

**School of Social Sciences
Economics Division
University of Southampton
Southampton SO17 1BJ, UK**

**Discussion Papers in
Economics and Econometrics**

Social Learning with Local Interactions

Antonio Guarino & Antonella Ianni
No. 1011

**This paper is available on our website
<http://www.southampton.ac.uk/socsci/economics/research/papers>**

ISSN 0966-4246

Social Learning with Local Interactions

123

Antonio Guarino
University College London

Antonella Ianni
University of Southampton

June 11, 2010

¹Most of the work was carried out while A. Ianni was at the Università Ca' Foscari di Venezia, (funded by the MIUR Rientro dei Cervelli) and visiting the E.U.I..

²**JEL: 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 a simple dynamic model of social learning with local informational externalities. There is a large population of agents, who repeatedly have to choose one, out of two, reversible actions, each of which is optimal in one, out of two, unknown states of the world. Each agent chooses rationally, on the basis of private information (s)he receives by a symmetric binary signal on the state, as well as the observation of the action chosen among their nearest neighbours. Actions can be updated at revision opportunities that agents receive in a random sequential order. Strategies are stationary, in that they do not depend on time, nor on location.

We show that:

if agents receive equally informative signals, and observe both neighbours, then the social learning process is not adequate and the process of actions converges exponentially fast to a configuration where some agents are permanently wrong;

if agents are unequally informed, in that their signal is either fully informative or fully uninformative (both with positive probability), and observe one neighbour, then the social learning process is adequate and everybody will eventually choose the action that is correct given the state. Convergence, however, obtains very slowly, namely at rate \sqrt{t} .

We relate the findings with the literature on social learning and discuss the property of efficiency of the information transmission mechanism under local interaction.

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 private information that others hold. If there is uncertainty over which decision is more profitable, 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 change our original decision just to follow.

The process of learning from others has first been studied by Bikhchandani et al. (1992), and Banerjee (1992). 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 to simplify 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 and 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 decision, and those of their friends and neighbours, that we cannot 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 a subset of actions taken by others, namely, the actions of their *neighbours*. We are interested in understanding how the local interaction among agents develops in the process of social learning. The first issue that we tackle here is whether the transmission of information through local observation can lead to converge in beliefs and in actions in the whole population. The second is the characterization of the speed at which information is transmitted.

The seminal papers by Bikhchandani et al. (1992), and Banerjee (1992) 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 predecessors. Conformism eventually prevails, and it may well be that the entire population settles on the wrong action and that the beliefs never converge to the truth. It is worth mentioning that this occurs despite the population is only formed of rational, Bayesian 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 converge eventually to the right decision? And in the positive case, will the convergence 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 in society 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.

We depart from the canonical social learning model in different ways. First, as we said, we assume that agents can only observe the behavior of their neighbours, i.e. a subset of other individuals who live close by, on an appropriately characterized spatial structure. Second, we let each agent revise her or his original decision repeatedly, and specifically postulate that updating takes place in a random sequential order. We see these assumptions as particularly appropriate to the set up of our model that involves a large population of individuals, who only have limited information about the environment where they are to take decisions. These different features affect the process of social learning.

In particular, we address two complementary issues. The first concerns the social learning process in terms of its asymptotics. Starting from an initial random configuration of beliefs in the population, will social learning eventually show consensus in the aggregate, or will different actions coexist indefinitely? Is social learning complete, in the sense that beliefs will converge to the truth? Is it at least adequate, in the sense that all agents will choose the action that is optimal given the state of the world? Second, we aim at evaluating the social learning process in terms of its convergence rates. We believe that this issue is particularly relevant when analyzing processes of information transmission and social learning, 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. For the setting we use, we are able to provide analytical solution for the speed of convergence of the dynamic 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) we also study the process of cluster formation.

