$ or $\eta(x) = 1$). Notice that, by construction, $\eta(x - 2) \in \{0, 1\}$ is given at time τ_x .

Suppose $y = x + 1$, $\eta(x + 1) = 0$ and $\eta(x) = 1$.

Then:

$$\begin{aligned} \Pr[\eta(x + 1) = 0 \mid \omega = 1] &= (1 - r)\{(1 - s)\frac{1}{2} + s[\frac{1}{2}(1 - \eta(x + 2))]\} \\ \Pr[\eta(x + 1) = 0 \mid \omega = 0] &= r + (1 - r)\{(1 - s)\frac{1}{2} + s[\frac{1}{2}(1 - \eta(x + 2))]\} \end{aligned}$$

where the term in square brackets refers to the possibility that $x + 1$ might have chosen $\eta(x + 1) = 0$ as a result of S2 (and hence observed either $\eta(x + 2) = 0$ or $\eta(x) = 1$). Notice that, by construction, $\eta(x + 2) \in \{0, 1\}$ is given at time τ_x .

As a result,

$$\begin{aligned} \lambda(x) &\equiv \log \frac{\Pr[\eta(y) \mid \omega = 1] \Pr[\omega = 1]}{\Pr[\eta(y) \mid \omega = 0] \Pr[\omega = 0]} \\ &= \begin{cases} \log \frac{r + (1 - r)\{(1 - s)\frac{1}{2} + s[\eta(x - 2)\frac{1}{2} + \frac{1}{2}]\}}{(1 - r)\{(1 - s)\frac{1}{2} + s[\eta(x - 2)\frac{1}{2} + \frac{1}{2}]\}} > 0 & \text{if } y = x - 1, \eta(y) = 1 = \eta(x) \\ \log \frac{(1 - r)\{(1 - s)\frac{1}{2} + s[\frac{1}{2}(1 - \eta(x + 2))]\}}{r + (1 - r)\{(1 - s)\frac{1}{2} + s[\frac{1}{2}(1 - \eta(x + 2))]\}} < 0 & \text{if } y = x + 1, \eta(y) = 0 \neq \eta(x) \end{cases} \end{aligned}$$

Since $\Pr[y = x - 1] = \Pr[y = x + 1] = 0.5$, this shows that S2 is optimal at any time τ_x . ■

Remark 9 *If agents use strategy S2, the characterization of the limit behaviour of the social learning process is as follows.*

Let η^ω be the configurations where $\eta(x) = \omega$ for all $x \in X$. Then, starting from any given initial condition, μ^ω , the process converges in probability to configuration η^ω :

$$P^{\mu^\omega}[\lim_{t \rightarrow \infty} \eta_t = \eta^\omega] = 1$$

Convergence obtains slowly, namely at rate \sqrt{t} :

$$P^{\mu^\omega}[\eta_t \neq \eta^\omega] \propto \frac{1}{\sqrt{t}}$$

Proof. Let us denote the population of agents as $X \cup \mathbf{X}$, where $\mathbf{x} \in \mathbf{X}$ are the informed agents and $x \in X$ are the uninformed agents. By construction, the flip rates for this process $\{c(\mathbf{x}, \eta), c(x, \eta)\}$ are:

$$c(\mathbf{x}, \eta) = 0 \tag{3}$$

$$c(x, \eta) = \begin{cases} \frac{1}{2} \sum_{y \in \{x \pm 1\}} \eta(y) & \eta(x) = 0 \\ \frac{1}{2} \sum_{y \in \{x \pm 1\}} (1 - \eta(y)) & \eta(x) = 1 \end{cases} \tag{4}$$

By simple inspection, it is clear that only the state for which $\eta(x) = \omega$ for all x in X is stationary for this process. However, since the process η defines a continuous

time Markov chain on the state-space $S = Z^1$ which is countable, but infinite, we need to prove that the process is ergodic, that is, that starting from any initial distribution μ^ω , the process will converge to η^ω with probability one (first part of the assert).

We proceed as follows. Let S_N be finite sets that increase to S , such that $\lim_{N \rightarrow \infty} S_N = S$. Define the following flip rates:

$$c_i^N = \begin{cases} \{c(\mathbf{x}, \eta), c(x, \eta)\} & \text{if } x, \mathbf{x} \in S_N \\ 0 & \text{if } x, \mathbf{x} \notin S_N \text{ and } \eta(x) = i \\ 1 & \text{if } x, \mathbf{x} \notin S_N \text{ and } \eta(x) \neq i \end{cases}$$

Let us call the process defined by these flip rates $S_{i,N}(t)$. Notice that this process is equal to the original process for all x, \mathbf{x} in S_N and characterized by all coordinates set equal to i for x, \mathbf{x} not in S_N .

