

Bayesian Learning in Social Networks

DARON ACEMOGLU and MUNTHER A. DAHLEH

Massachusetts Institute of Technology

ILAN LOBEL

New York University

and

ASUMAN OZDAGLAR

Massachusetts Institute of Technology

First version received November 2008; final version accepted December 2010 (Eds.)

We study the (perfect Bayesian) equilibrium of a sequential learning model over a general social network. Each individual receives a signal about the underlying state of the world, observes the past actions of a stochastically generated neighbourhood of individuals, and chooses one of two possible actions. The stochastic process generating the neighbourhoods defines the network topology. We characterize pure strategy equilibria for arbitrary stochastic and deterministic social networks and characterize the conditions under which there will be asymptotic learning—convergence (in probability) to the right action as the social network becomes large. We show that when private beliefs are unbounded (meaning that the implied likelihood ratios are unbounded), there will be asymptotic learning as long as there is some minimal amount of “expansion in observations”. We also characterize conditions under which there will be asymptotic learning when private beliefs are bounded.

Key words: Information aggregation, Learning, Social networks, Herding, Information cascades

JEL Codes: C72, D83

1. INTRODUCTION

How is dispersed and decentralized information held by a large number of individuals aggregated? Imagine a situation in which each of a large number of individuals has a noisy signal about an underlying state of the world. This state of the world might concern, among other things, earning opportunities in a certain occupation, the quality of a new product, the suitability of a particular political candidate for office or pay-off-relevant actions taken by the government. If signals are unbiased, the combination—aggregation—of the information of the individuals will be sufficient for the society to “learn” the true underlying state. The above question can be formulated as the investigation of what types of behaviours and communication structures will lead to this type of information aggregation.

Condorcet’s Jury theorem provides a natural benchmark, where sincere (truthful) reporting of their information by each individual is sufficient for aggregation of information by a law of large numbers argument (Condorcet, 1785). Against this background, a number of papers, most

notably Bikhchandani, Hirshleifer and Welch (1992), Banerjee (1992), and Smith and Sorensen (2000), show how this type of aggregation might fail in the context of the (perfect) Bayesian equilibrium of a dynamic game: when individuals act sequentially and observe the actions of all previous individuals (agents), there may be “herding” on the wrong action, preventing the efficient aggregation of information.

An important modelling assumption in these papers is that each individual observes all past actions. In practice, individuals are situated in complex social networks, which provide their main source of information.¹ In this paper, we address how the structure of social networks, which determines the information that individuals receive, affects equilibrium information aggregation in a sequential learning environment.

As a motivating example, consider a group of consumers deciding which one of two possible new smartphones to switch to. Each consumer makes this decision when her existing service contract expires, so that there is an exogenous sequence of actions determined by contract expiration dates. Each consumer observes the choices of some of her friends, neighbours and coworkers, and the approximate timing of these choices (as she sees when they start using the new phone). She does not, however, observe who else these friends, neighbours, and coworkers have themselves observed. Additional information from direct communication is limited since consumers cannot easily identify and communicate their valuations soon after a new purchase. A related example would be the choice of a firm between two new technologies after the (exogenous) breakdown of its current technology. The firm observes the nature and timing of the choices of some of the nearby firms, but not what other information these firms had at the time they made their decisions. In this case also, the main source of information would be observation of past actions rather than direct communication.

Similar to these examples, in our model, a large number of agents sequentially choose between two actions. An underlying state determines the pay-offs of these two actions. Each agent receives a signal on which of these two actions yield a higher pay-off. Preferences of all agents are aligned in the sense that, given the underlying state of the world, they all prefer the same action. In addition to his own signal, each agent observes the choices of others in a stochastically generated neighbourhood. In line with the examples above, each individual knows the identity of the agents in his neighbourhood. But he does not observe their private signal or necessarily know what information these agents had access to when making their own decisions.

More formally, this dynamic game of incomplete information is characterized by two features: (1) the signal structure, which determines how informative the signals received by the individuals are and (2) the (social) network topology. We represent the network topology by a sequence of probability distributions (one for each agent) over subsets of past actions. The environment most commonly studied in the previous literature, the full observation network topology, is the special case where all past actions are observed. Another deterministic special case is the network topology where each agent observes the actions of the most recent $M \geq 1$ individuals. Other relevant networks include stochastic topologies in which each agent observes a random subset of past actions, as well as those in which, with a high probability, each agent observes the actions of some “influential” group of agents, who may be thought of as “informational leaders” or the media. In addition to these examples, our representation of social networks is sufficiently general to nest several commonly studied models of stochastic networks, including small-world models, and observation structures in which each individual sees the actions of one

1. Granovetter (1973), Montgomery (1991), Munshi (2003), and Ioannides and Loury (2004) document the importance of information obtained from the social network of an individual for employment outcomes. Besley and Case (1994), Foster and Rosenzweig (1995), Munshi (2004), and Udry and Conley (2001) show the importance of the information obtained from social networks for technology adoption. Jackson (2006, 2007) provides excellent surveys.

or several agents randomly drawn from the entire or the recent past. Our representation also does not impose any restriction on the degree distribution (cardinality) of the agents' neighbourhoods or the degree of clustering in the network.