The analysis of the rate of convergence is particularly important in the study of social learning. It is now clear that the results of the first models of herding on the complete blockage of information (informational cascade) is essentially an extreme example of an informational inefficiency. This type of extreme inefficiency depends on specific features of the models, in particular, on the discreteness of the action space. While this result may not be very robust to perturbations of the model of learning, the finding that the actions by early agents can have a disproportionate effect on the decisions of successors, and that there can be informational inefficiencies in the process of social learning seems, instead, independent of the technicalities of

the model. This finding is, indeed, the main lesson that we learn from this literature (Gale (1996); Chamley (2004); Bikhchandani et al. (2005)). After all, even if there is not a complete blockage of information, but the dissemination of information is very slow, the result is that the society may be choosing the wrong decision for a very long while. Therefore, being able to characterize the speed of learning is of crucial importance for our understanding of the process of information dissemination.

The main result of the paper shows that, in our model of social learning with local interactions, whenever equally informed individuals use optimal stationary strategies, i.e. strategies that only depend on their current information (and not on history, nor on location) learning cannot be complete and the process of actions converges exponentially fast to a configuration where somebody is permanently wrong. The intuition for this is that since information is highly decentralized, it can happen that inward looking groups of agents who received the same signal may continue choosing the same, perhaps incorrect, action, indefinitely. We then study a variation of the model where agents are unequally informed, in the sense that some receive a fully informative signal and some receive a completely uninformative signal (which they rationally disregard). We find that in this case learning is adequate and the process of social learning converges to a state where everybody is choosing the correct action. Convergence, however, obtains very slowly.

The paper is organized as follows. Sections 2 and 3 describe the general framework of the set up. Sections 3.1 and 3.2 analyze the properties of the dynamic processes of learning and contain the main results of the paper. Section 4 relates these findings to the existing literature. Section 5 concludes.

2 The Economy

We consider a set $\Omega = \{0, 1\}$ of possible states of nature. In the economy there is a set X of countably many agents. Each agent $x \in X$ has to choose an action in the set $A = \{0, 1\}$. Time runs continuously and each agent can be called to choose an action more than once. In particular, each individual may choose a new action at a random exponential time, with mean 1. In any small time interval at most one agent can reassess her or his decision, and every agent is equally likely to receive an updating opportunity. As a result of these assumptions, agents choose actions in a random sequential order. We denote the action chosen by individual x at time t by $\eta_t(x) \in A$.

Information

Initially, each agent has uniform priors, i.e., (s)he believes that both states of the world are equally likely. Then, (s)he gathers information on the state of the world in two ways. At time $t = 0$ (s)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 $\sigma \in \{w, s\}$ refers to the signal precision. An agent can receive a strongly-informative signal, with precision $q^s \equiv \Pr[\theta^s(x) = 1 \mid \omega = 1] = \Pr[\theta^s(x) =$

$0 \mid \omega = 0] \in (0.5, 1]$ or a weakly-informative signal with precision $q^w \equiv \Pr[\theta^w(x) = 1 \mid \omega = 1] = \Pr[\theta^w(x) = 0 \mid \omega = 0] \in [0.5, q^s]$. Note that, conditional on a state of nature, the signals that agents receive are independently distributed. The level of information in the economy depends of course on the proportion of agents receiving the strongly-informative or the weakly-informative signal. We denote the probability that each agent receives a strongly-informative signal by r .

Having observed the signal at time 0, each agent makes her or his first choice, $\eta_0(x)$. When the agent receives another opportunity to take a decision (i.e., to revise the choice previously made), (s)he observes the decisions taken by a subset of other agents in the population. This is the second way in which (s)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 (s)he can only interact with the set of agents who live in her or his vicinity. Formally, we take $X \subseteq Z^1$ and define the set of x 's nearest neighbours as $N(x) = \{y : \|y - x\| = 1\}$, i.e., the set of 2 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, as $I(x)$.