Let $\mu^0 S_{0,N}(t)$ be the law of the process characterized by flip rates c_0^N when the initial distribution is given by all 0 at time 0 and let $\mu^1 S_{1,N}(t)$ be the law of the process characterized by flip rates c_1^N when the initial distribution is given by all 1 at time 0. As the original process is attractive⁷, so are the processes c_i^N and, by Theorem 2.7 in Liggett [26]:

$$\mu^0 S_{0,N}(t) \leq \mu^\theta S(t) \leq \mu^1 S_{1,N}(t)$$

for $\theta \in (0, 1)$, and

$$\begin{aligned} \lim_{N \rightarrow \infty} \lim_{t \rightarrow \infty} \mu^0 S_{0,N}(t) &= \lim_{t \rightarrow \infty} \mu^0 S(t) \\ \lim_{N \rightarrow \infty} \lim_{t \rightarrow \infty} \mu^1 S_{1,N}(t) &= \lim_{t \rightarrow \infty} \mu^1 S(t) \end{aligned}$$

WLOG suppose $\omega = 1$. Then $\lim_{t \rightarrow \infty} \mu^0 S_{0,N}(t) = \lim_{t \rightarrow \infty} \mu^1 S_{1,N}(t) = \mu^{1,N}$, that is, as $t \rightarrow \infty$, independently of the initial distribution, the process restricted on S_N converges to a configuration all ones. In fact $S_{i,N}(t)$ is a finite Markov chain over S_N , and as there is a unique absorbing state ($\eta_N^1 \equiv \{\eta(x) = 1 \text{ for all } x \in S_N\}$) we know that the unique ergodic distribution posits pointmass one on this state. As $\lim_{N \rightarrow \infty} S_N = S$, it follows that

$$\lim_{N \rightarrow \infty} \lim_{t \rightarrow \infty} \mu^0 S_{0,N}(t) = \lim_{N \rightarrow \infty} \lim_{t \rightarrow \infty} \mu^1 S_{1,N}(t) = \lim_{N \rightarrow \infty} \mu^{1,N} = \mu^1$$

and the first part of the assert follows.

To prove the second part of the statement, we need to compute the rate of convergence for this process. Notice that the rate at which social learning takes place

⁷We say that, for $\eta, \zeta \in \{0, 1\}^{Z^1}$, $\eta \leq \zeta$ if $\eta(x) \leq \zeta(x)$ for all $x \in Z^1$. Then a process is defined to be *attractive (or monotonic)* if, whenever $\eta \leq \zeta$ flip rates satisfy the following:

$$\begin{aligned} c(x, \eta) &\leq c(x, \zeta) & \text{if } \eta(x) = \zeta(x) = 0 \\ c(x, \eta) &\geq c(x, \zeta) & \text{if } \eta(x) = \zeta(x) = 1 \end{aligned}$$

is given by the speed with which those uninformed agents flip. Hence we need to study the dynamics of choices of the individuals in X . We notice that this dynamics is analog to that of the Voter's model (Liggett [26], Section 1 and 3, Chapter V or in Bramson and Griffeath [9]), well studied in the statistical literature. In the Voter's model, a voter at $x \in Z^d$ changes his opinion at an exponential rate (with mean one) proportional to the number of $2d$ nearest neighbours with the opposite opinion. If $2d$ neighbours disagree with the person at x , the flip rate is 1. It can be seen by equation (3) that this is exactly the dynamics of the uninformed agents in our model. Hence, although the asymptotics of our model are substantially different from those of the Voter's model, the dynamics is exactly the same.

To show that learning occurs at rate \sqrt{t} we proceed as follows. As the process is defined in the two dimensions of time and space, we shall find it useful to relate these two dimensions in a space-time analysis. In particular, we characterize a *clustering* process, by relying on the local specification of the model. With the term "cluster" we mean the length of a segment with all connected individuals choosing the same action. In order to see how the size of a cluster increases with time, we shall later express the length of a cluster as a function of t . Formally, given a configuration, η , we define a *cluster* as the connected components of $\{x : \eta(x) = 0\}$ or $\{x : \eta(x) = 1\}$; the size of a cluster of ones in a segment of side l around the origin as:

$$|\eta_l| = |\{x : \eta(x) = 1; x \in [-l, l]\}|$$

and the mean cluster size of η around the origin as:

$$C(\eta) = \lim_{l \rightarrow \infty} \frac{2l}{\text{number of clusters of } \eta \text{ in } [-l, l]}$$

whenever this limit exists.

Given the asymptotics described, we already know that the mean cluster size tends to grow indefinitely. To prove the statement, we need to show that the mean cluster size, $C^{\mu_\omega}(\eta_t)$, grows in probability at rate \sqrt{t} , in the sense that:

$$\frac{C^{\mu_\omega}(\eta_t)}{t^{1/2}} \rightarrow_p K$$

where K is a positive constant depending on ω . Since, as stated before, this model reproduces the of the Voter's model, this statement is proved in Bramson and Griffeath [9]. In fact, Theorem 7, p. 211 of that paper also provides the following estimate for the lower and upper bound of the limit expected value of the above quantity (re-written with our parametrization):

$$\sqrt{\pi} \left(\frac{1}{2 \frac{1+r}{2} \frac{1-r}{2}} \right) \leq \lim_{t \rightarrow \infty} E \left[\frac{C^{\mu_\omega}(\eta_t)}{t^{1/2}} \right] \leq 2 \left(\frac{\left(\frac{1+r}{2}\right)^2 + \left(\frac{1-r}{2}\right)^2}{\left(\frac{1+r}{2}\right) \left(\frac{1-r}{2}\right)} \right) \sqrt{\pi} \quad (5)$$

where $\pi = 3.1416$. ■

Remark 10 *If agents use strategy S2 , the process of social learning with local interactions is adequate, in that*

$$\lim_{t \rightarrow \infty} \Pr[\eta_t(x) = \omega \text{ for all } x \text{ in } X] = 1$$