We provide a systematic characterization of the conditions under which there will be asymptotic learning in this model. We say that there is *asymptotic learning* if as the size of the society becomes arbitrarily large, equilibrium actions converge (in probability) to the action that yields the higher pay-off. Conversely, asymptotic learning fails if, as the society becomes large, the correct action is not chosen (or more formally, the liminf of the probability that the right action is chosen is strictly less than 1).

Two concepts turn out to be central in the study of sequential learning in social networks. The first is whether the likelihood ratio implied by individual signals is always finite and bounded away from 0.² Smith and Sorensen (2000) refer to beliefs that satisfy this property as bounded (private) beliefs. With bounded beliefs, there is a maximum amount of information in any individual signal. In contrast, when there exist signals with arbitrarily high and low likelihood ratios, (private) beliefs are unbounded. Whether bounded or unbounded beliefs provide a better approximation to reality is partly an interpretational and partly an empirical question. The main result of Smith and Sorensen is that when each individual observes all past actions and private beliefs are unbounded, information will be aggregated and the correct action will be chosen asymptotically. In contrast, the results in Bikhchandani, Hirshleifer and Welch (1992), Banerjee (1992), and Smith and Sorensen (2000) indicate that with bounded beliefs, there will not be asymptotic learning (or information aggregation). Instead, as emphasized by Bikhchandani, Hirshleifer and Welch (1992) and Banerjee (1992), there will be "herding" or "informational cascades", where individuals copy past actions and/or completely ignore their own signals.

The second key concept is that of a network topology with expanding observations. To describe this concept, let us first introduce another notion: a finite group of agents is *excessively influential* if there exists an infinite number of agents who, with probability uniformly bounded away from 0, observe only the actions of a subset of this group. For example, a group is excessively influential if it is the source of all information (except individual signals) for an infinitely large component of the social network. If there exists an excessively influential group of individuals, then the social network has *non-expanding observations*, and conversely, if there exists no excessively influential group, the network has *expanding observations*. This definition implies that most reasonable social networks have expanding observations, and in particular, a minimum amount of "arrival of new information" in the social network is sufficient for the expanding observations property.³ For example, the environment studied in most of the previous work in this area, where all past actions are observed, has expanding observations. Similarly, a social network in which each individual observes one uniformly drawn individual from those who have taken decisions in the past or a network in which each individual observes her immediate neighbour all feature expanding observations. A simple, but typical, example of a network with non-expanding observations is the one in which all future individuals only observe the actions of the first $K < \infty$ agents.

Our main results in this paper are presented in four theorems. In particular, Theorems 2 and 4 are the most substantive contributions of this paper. Theorem 1 shows that there is no asymptotic learning in networks with non-expanding observations. This result is not surprising since information aggregation is not possible when the set of observations on which (an infinite subset of) individuals can build their decisions remains limited forever.

2. The likelihood ratio is the ratio of the probabilities or the densities of a signal in one state relative to the other.

3. Here, "arrival of new information" refers to the property that the probability of each individual observing the action of some individual from the recent past converges to one as the social network becomes arbitrarily large.

Theorem 2 shows that when (private) beliefs are unbounded and the network topology is expanding, there will be asymptotic learning. This is a strong result (particularly if we consider unbounded beliefs to be a better approximation to reality than bounded beliefs) since almost all reasonable social networks have the expanding observations property. This theorem, *e.g.* implies that when some individuals, such as “informational leaders”, are overrepresented in the neighbourhoods of future agents (and are thus “influential”, though not excessively so), learning may slow down, but asymptotic learning will still obtain as long as private beliefs are unbounded.

Theorem 3 presents a partial converse to Theorem 2. It shows that for many common deterministic and stochastic networks, bounded private beliefs are incompatible with asymptotic learning. It therefore generalizes existing results on asymptotic learning, *e.g.* those in Bikhchandani, Hirshleifer and Welch (1992), Banerjee (1992), and Smith and Sorensen (2000), to general networks.

Our final main result, Theorem 4, establishes that there is asymptotic learning with bounded private beliefs for a sizable class of stochastic network topologies. Suppose there exists a subset S of the agents such that agents in S have probability $\varepsilon > 0$ of observing the entire history of actions and that an infinite subset of the agents in S makes decisions (partially) based on their private signals because they have neighbours whose actions are not completely informative. The remaining agents in the society have expanding observations with respect to S (in the sense that they are likely to observe some recent actions from S). Then, Theorem 4 shows that asymptotic learning occurs for any distribution of private signals.

This result is particularly important since it shows how moving away from simple network structures, which have been the focus of prior work, has major implications for equilibrium learning dynamics. In particular, the network structure that leads to asymptotic learning even in the absence of strong signals is a combination of three features: (1) the presence of a set of agents acting based on their private signals, (2) the presence of a different set of agents observing all actions who can thus piece together the state of the world from the actions of the agents that act according to their signals, and (3) a variant of expanding observations that ensures the existence of information paths the agents that have access to the entire history of actions to most other agents.