Payoff

Agent x has the following payoff function, depending on the action chosen and on the state of nature:

$$U(\eta_t(x), \omega) = \begin{cases} 2\omega - 1 & \text{if } \eta_t(x) = 1, \\ 0 & \text{if } \eta_t(x) = 0. \end{cases} \quad (1)$$

This formulation rules away any potential strategic effect on the part of agents, as the choice of $\eta_t(x)$ only depends on the assessment of the probability that $\omega \in \Omega$ is the true state of the world. On the basis of the information available, x at time t chooses $\eta_t(x)$ to maximize $E[U(\eta_t(x), \omega) \mid I_t(x)]$ and sticks to this decision until a new updating opportunity arises.

If we denote the belief at time t that $\omega = 1$ by $\pi_t(x) = \Pr[\omega = 1 \mid I_t(x)]$, then the optimal strategy for x at t is

$$\eta_t^*(x) = \begin{cases} 1 & \text{if } \pi_t(x) > 0.5, \\ \{0, 1\} & \text{if } \pi_t(x) = 0.5, \\ 0 & \text{if } \pi_t(x) < 0.5. \end{cases} \quad (2)$$

or, equivalently,

$$\eta_t^*(x) = \begin{cases} 1 & \lambda_t(x) > 0, \\ \{0, 1\} & \lambda_t(x) = 0, \\ 0 & \lambda_t(x) < 0. \end{cases} \quad (3)$$

where $\lambda_t(x) \equiv \log \frac{\Pr[\omega=1 \mid I_t(x)]}{\Pr[\omega=0 \mid I_t(x)]}$ denotes the log-likelihood ratio (LLR) for agent x at time t and clearly $\pi_t(x) \equiv \exp[\lambda_t(x)] / [1 + \exp[\lambda_t(x)]]$.

Equilibrium

An equilibrium of a social learning process is a profile of optimal strategies, one for each agent.

Definition 1 (Equilibrium) *An equilibrium of the social learning process is a profile of strategies $\{\eta^*(x)\}_{x \in X}$ such that, for all $x \in X$, $\eta^*(x) : I(x) \rightarrow \{0, 1\}$ and $\eta^*(x) \in \arg \max_{\eta(x) \in \{0, 1\}} E[U(\eta(x), \omega) | I(x)]$.*

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 can revise it whenever (s)he has an updating opportunity. Let $\{\tau_{x_l}\}$, $l = 0, 1, 2, \dots$ be the sequence of times when x receives an updating opportunity, with $\tau_{x_0} \equiv 0$. Furthermore, let $\lambda_{\tau_{x_l}}(x)$ denote x 's LLR at time τ_{x_l} . Given $\lambda_{\tau_{x_l}}(x)$, agent x will optimally choose an action according to (2). Agent x 's dynamics of choices is then given by

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

We shall regard any state of the learning process as $\eta \in \{0, 1\}^X$ and we denote by $|\eta|$ the number of agents choosing action 1 in state η . We shall denote by η_t the state of the stochastic dynamic process at time t and we are interested in characterizing its evolution over time and over space. At the beginning of time, "Nature" chooses ω 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 the properties of these dynamics.

3 Social Learning Processes

Before proceeding, 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 in our framework. This will allow us to motivate the specific assumptions that lead to our results and choose appropriate benchmarks.

Consider the standard model of sequential social learning proposed by Bickchandani et al. (1992) and suppose that each agent can directly observe the signals received by the others. Suppose that $q^s = q^w \equiv q \in (0.5, 1)$, i.e., that each agent x receives a signal of the same precision to compute her or his initial LLR as:

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

Suppose that agent x could directly observe the private signals received by all other agents y , and for the sake of the argument suppose agents are numbered $1, 2, \dots, n, \dots$. The updated LLR on the basis of this information set $I(n) = \{\theta(n), \theta(y) \text{ for } y \leq n\}$ will be:

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

Notice that the order in which other agents' signals are observed does not matter, in that observations are exchangeable. Assume that the true state of the world is $\omega = 0$. Then, by assumption, the random variables $\theta(\cdot)$ have mean $1 - q$ and, as a result, $E[\lambda(\cdot) \mid \omega = 0] = 1 - 2q$ is strictly negative. This implies that $\lambda(n)$ tends to $-\infty$, as agent n observes all other agents' signals. 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 sequentially observe all signals, learning is *complete* in the sense that beliefs converge to the truth (and, hence, their actions to the correct decisions), i.e., $\Pr[\lim_{n \rightarrow \infty} \pi(n) = \eta(n) = \omega] = 1$ for all n . In particular, convergence obtains exponentially fast.