Proof. Recall that the initial condition is produced by a product measure: $\Pr^{\mu_{\omega=1}}[\eta : \eta(x) = 1] = \frac{1+r}{2}$ or $\Pr^{\mu_{\omega=0}}[\eta : \eta(x) = 1] = \frac{1-r}{2}$ respectively. Hence, by the previous Remark, the assert follows. ■

References

1. *Acemoglu, D.; Dahleh, M.A.; Lobel, I.; Ozdaglar, A. Bayesian Learning in Social Networks. mimeo, 2010, MIT.*
2. *Acemoglu, D.; Ozdaglar, A.; Parandeh-Greibi, A. Spread of (Mis)Information in Social Networks. 2009 MIT WP 09-15.*
3. *Bala V.; Goyal S. Learning from Neighbours. Review of Economic Studies, 1998, 65, 595-621.*
4. *Banerjee, A.V. A Simple Model of Herd Behavior. Quarterly Journal of Economics 1992, 107, 797-817.*
5. *Banerjee, A.; Fudenberg, D. Word of Mouth Learning. Games and Economic Behavior 2004, 46, 1-22.*
6. *Birkhchandani, S.; Hirshleifer, D.; Welch, I. A Theory of Fads, Fashions, Custom, and Cultural Change as Information Cascades. Journal of Political Economy 1992, 100, 992-1026.*
7. *Birkhchandani, S.; Hirshleifer, D.; Welch, I. Informational Cascades and Observational Learning, The New Palgrave Dictionary of Economics, 2006.*
8. *Bohren, A. Information-Processing Uncertainty in Social Learning. mimeo, 2009 University of San Diego.*
9. *Bramson, M.; Griffeath, D. Clustering and Dispersion Rates for Some Interacting Particle Systems on Z^1 . Annals of Probability, 1980, 8, 183-213.*
10. *Chamley, C.P. Rational Herds: Economic Models of Social Learning; Cambridge University Press, Cambridge, UK, 2004.*
11. *Çelen, B.; Kariv, S. Observational Learning under Imperfect Information. Games and Economic Behavior, 2004, 47, 72-86.*

12. Collander S.; Hörner J. *The Wisdom of the Minority*. *Journal of Economic Theory*, **2009**, 144, 1421-1439.
13. De Marzo, P.; Vayanos, D.; Zwiebel, J. *Persuasion Bias, Social Influence, and Unidimensional Opinions*. *Quarterly Journal of Economics*, **2003**, 118, 909-968.
14. Devenow, A.; Welch, I. *Rational Herding in Financial Economics*. *European Economic Review*, **1996**, 40, 603-615.
15. Duffie, D.; Manso, G. *Information Percolation in Large Markets*. *American Economic Review Papers and Proceedings*, **2007**, 97, 203-209.
16. Durrett, R.; Steiff, J.E. *Fixation Results for Threshold Voter*. *Annals of Probability*, **1993**, 21, 232-247.
17. Ellison G.; Fudenberg D. *Rules of Thumb for Social Learning*. *Journal of Political Economy*, **1993**, 111, 612-643.
18. Eyster, E. and M. Rabin *Naive Herding*, *mimeo, LSE, 2008*.
19. Gale, D. *What Have We Learned from Social Learning*. *European Economic Review*, **1996**, 40, 617-628.
20. Gale, D.; Kariv, S. *Bayesian Learning in Social Networks*. *Games and Economic Behavior*, **2003**, 45, 329-346.
21. Guarino, A.; Harmgart, H.; Huck S. *Aggregate Information Cascades*, *mimeo, 2008, UCL*.
22. Guarino, A.; Jehiel, P. *Social Learning with Coarse Inference*, 2009, *ELSE Working Papers 337*.
23. Hirshleifer, D.; Theo S. *Herd Behaviour and Cascading in Capital Markets: a Review and Synthesis*. *European Financial Management*, **2003**, 9, 25-66.
24. Larson, N. *In with the New Or Out with the Old? Bottlenecks in Social Learning*. *mimeo, 2008, University of Virginia*.
25. Lee, I.H. *On the Convergence of Informational Cascades* *Journal of Economic Theory*, **1993**, 61, 395-411.
26. Liggett, T.M. *Interacting Particle Systems*; Springer-Verlag: Berlin, Germany, 1985.
27. Smith, L.; Soerensen, P. *Rational Social Learning with Random Sampling*, *mimeo, 2008, University of Michigan*.

28. *Vives, X. Information and Learning in Markets; Princeton University Press: Princeton, 2008.*