Our paper contributes to the large and growing literature on social learning. Bikhchandani, Hirshleifer and Welch (1992) and Banerjee (1992) started the literature on learning in situations in which individuals are Bayesian and observe past actions. Smith and Sorensen (2000) provide the most comprehensive and complete analysis of this environment. Their results and the importance of the concepts of bounded and unbounded beliefs, which they introduced, have already been discussed in the introduction and will play an important role in our analysis in the rest of the paper. Other important contributions in this area include, among others, Welch (1992), Lee (1993), Chamley and Gale (1994), and Vives (1997). An excellent general discussion is contained in Bikhchandani, Hirshleifer and Welch (1998). These papers typically focus on the special case of full observation network topology in terms of our general model.

The two papers most closely related to ours are Banerjee and Fudenberg (2004) and Smith and Sorensen (2008). Both of these papers study social learning with sampling of past actions. In Banerjee and Fudenberg, there is a continuum of agents and the focus is on proportional sampling (whereby individuals observe a “representative” sample of the overall population). They establish that asymptotic learning is achieved under mild assumptions as long as the sample size is no smaller than two. The existence of a continuum of agents is important for this result since it ensures that the fraction of individuals with different posteriors evolves deterministically. Smith and Sorensen, on the other hand, consider a related model with a countable number of agents. In their model, as in ours, the evolution of beliefs is stochastic. Smith and Sorensen provide conditions under which asymptotic learning takes place.

A crucial difference between the study of Banerjee and Fudenberg and Smith and Sorensen, on the one hand, and our work, on the other, is the information structure. These papers assume that “samples are unordered” in the sense that individuals do not know the identity of the agents they have observed. In contrast, as mentioned above, our setup is motivated by a social network and assumes that individuals have stochastic neighbourhoods, but know the identity of the agents in their realized neighbourhood. We view this as a better approximation to sequential learning in social networks. While situations in which an individual observes the actions of “strangers” would naturally correspond to “unordered samples”, in most social networks situations, individuals have some idea about where others are situated in the network. For example, an individual would have some idea about which of their friends and coworkers are likely to have observed the choices of many others, which in the context of our model corresponds to “ordered samples”. In addition to its descriptive realism, this assumption leads to a sharper characterization of the conditions under which asymptotic learning occurs. For example, in the environment of Smith and Sorensen, asymptotic learning fails whenever an individual is “oversampled”, in the sense of being overrepresented in the samples of future agents. In contrast, in our environment, asymptotic learning occurs when the network topology features expanding observations (and private beliefs are unbounded). Expanding observations is a much weaker requirement than “non-oversampling”. For example, when each individual observes Agent 1 and a randomly chosen agent from his predecessors, the network topology satisfies expanding observations, but there is oversampling.⁴

Other recent work on social learning includes Celen and Kariv (2004), who study Bayesian learning when each individual observes his immediate predecessor, Callander and Horner (2009), who show that it may be optimal to follow the actions of agents that deviate from past average behaviour, and Gale and Kariv (2003), who generalize the pay-off equalization result of Bala and Goyal (1998) in connected social networks (discussed below) to Bayesian learning.⁵

The second branch of the literature focuses on non-Bayesian learning, typically with agents using some reasonable rules of thumb. This literature considers both learning from past actions and from pay-offs (or directly from beliefs). Early papers in this literature include Ellison and Fudenberg (1993, 1995), which show how rule-of-thumb learning can converge to the true underlying state in some simple environments. The papers most closely related to our work in this genre are Bala and Goyal (1998, 2001), DeMarzo, Vayanos and Zwiebel (2003), and Golub and Jackson (2010). These papers study non-Bayesian learning over an arbitrary connected social network. Bala and Goyal (1998) establish the important and intuitive pay-off equalization result that, asymptotically, each individual must receive a pay-off equal to that of an arbitrary individual in his “social network” since otherwise he could copy the behaviour of this other individual. Our paper can be viewed as extending the results of Bala and Goyal to a sequential setting with Bayesian learning. A similar “imitation” intuition plays an important role in our proof of asymptotic learning with unbounded beliefs and unbounded observations. A key difference between our results and those in the studies of Bala and Goyal; DeMarzo, Vayanos, and Zwiebel; and Golub and Jackson concerns the effects of “influential” groups on learning. In these non-Bayesian papers, an influential group or what Bala and Goyal refer to as a royal family, which is highly connected to the rest of the network, prevents aggregation of information. In contrast, in our model, Bayesian updating ensures that such individuals do not receive disproportionate

4. This also implies that, in the terminology of Bala and Goyal, a “royal family” precludes learning in the model of Smith and Sorensen, but as we show below, not in ours.

5. Gale and Kariv show that the behavior of agents in a connected social network who repeatedly act according to their posterior beliefs eventually converge. This does not constitute asymptotic learning according to our definition since behavior does not necessarily converge to the optimal action given the available information.

weight and their presence does not preclude efficient aggregation of information. Only when there is an excessively influential group, *i.e.* a group that is the *sole* source of information for an infinite subset of individuals, that asymptotic learning breaks down.⁶

The rest of the paper is organized as follows. Section 2 introduces our model. Section 3 characterizes the (pure strategy) perfect Bayesian equilibria and introduces the concepts of bounded and unbounded beliefs. Section 4 presents our main results, Theorems 1–4, and discusses some of their implications (as well as presenting a number of corollaries to facilitate interpretation). Section 5 concludes. Appendix A contains the main proofs, while the Appendix B in Supplementary Material contains additional results and omitted proofs.