As a result, in a typical model of social learning, the observability of private information is by itself enough to guarantee that the outcome of the learning process is informationally efficient. Actions are clearly not as informative as signals are in this setting: potential inefficiencies may arise when agents only observe the actions taken by others and not the signals that led to those choices. Indeed, in the classical model of sequential social learning these inefficiencies 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, i.e. assume that $I(x) = \{\theta(x), \theta(x \pm 1)\}$. Since information is limited, agents' beliefs, as measured by their LLR, are necessarily bounded, as each agent can observe at most three signals. As a result, we cannot expect any convergence in beliefs to the truth in this case: since information is highly decentralized, it can happen that inward looking groups of agents who received the same signal may continue choosing the same action (perhaps the incorrect one) indefinitely.

Essentially, in a model of social learning with local interactions, the public observability of the neighbours' signals is not enough to achieve efficiency in the process of information transmission. Observing actions may in fact improve efficiency, as actions taken by one's neighbour may convey information on signals received by that neighbour's neighbours and so on. This transfer of information can potentially trickle to and from any agent and enhance the efficiency of the social learning process. However, unlike signals, actions necessarily show a degree of correlation which agents ought to rationally account for, when making their choices.

In order for this spatial correlation to be recognized and fully exploited in the inference process, one may think that endowing agents with the ability to convey more information might suffice. In fact, it turns out that this could be the case only in the very special case of the following example. Suppose agents are numbered $1, 2, \dots, n, \dots$, updating opportunities are assigned sequentially, in exactly that given order, and each agent can only observe her or his predecessor. Suppose further that, when observed, agents could also show their LLR, or in other words, their posterior,

i.e. $I(n) = \{\lambda(n-1), \theta(n)\}$. As LLR are additive, clearly:

$$\lambda(n) = \lambda(n-1) + 2 \left(\log \frac{q}{1-q} \right) \left(\theta(n) - \frac{1}{2} \right)$$

and by exactly the same logic used at the beginning of this Section, one can show that the social learning process is *complete*, in the sense that beliefs converge to the truth. The key feature of this example is that $\lambda(n-1)$ constitutes a sufficient statistic of all the signals upon which it is based (i.e. all the signals that agent $n-1$ has observed before being observed by agent n). This is no longer true in a locally interactive model where agents update repeatedly and in a random order, and where, by analogy, $I(x) = \{\lambda(x \pm 1), \theta(x)\}$. In fact, in this case, the information contained in $\lambda(x \pm 1)$, which in principle could indirectly reveal information on $x \pm 2$, $x \pm 3$ and so forth, would already contain the information revealed by $\theta(x)$ through the observation of x on the part of $x \pm 1$. In essence, this is due to the fact that, whenever interaction is local, information sets are not disjoint¹.

In order for agents to be able to correctly draw inference upon the observation of their neighbours' action, one should endow each agent with a very rich information set, that should include not only the history of actions chosen in their neighbourhood, but also, the exact order with which updating opportunities have been assigned until that point. Although our model assumes that agents are able and willing to perform Bayesian updating, we take the view that these requirements are unreasonable in the set up of a large population of agents. As a result, we introduce the following modeling assumption.