2. MODEL

A countably infinite number of agents (individuals), indexed by $n \in \mathbb{N}$, sequentially make a single decision each. The pay-off of agent n depends on an underlying state of the world θ and his decision. To simplify the notation and the exposition, we assume that both the underlying state and decisions are binary. In particular, the decision of agent n is denoted by $x_n \in \{0, 1\}$ and the underlying state is $\theta \in \{0, 1\}$. The pay-off of agent n is

$$u_n(x_n, \theta) = \begin{cases} 1 & \text{if } x_n = \theta, \\ 0 & \text{if } x_n \neq \theta. \end{cases}$$

Again to simplify notation, we assume that both values of the underlying state are equally likely, so that $\mathbb{P}(\theta = 0) = \mathbb{P}(\theta = 1) = 1/2$.

The state θ is unknown. Each agent $n \in \mathbb{N}$ forms beliefs about this state from a private signal $s_n \in \bar{S}$ (where \bar{S} is a metric space or simply a Euclidean space) and from his observation of the actions of other agents. Conditional on the state of the world θ , the signals are independently generated according to a probability measure \mathbb{F}_θ . We refer to the pair of measures $(\mathbb{F}_0, \mathbb{F}_1)$ as the signal structure of the model. We assume that \mathbb{F}_0 and \mathbb{F}_1 are absolutely continuous with respect to each other, which immediately implies that no signal is fully revealing about the underlying state. We also assume that \mathbb{F}_0 and \mathbb{F}_1 are not identical, so that some signals are informative. These two assumptions on the signal structure are maintained throughout the paper and will not be stated in the theorems explicitly.

In contrast to much of the literature on social learning, we assume that agents do not necessarily observe all previous actions. Instead, they observe the actions of other agents according to the structure of the social network. To introduce the notion of a social network, let us first define a neighbourhood. Each agent n observes the decisions of the agents in his (stochastically generated) neighbourhood, denoted by $B(n)$.⁷ Since agents can only observe actions taken previously, $B(n) \subseteq \{1, 2, \dots, n-1\}$. Each neighbourhood $B(n)$ is generated according to an arbitrary probability distribution \mathbb{Q}_n over the set of all subsets of $\{1, 2, \dots, n-1\}$. We impose no special assumptions on the sequence of distributions $\{\mathbb{Q}_n\}_{n \in \mathbb{N}}$ except that the draws from each \mathbb{Q}_n are independent from each other for all n and from the realizations of private signals. The sequence $\{\mathbb{Q}_n\}_{n \in \mathbb{N}}$ is the network topology of the social network formed by the agents. The network topology is common knowledge, whereas the realized neighbourhood $B(n)$ and the private signal s_n

6. There is also a literature in engineering, which studies related problems, especially motivated by aggregation of information collected by decentralized sensors. These include Cover (1969), Papastavrou and Athans (1992), Lorenz, Marcinszyn and Steger (2007), and Tay, Tsitsiklis and Win (2008). The work by Papastavrou and Athans contains a result that is equivalent to the characterization of asymptotic learning with the observation of the immediate neighbour.

7. If $n' \in B(n)$, then agent n not only observes the action of n' but also knows the identity of this agent. Crucially, however, n does not observe $B(n')$ or the actions of the agents in $B(n')$.

are the private information of agent n . We say that $\{Q_n\}_{n \in \mathbb{N}}$ is a deterministic network topology if the probability distribution Q_n is a degenerate (Dirac) distribution for all n . Otherwise, *i.e.* if $\{Q_n\}$ for some n is non-degenerate, $\{Q_n\}_{n \in \mathbb{N}}$ is a stochastic network topology.

A social network consists of a network topology $\{Q_n\}_{n \in \mathbb{N}}$ and a signal structure (F_0, F_1) .

Example 1. Here are some examples of network topologies.

1. If $\{Q_n\}_{n \in \mathbb{N}}$ assigns probability 1 to neighbourhood $\{1, 2, \dots, n - 1\}$ for each $n \in \mathbb{N}$, then the network topology is identical to the canonical one studied in the previous literature where each agent observes all previous actions (e.g. Banerjee, 1992; Bikhchandani, Hirshleifer and Welch, 1992; Smith and Sorensen, 2000).
2. If $\{Q_n\}_{n \in \mathbb{N}}$ assigns probability $1/(n - 1)$ to each one of the subsets of size 1 of $\{1, 2, \dots, n - 1\}$ for each $n \in \mathbb{N}$, then we have a network topology of random sampling of one agent from the past.
3. If $\{Q_n\}_{n \in \mathbb{N}}$ assigns probability 1 to neighbourhood $\{n - 1\}$ for each $n \in \mathbb{N}$, then we have a network topology where each individual only observes his immediate neighbour.
4. If $\{Q_n\}_{n \in \mathbb{N}}$ assigns probability 1 to neighbourhoods that are subsets of $\{1, 2, \dots, K\}$ for each $n \in \mathbb{N}$ for some $K \in \mathbb{N}$. In this case, all agents observe the actions of at most K agents.
5. Figure 1 depicts an arbitrary stochastic topology until Agent 7. The thickness of the lines represents the probability with which a particular agent will observe the action of the corresponding preceding agent.

Note that our framework is general enough to nest many social network models studied in the literature, including the popular preferential attachment and small-world network structures. For example, a preferential attachment network can be formed by choosing a stochastic network




FIGURE 1

The figure illustrates the world from the perspective of Agent 7. Agent 7 knows her private signal s_7 , her realized neighbourhood, $B(7) = \{4, 6\}$, and the decisions of Agents 4 and 6, x_4 and x_6 . She also knows the probabilistic model

$$\{Q_n\}_{n < 7}$$

topology $\{\mathbb{Q}_n\}_{n \in \mathbb{N}}$ where all observations are independent of each other and $\mathbb{Q}_n(m \in B(n)) = \alpha_m$ for a sequence of numbers $\{\alpha_m\}_{m \in \mathbb{N}}$. In this network, agents with a high α will be observed by many peers, while agents with low α will not. A small-world network structure can be formed by choosing a partition $\{S_j\}$ of \mathbb{N} such that for every $m < n$, the probability that $m \in B(n)$ is high if both m and n belong to the set S_j for some j and low, but positive, if both m and n belong to different sets S_j and $S_{j'}$. More generally, any network structure can be represented by a judicious choice of $\{\mathbb{Q}_n\}_{n \in \mathbb{N}}$ provided that we keep the assumption that the realizations of $\{\mathbb{Q}_n\}_{n \in \mathbb{N}}$ are independent.⁸ The independence assumption on the neighbourhoods does not impose a restriction on the degree distribution (cardinality) of the agents or on the degree of clustering of the agents. To observe this, note that any given deterministic network topology satisfies the independence assumption and it can be selected to have an arbitrary degree distribution or level of clustering.

3. EQUILIBRIUM STRATEGIES

In this section, we introduce the definitions of equilibrium and asymptotic learning and we provide a characterization of equilibrium strategies. In particular, we show that equilibrium decision rules of individuals can be decomposed into two parts, one that only depends on an individual's private signal and the other that is a function of the observations of past actions. We also show why a full characterization of individual decisions is non-trivial and motivates an alternative proof technique, relying on developing bounds on improvements in the probability of the correct decisions, that will be used in the rest of our analysis.

3.1. Perfect Bayesian equilibrium and asymptotic learning

Given the description above, it is evident that the information set I_n of agent n is given by her signal s_n , her neighbourhood $B(n)$, and all decisions of agents in $B(n)$, *i.e.*

$$I_n = \{s_n, B(n), x_k \text{ for all } k \in B(n)\}. \quad (3.1)$$

The set of all possible information sets of agent n is denoted by \mathcal{I}_n . The agent also knows the network topology $\{\mathbb{Q}_n\}_{n \in \mathbb{N}}$, as it is common knowledge. A strategy for individual n is a mapping $\sigma_n : \mathcal{I}_n \rightarrow \{0, 1\}$ that selects a decision for each possible information set. A strategy profile is a sequence of strategies $\sigma = \{\sigma_n\}_{n \in \mathbb{N}}$. We use the standard notation $\sigma_{-n} = \{\sigma_1, \dots, \sigma_{n-1}, \sigma_{n+1}, \dots\}$ to denote the strategies of all agents other than n and also (σ_n, σ_{-n}) for any n to denote the strategy profile σ . Given a strategy profile σ , the sequence of decisions $\{x_n\}_{n \in \mathbb{N}}$ is a stochastic process and we denote the measure generated by this stochastic process by \mathbb{P}_σ .

Definition 1. *A strategy profile σ is a pure strategy perfect Bayesian equilibrium of this game of social learning if for each $n \in \mathbb{N}$, σ_n maximizes the expected pay-off of agent n given the strategies of other agents σ_{-n} .*

In the rest of the paper, we focus on pure strategy perfect Bayesian equilibria and simply refer to them as “equilibria” (without the pure strategy and the perfect Bayesian qualifiers). We refer to the set of equilibria as Σ .

8. The independence assumption rules out generative preferential attachment models, such as Jackson and Rogers (2007), in which a particular individual being observed more frequently in the past increases the likelihood that he will be observed in the future. Nevertheless, as noted in the text, networks with a preferential attachment structure can be cast as special cases of our model, without abandoning the independence assumption, by fixing *ex ante* which agents are going to be “highly connected”.

Given a strategy profile σ , the expected pay-off of agent n from action $x_n = \sigma_n(I_n)$ is simply $\mathbb{P}_\sigma(x_n = \theta | I_n)$. Therefore, for any equilibrium σ , we have

$$\sigma_n(I_n) \in \operatorname{argmax}_{y \in \{0,1\}} \mathbb{P}_{(y, \sigma_{-n}^*)}(y = \theta | I_n). \tag{3.2}$$

We denote the set of equilibria (pure strategy perfect Bayesian equilibria) of the game by Σ . It is clear that Σ is non-empty. Given the sequence of strategies $\{\sigma_1, \dots, \sigma_{n-1}\}$, the maximization problem in equation (3.2) has a solution for each agent n and each $I_n \in \mathcal{I}_n$. Proceeding inductively, and choosing either one of the actions in case of indifference determines an equilibrium. We note the existence of equilibrium here.