Definition 2 (Limited Memory) *If at time t agent x gets an opportunity, (s)he observes the action taken in her neighbourhood at that time. As a result, at each revision opportunity, agent x takes decisions on the basis of the following information set:*

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

It is important to note that we are assuming that the agent observes only her or his neighbours' current decisions. (S)he has limited memory, in that she remembers her signal, but she possibly forgets past decisions and observations. As, by construction, information is limited both over space (only nearest neighbours can be seen) and over time (only current actions can be observed), we focus on strategies that are *stationary*:

$$\eta(x) : I(x) \rightarrow \{0, 1\}$$

¹A very neat analysis of a model where agents repeatedly encounter and fully reveal their posterior beliefs to each other over time is provided in Duffie et al. (2009). Their model does not have a spatial dimension and the authors are able to characterize explicitly the convergence of the cross-sectional distribution of beliefs to a common posterior.

where, for each t , we regard $I_t(x) \equiv I(x) \subseteq \{0, 1\}^3$.

We are interested in characterizing the properties of social learning processes with local interactions in terms of the degree of informational efficiency they achieve. For the reasons we mentioned above, requiring complete learning would be too demanding in a model of social learning with local interactions, since beliefs are, by construction, bounded, and the local interactive structure of the model embeds an amount of spatial correlation in actions. A weaker requirement is, instead, that all agents eventually choose the correct action, as in the following definition:

Definition 3 (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, $0 \leq \lim_{t \rightarrow \infty} \Pr[\eta_t(x) = \omega] \leq 1$ is the limit measure of agents who are correct (in the sense that they choose the action appropriate for the true state of the world) and it can be thought of as a measure of how *informationally* efficient the social learning process is.

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

3.1 Equally Informed Agents

In this SubSection, we focus on the case in which $q^s = q^w \equiv q \in (0.5, 1)$. In this case, at the beginning of time, each agent receives a signal commonly known to have precision $q > 0.5$. Note that this is the only case contemplated by the standard model of social learning.

The next result shows that in this set up, the use of optimal stationary strategies leads to extreme inefficiency in the process of information transmission under local interactions.

Theorem 4 *If each agent x receives a symmetric binary signal with precision $q > 0.5$ and can observe both neighbours, then the process of social learning with local interactions 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. The logic of the proof relies on the explicit characterization of the process of inference underlying agents' optimal choices. We recall that, in our model, *stationary* strategies are of the form $\eta(x) : \{0, 1\}^3 \rightarrow \{0, 1\}$, where x chooses on the basis of the configuration of actions within her or his neighbourhood, as well as on the realization of their signal, i.e.

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

where τ_x denote any random time at which x receives an updating opportunity.

Let τ^i ($i = 1, 2, \dots$) denote 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 or her choice, τ^2 is the second time, etc. Consider the following stationary strategy:

S1: *At any time τ_x :*
 if $\eta(x-1) = \eta(x+1)$, then choose $\eta(x) = \eta(x \pm 1)$
 if $\eta(x-1) \neq \eta(x+1)$, then stick to $\eta(x)$

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, i.e., for any $\tau_x = \tau^i$, and at any stage of the revision process, i.e., for any τ_{x_l} , $l = 1, 2, \dots$

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. To see this, note that at time $t = 0$, upon receiving the signal $\theta(x)$, 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, agent x 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, Remark 7 shows that this strategy is optimal as long as it is followed by all other agents. On the basis of this result, Remark 8 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 9 emphasizes that this social learning process with local interactions gives rise to an extreme form of informational inefficiency.

The above result shows that if agents use optimal stationary strategies the probability that the whole population learns to behave optimally, given the state of the world, is zero. The process of social learning may in fact 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 see this, 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 they do, none 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 by analyzing a model where agents have relatively less information about their neighbours, but are heterogeneous in terms of the quality of their private information, with some agents being perfectly informed and some perfectly uninformed. One implication of these assumptions is that those agents who are aware of being uninformed will disregard their private information and be more prone to changing actions. As we shall show below this modeled heterogeneity significantly improves the efficiency of the mechanism of information transmission.