Proposition 1. *There exists a pure strategy perfect Bayesian equilibrium.*

Our main focus is whether equilibrium behaviour will lead to information aggregation. This is captured by the notion of *asymptotic learning*, which is introduced next.

Definition 2. *Given a signal structure $(\mathbb{F}_0, \mathbb{F}_1)$ and a network topology $\{\mathbb{Q}_n\}_{n \in \mathbb{N}}$, we say that asymptotic learning occurs in equilibrium σ if x_n converges to θ in probability (according to measure \mathbb{P}_σ), i.e.*

$$\lim_{n \rightarrow \infty} \mathbb{P}_\sigma(x_n = \theta) = 1.$$

Note that asymptotic learning requires that the probability of taking the correct action converges to 1.⁹ Therefore, asymptotic learning will fail when, as the network becomes large, the limit inferior of the probability of all individuals taking the correct action is strictly less than 1.

Our goal in this paper is to characterize conditions on social networks—on signal structures and network topologies—that ensure asymptotic learning.

3.2. Characterization of individual decisions

Our first result shows that individual decisions can be characterized as a function of the sum of two posteriors. These posteriors play an important role in our analysis. We will refer to these posteriors as the individual’s private belief and the social belief.

Proposition 2. *Let $\sigma \in \Sigma$ be an equilibrium of the game. Let $I_n \in \mathcal{I}_n$ be an information set of agent n . Then, the decision of agent n , $x_n = \sigma_n(I_n)$, satisfies*

$$x_n = \begin{cases} 1, & \text{if } \mathbb{P}_\sigma(\theta = 1 | s_n) + \mathbb{P}_\sigma(\theta = 1 | B(n), x_k, k \in B(n)) > 1, \\ 0, & \text{if } \mathbb{P}_\sigma(\theta = 1 | s_n) + \mathbb{P}_\sigma(\theta = 1 | B(n), x_k, k \in B(n)) < 1, \end{cases}$$

and $x_n \in \{0, 1\}$ otherwise.

Proof. See Appendix A. ||

This proposition establishes an additive decomposition in the equilibrium decision rule between the information obtained from the private signal of the individual and from the observations of others’ actions (in his neighbourhood). The next definition formally distinguishes between the two components of an individual’s information.

9. It is also clear that asymptotic learning is equivalent to the posterior beliefs converging to a distribution putting probability 1 on the true state.

Definition 3. We refer to the probability $\mathbb{P}_\sigma(\theta = 1 \mid s_n)$, as the private belief of agent n , and the probability

$$\mathbb{P}_\sigma(\theta = 1 \mid B(n), x_k \text{ for all } k \in B(n)),$$

as the social belief of agent n .

Proposition 2 and Definition 3 imply that the equilibrium decision rule for agent $n \in \mathbb{N}$ is equivalent to choosing $x_n = 1$ when the sum of his private and social beliefs is greater than 1. Consequently, the properties of private and social beliefs will shape equilibrium learning behaviour.

Note that the social belief depends on n since it is a function of the (realized) neighbourhood of agent n . In most learning models, social beliefs have natural monotonicity properties. For example, a greater fraction of individuals choosing action $x = 1$ in the information set of agent n would increase the social belief of agent n . It is straightforward to construct examples, where such monotonicity properties do not hold under general social networks (see Appendix B in Supplementary Material). For this reason, we will use a different line of attack, based on developing lower bounds on the probability of taking the correct action, for establishing our main results, which are presented in the next section.

3.3. Bounded and unbounded private beliefs

The private belief of an individual is a function of his private signal $s \in S$ and is not a function of the strategy profile σ since it does not depend on the decisions of other agents. We represent probabilities that do not depend on the strategy profile by \mathbb{P} . We use the notation p_n to represent the private belief of agent n , *i.e.*

$$p_n = \mathbb{P}(\theta = 1 \mid s_n).$$

A straightforward application of Bayes' rule implies that for any n and any signal $s_n \in \bar{S}$, the private belief p_n of agent n is given by¹⁰

$$p_n = \left(1 + \frac{d\mathbb{F}_0}{d\mathbb{F}_1}(s_n)\right)^{-1}. \quad (3.3)$$

We next define the support of a private belief. In our subsequent analysis, we will see that properties of the support of private beliefs play a key role in asymptotic learning behaviour. Since the p_n are identically distributed for all n (which follows by the assumption that the private signals s_n are identically distributed), in the following, we will use Agent 1's private belief p_1 to define the support and the conditional distributions of private beliefs.

Definition 4. The support of the private beliefs is the interval $[\underline{\beta}, \bar{\beta}]$, where the end points of the interval are given by

$$\underline{\beta} = \inf\{r \in [0, 1] \mid \mathbb{P}(p_1 \leq r) > 0\} \quad \text{and} \quad \bar{\beta} = \sup\{r \in [0, 1] \mid \mathbb{P}(p_1 \leq r) < 1\}.$$

The signal structure has bounded private beliefs if $\underline{\beta} > 0$ and $\bar{\beta} < 1$ and unbounded private beliefs if $\underline{\beta} = 1 - \bar{\beta} = 1$.

10. If the probability measures \mathbb{F}_0 and \mathbb{F}_1 have densities f_0 and f_1 , respectively, then $\frac{d\mathbb{F}_0}{d\mathbb{F}_1}(s_n) = \frac{f_0(s_n)}{f_1(s_n)}$.