3.2 Unequally Informed Agents

We now move to a different scenario, in which agents receive signals of a different precision. 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 $q^s = 1$, $q^w = 0.5$ and that the probability that each agent receives a strongly-informative signal or a weakly-informative signal are both positive, i.e., $r \in (0, 1)$. Also, we assume that at each time x is to take a decision, (s)he observes the action currently chosen by only one of her or his two neighbours, drawn at random 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 5 *If each agent is perfectly informed with probability r (and perfectly uninformed otherwise), and can observe one neighbour (randomly chosen), 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 the previous Section. 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,

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

Consider the following stationary strategy for agent x :

S2: *At any time τ_x : choose $\eta(x) = \eta(y)$*

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 (s)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 11 shows that strategy S2 is optimal for any agent x , as long as it is followed by all other uninformed agents. Remark 12 characterizes the limit behaviour of this social learning process and Remark 13 evaluates the degree of informational efficiency of this model.

4 Related Literature

The literature on social learning has been growing very fast over the last decades. A variety of models have been used to shed light on the way in which information is transmitted among economic agents who have to take decisions under uncertainty. Within this literature, social learning refers to the fact that agents learn from observing actions taken by other individuals. Early models known as herding and informational cascades show that individuals who take choices sequentially and observe the choices made by others, may actually ignore their private information and base their decisions entirely on what is publicly observed. Herds may occur because the informational content of the history of choices of agents in the economy overwhelms the information contained in their private signals. Since in this case agents' information is not revealed through their actions, the social learning process may never converge to the truth.

Recent contributions to the theory of social learning have extended these models to account for situations where agents do not observe the entire history of the actions chosen. Though the amount of information they receive is limited in this case, the way in which this is disseminated in society 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. The overall effect on social learning is hence unclear and results to date in the literature on social learning with limited memory lead to different conclusions.

The intuition that a small sample size may be efficient because it can enhance the diffusion of private information is formalized in Smith and Soerensen (2008), who study a model of sequential herding with random sampling. They show that

under some circumstances (unbounded beliefs) learning is complete, while under other circumstances (bounded beliefs) informational cascades can occur in finite time.

The consideration that more information is revealed in a larger sample sizes is instead dominant in Banerjee and Fudenberg (2004), who analyze a model of rational “Word of Mouth” communication in a large population. They show that under such conditions learning is complete *even* if beliefs are bounded if the sample size is at least 2 and, furthermore, private information becomes irrelevant if the sample size is greater than 3.

Imperfect information in this set up, might however lead beliefs and actions to cycle forever (see, for example, Celen and Kariv (2004)).

Gale and Kariv (2004) analyze a situation where observability is limited to the actions taken by their neighbours, as agents belong to a social network and can only observe the decisions of the other agents to whom they are connected. 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. (2010) analyze a model in which agents observe past actions of a stochastically-generated neighbourhood of individuals, where each agent knows the identity of their neighbours. They show that, when beliefs are unbounded and there is some minimal amount of *expansion in observations*, asymptotic learning obtains (meaning that actions converge to the correct one). The authors also provide conditions under which the same is true even when private beliefs are bounded, for a large class of stochastic network topologies.

Finally, the issue of imperfect observability is also discussed in Eyster and Rabin (2008) and in Guarino and Jehiel (2009) in contexts in which agents are not fully rational. The fact that observability is imperfect can actually alleviate some biases that bounded rationality produces in a classical model of learning with a continuous action space (as in Lee (1992)).

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, i.e. 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.

Neither the heterogeneity in the quality of private information, nor the existence of someone who is perfectly informed, are, per se, sufficient to guarantee that the social learning process is adequate. To see this, consider the model of Section 3.1 and suppose $1 \geq q^s > q^w \in (0.5, q^s)$, i.e. assume that all agents receive an informative signal, but some agents are better informed than others. We conjecture that for a non empty set of parameters (q^s, q^w) in this range the result of Theorem 4 would carry on in this case as well. In fact, consider the extreme case in which the better informed agents receive a perfectly informative signal, i.e. $q^s = 1$. Any such an agent, say for example agent 0, knows the true state of the world and sticks to $\eta(0) = \omega$ independently of the actions adopted in her or his neighbourhood. This relevant information cannot however be transmitted to others, in that, since agents ± 1 are themselves informed, it could be enough for them to have their 'other' neighbour, ± 2 respectively, be choosing the wrong action, $1 - \omega$, in order for them to be permanently wrong in their choice of actions. 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.

Heterogeneity of information, coupled with the existence of some perfectly uninformed agents, i.e. for $1 > q^s > q^w = 0.5$ is not per se sufficient to guarantee that the social learning process is adequate. To see this consider the model of Section 3.2: it can be shown that strategy S2 would still be optimal for the uninformed agents in this case. Also, for a non empty set of values of q^s in this range, a better informed agent, say 0, would play her or his signal independently of their neighbours. As a result, with positive probability uninformed agents ± 1 will learn agent 0's action,

and transmit it to agents ± 2 etc. The social learning process could however fail to be adequate, because Theorem 10 may not hold, since with positive probability some better informed agents could receive the wrong signal and never correct their initial decision.

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 poorly informed, are willing to learn it. In this case, a process of slow clustering on the correct decision ensues. Our explicit characterization of the convergence rates shows that, perhaps surprisingly, the speed at which this cluster grows does not depend on the proportion of perfectly informed agents in the population: the estimate of which in (6) holds in fact unaltered for any value of the parameter $0 \leq r < 1$ ².

²Clearly, 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.

Appendix A

Theorem 6 *If each agent receives a symmetric binary signal with precision $q > 0.5$, and can observe both neighbours, then the process of social learning with local interactions 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 split into a few Remarks: Remark 7 shows that the model admits an equilibrium; Remark 8 characterizes limit behaviour and convergence rates of this process of social learning with local interactions and finally Remark 9 evaluates the degree of informational efficiency of the process.

Remark 7 *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 as proved above, 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 (s)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,

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

If $\tau_{(x-1)_1} > \tau_{x_1}$, this probability is one, as $x-1$ is playing her or 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 (s)he herself is choosing, infers that neighbour is playing her 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 (s)he herself is choosing, infers that the neighbour is more likely to be playing her 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:

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

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 $E[\theta(x-1) \mid I(x)] < 0.5$. Hence:

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

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 $E[\theta(x-1) \mid I(x)] < 0.5$ and $E[\theta(x+1) \mid I(x)] = 1$. Hence:

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

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 8 *If agents use strategy S1, the characterization of the limit behaviour of the social learning process is as follows³.*

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 7, 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 $X = \mathbb{Z}^1 = \{\dots, -2, -1, 0, +1, +2, \dots\}$. Define the following function:

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

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 simplicity⁴ take $x = 0$, with neighbours $y \in \{-1, +1\}$:

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

³In 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.

⁴This is done WLOG, since the initial distribution that determines the initial condition is (a product measure and hence) translation invariant.

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

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

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

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

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 (3) and learning is not adequate. ■

Appendix B

Theorem 10 *If each agent is perfectly informed with probability r and perfectly uninformed otherwise, and can only 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 of the Theorem is split into three Remarks: Remark 11 shows that an equilibrium exists, Remark 12 characterizes the limit behaviour and Remark 13 evaluates the degree of informational efficiency of this model. The logic of the proof parallels that of the previous Section.

Remark 11 *Suppose all agents $y \neq x$ choose strategy S2. Then this strategy is also optimal for x at any τ_{x_l} .*

Proof. We follow exactly the same logic as that of Remark 7 and describe the process of inference undertaken by agent x at time τ_{x_l} , (we drop the time subscript for notational convenience). Notice that, within this model, agent x cannot draw any inference from his or her signal. Also, agent x cares about his or her 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:

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

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 12 *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:

$$c(\mathbf{x}, \eta) = 0 \tag{4}$$

$$c(x, \eta) = \begin{cases} \frac{1}{2} \sum_{y \in \{x \pm 1\}} \eta(y) & \eta(x) = 0 \\ \frac{1}{2} \sum_{y \in \{x \pm 1\}} (1 - \eta(y)) & \eta(x) = 1 \end{cases} \tag{5}$$