When private beliefs are bounded, there is a maximum informativeness to any signal. When they are unbounded, agents may receive arbitrarily strong signals favouring either state (this follows from the assumption that $(\mathbb{F}_0, \mathbb{F}_1)$ are absolutely continuous with respect to each other).

The conditional distribution of private beliefs given the underlying state $j \in \{0, 1\}$ can be directly computed as

$$\mathbb{G}_j(r) = \mathbb{P}(p_1 \leq r \mid \theta = j). \tag{3.4}$$

The signal structure $(\mathbb{F}_0, \mathbb{F}_1)$ can be equivalently represented by the corresponding private belief distributions $(\mathbb{G}_0, \mathbb{G}_1)$, and in what follows, it will typically be more convenient to work with $(\mathbb{G}_0, \mathbb{G}_1)$ rather than $(\mathbb{F}_0, \mathbb{F}_1)$. It is straightforward to verify that $\mathbb{G}_0(r)/\mathbb{G}_1(r)$ is non-increasing in r and $\mathbb{G}_0(r)/\mathbb{G}_1(r) > 1$ for all $r \in (\underline{\beta}, \bar{\beta})$ (see Lemma A1 in Appendix A).

4. MAIN RESULTS

In this section, we present our main results on asymptotic learning and provide the main intuition for the proofs.

4.1. Expanding observations

We start by introducing the key properties of network topologies and signal structures that impact asymptotic learning. Intuitively, for asymptotic learning to occur, the information that each agent receives from other agents should not be confined to a bounded subset of agents. This property is established in the following definition. For this definition and throughout the paper, if the set $B(n)$ is empty, we set $\max_{b \in B(n)} b = 0$.

Definition 5. *The network topology has expanding observations if for all $K \in \mathbb{N}$, we have*

$$\lim_{n \rightarrow \infty} \mathbb{Q}_n \left(\max_{b \in B(n)} b < K \right) = 0.$$

If the network topology does not satisfy this property, then we say it has non-expanding observations.

Recall that the neighbourhood of agent n is a random variable $B(n)$ (with values in the set of subsets of $\{1, 2, \dots, n - 1\}$) and distributed according to \mathbb{Q}_n . Therefore, $\max_{b \in B(n)} b$ is a random variable that takes values in $\{0, 1, \dots, n - 1\}$. The expanding observations condition can be restated as the sequence of random variables $\{\max_{b \in B(n)} b\}_{n \in \mathbb{N}}$ converging to infinity in probability. Similarly, it follows from the preceding definition that the network topology has non-expanding observations if and only if there exists some $K \in \mathbb{N}$ and some scalar $\varepsilon > 0$ such that

$$\limsup_{n \rightarrow \infty} \mathbb{Q}_n \left(\max_{b \in B(n)} b < K \right) \geq \varepsilon.$$

An alternative restatement of this definition might clarify its meaning. Let us refer to a finite set of individuals C as *excessively influential* if there exists a subsequence of agents who, with probability uniformly bounded away from 0, observe the actions of a subset of C . Then, the network topology has non-expanding observations if and only if there exists an excessively influential group of agents. Note also that if there is a minimum amount of arrival of new information in the network, so that the probability of an individual observing some other individual from the recent past goes to one as the network becomes large, then the network topology will feature expanding observations. This discussion therefore highlights that the requirement that a network

topology has expanding observations is quite mild and most social networks, including all of those discussed above, satisfy this requirement.

When the topology has non-expanding observations, there is a subsequence of agents that draws information from the first K decisions with positive probability (uniformly bounded away from 0). It is then intuitive that network topologies with non-expanding observations will preclude asymptotic learning. Our first theorem states and proves this result.

Theorem 1. *Assume that the network topology $\{\mathbb{Q}_n\}_{n \in \mathbb{N}}$ has non-expanding observations. Then, there exists no equilibrium $\sigma \in \Sigma$ with asymptotic learning.*

Proof. See Appendix A. \parallel

This theorem states the intuitive result that with non-expanding observations, asymptotic learning will fail. This result is not surprising since asymptotic learning requires the aggregation of the information of different individuals. But a network topology with non-expanding observations does not allow such aggregation. Intuitively, non-expanding observations, or equivalently the existence of an excessively influential group of agents, imply that infinitely many individuals will observe finitely many actions with positive probability and this will not enable them to aggregate the dispersed information collectively held by the entire social network.

4.2. Asymptotic learning with unbounded private beliefs

A central question is then whether, once we exclude network topologies with non-expanding observations, what other conditions need to be imposed to ensure asymptotic learning. The following theorem is one of the main results of the paper and shows that for general network topologies, unbounded private beliefs play a key role. In particular, unbounded private beliefs and expanding observations are sufficient for asymptotic learning in all equilibria.