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, i.e. 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 attractive⁵, so are the processes c_i^N and, by Theorem 2.7 in Liggett (1985):

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

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

⁵We 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:

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

given by the speed with which those uninformed agents who are choosing the incorrect action, flip to the correct one. Hence we need to study the dynamics of choices of the individuals in \tilde{X} . We notice that this dynamics is analog to that of the Voter's model (Liggett (1985), Section 1 and 3, Chapter V or in Bramson and Griffeath (1980)), 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 (4) 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 (1980). 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):

$$\sqrt{\pi} \left(\frac{1}{2^{\frac{1+r}{2} \frac{1-r}{2}}} \right) \leq \lim_{t \rightarrow \infty} E \left[\frac{C^{\mu_\omega}(\eta_t)}{t^{1/2}} \right] \leq 2 \left(\frac{\left(\frac{1+r}{2}\right)^2 + \left(\frac{1-r}{2}\right)^2}{\left(\frac{1+r}{2}\right) \left(\frac{1-r}{2}\right)} \right) \sqrt{\pi} \quad (6)$$

where $\pi = 3.1416$. ■

Remark 13 *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

- ACEMOGLU, D. M. DAHLEH, I. LOBEL AND A. OZDAGLAR (2010), "Bayesian Learning in Social Networks," , mimeo, MIT.
- BANERJEE, A.V. (1992), "A Simple Model of Herd Behavior," *Quarterly Journal of Economics*, **107**, 797-817.
- BANERJEE, A. AND D. FUDENBERG (2004), "Word of Mouth Learning," *Games and Economic Behavior*, **46**, 1-22.
- BIRKCHANDANI, S., D. HIRSHLEIFER AND I. WELCH (1992), "A Theory of Fads, Fashions, Custom, and Cultural Change as Information Cascades," *Journal of Political Economy*, **100**, 992-1026.
- BIRKCHANDANI, S., D. HIRSHLEIFER AND I. WELCH (2006), "*Informational Cascades and Observational Learning*," The New Palgrave Dictionary of Economics.
- BRAMSON M. AND D. GRIFFEATH (1980), "Clustering and Dispersion Rates for Some Interacting Particle Systems on Z^1 ," *Annals of Probability*, **8**, 183-213.
- CHAMLEY, C.P. (2004), "*Rational Herds: Economic Models of Social Learning*," Cambridge University Press.
- CHEW, M. S. Y. (2000), "Communication and Coordination in Social Networks," *The Review of Economic Studies*, **67**, 1-16.
- CELEN, B. AND S. KARIV (2004), "Observational Learning under Imperfect Information," *Games and Economic Behavior*, **47**, 72-86.
- DEVENOW, A. AND I. WELCH (1996), "Rational Herding in Financial Economics," *European Economic Review*, **40**, 603-615.
- DUFFIE, D. AND G. MANSO (2007), "Information Percolation in Large Markets," *American Economic Review Papers and Proceedings*, **97**, 203-209.
- DURRETT, R. AND J.E. STEIFF (1993), "Fixation Results for Threshold Voter," *Annals of Probability*, **21**, 232-247.
- EYSTER, E. AND M. RABIN (2008), "Naive Herding," *mimeo*, LSE.
- GALE, D. (1996), , "What Have We Learned from Social Learning," *European Economic Review*, **40**, 617-628.
- GALE, D. AND S. KARIV (2003), "Bayesian Learning in Social Networks," *Games and Economic Behavior*, **45**, 329-346.

- GUARINO, A. AND P. JEHIEL (2009), "Social Learning with Coarse Inference," , ELSE Working Papers 337.
- LIGGETT, T.M. (1985), *Interacting Particle Systems*, Springer-Verlag.
- SMITH, L. AND P. SOERENSEN (2008), "Rational Social Learning with Random Sampling," *mimeo*, University of Michigan.