Theorem 2. *Assume that the signal structure $(\mathbb{F}_0, \mathbb{F}_1)$ has unbounded private beliefs and the network topology $\{\mathbb{Q}_n\}_{n \in \mathbb{N}}$ has expanding observations. Then, asymptotic learning occurs in every equilibrium $\sigma \in \Sigma$.*

Proof. See Appendix A. \parallel

Theorem 2 implies that unbounded private beliefs are sufficient for asymptotic learning for most (but not all) network topologies. In particular, the condition that the network topology has expanding observations is fairly mild and only requires a minimum amount of arrival of recent information to the network. Social networks in which each individual observes all past actions, those in which each observes just his neighbour and those in which each individual observes $M \geq 1$ agents independently and uniformly drawn from his predecessors are all examples of network topologies with expanding observations. Theorem 2 therefore implies that unbounded private beliefs are sufficient to guarantee asymptotic learning in social networks with these properties and many others.

This theorem also guarantees learning in the presence of agents who are highly influential, in the sense that their actions are visible to the entire society, but are *not* excessively influential, as they are not the only sources of information in the networks. Consider, *e.g.* a network where the actions of the first K agents are visible to everyone, but each agent also observes her immediate neighbour, *i.e.* $B(n) = \{1, 2, \dots, K, n-1\}$. This network topology satisfies expanding observations and thus leads to learning provided that private beliefs are unbounded. This is in contrast

with the predictions of non-Bayesian learning models such as DeMarzo, Vayanos and Zwiebel (2003) and Golub and Jackson (2010). To see this, suppose, as in these models, that agents determine their new beliefs as a weighted average of their own beliefs and the beliefs of the agents they observe and that the first K agents are influential in the sense that they are observed by all future agents and receive a weight of at least $\varepsilon > 0$. It is then straightforward that there will not be asymptotic learning, which contrasts with Theorem 2.

The proof of Theorem 2 is presented in Appendix A. Here, we provide a road map and the general intuition. As noted in the previous section, there is no monotonicity result linking the behaviour of an agent to the fraction of actions he or she observes. Instead, we prove Theorem 2 by making use of an informational monotonicity related to the (expected) welfare improvement principle in Banerjee and Fudenberg (2004) and in Smith and Sorensen (2008) and the imitation principle in Bala and Goyal (1998) and Gale and Kariv (2003). In particular, we first consider a special case in which each individual only observes one other from the past (*i.e.* $B(n)$ is a singleton for each n). We then establish the following *strong improvement principle*: with unbounded private beliefs, there exists a strict lower bound on the increase in the *ex ante* probability that an individual will make a correct decision over his neighbour's probability (recall that for now there is a single agent in each individual's neighbourhood, thus each individual has a single "neighbour"). Intuitively, each individual can copy the behaviour of their neighbour unless they have a very strong signal that points in a different direction. We show that this results in a strict improvement in the probability of taking the right action whenever this probability is not equal to 1.

We then prove a *generalized strong improvement principle* for arbitrary social networks by showing that each individual can obtain such an improvement even if they ignore all but one of the individuals in their information set. The overall improvement in the probability that each individual will take the right action is a priori greater than this lower bound. The proof of Theorem 2 then follows by using this generalized strong improvement principle to construct a subsequence of informational improvements at each point and showing that there is a strict improvement at each step of the subsequence. This combined with the expanding observations assumption on the network topology establishes the asymptotic learning result in the proof of the theorem.

The following corollary to Theorems 1 and 2 shows that for an interesting class of stochastic network topologies, there is a critical topology at which there is a phase transition—*i.e.* for all network topologies with greater expansion of observations than this critical topology, there will be asymptotic learning and for all topologies with less expansion, asymptotic learning will fail.

Corollary 1. *Assume that the signal structure $(\mathbb{F}_0, \mathbb{F}_1)$ has unbounded private beliefs. Assume also that the network topology is given by $\{\mathbb{Q}_n\}_{n \in \mathbb{N}}$ such that*

$$\mathbb{Q}_n(m \in B(n)) = \frac{a}{(n-1)^c} \quad \text{for all } n \text{ and all } m < n,$$

where, given n , the draws for $m, m' < n$, are independent and a and c are positive constants. If $c < 1$, then asymptotic learning occurs in all equilibria. If $c \geq 1$, then asymptotic learning does not occur in any equilibrium.

Proof. See Appendix A. ||

Given the class of network topologies in this corollary, $c < 1$ implies that as the network becomes large, there will be sufficient expansion of observations. In contrast, for $c \geq 1$, stochastic process \mathbb{Q}_n does not place enough probability on observing recent actions and the network

topology is non-expanding. Consequently, Theorem 1 applies and there is no asymptotic learning.

To highlight the implications of Theorems 1 and 2 for deterministic network topologies, let us introduce the following definition.

Definition 6. *Assume that the network topology is deterministic. Then, we say a finite sequence of agents π is an information path of agent n if for each i , $\pi_i \in B(\pi_{i+1})$, and the last element of π is n . Let $\bar{\pi}(n)$ be an information path of agent n that has maximal length. Then, we let $L(n)$ denote the number of elements in $\bar{\pi}(n)$ and call it agent n 's information depth.*

Intuitively, the concepts of information path and information depth capture the intuitive notion of how long the “trail” of the information in the neighbourhood of an individual is. For example, if each individual observes only his immediate neighbour (*i.e.* $B(n) = \{n-1\}$ with probability 1), each will have a small neighbourhood, but the information depth of a high-indexed individual will be high (or the “trail” will be long) because the immediate neighbour’s action will contain information about the signals of all previous individuals. The next corollary shows that with deterministic network topologies, asymptotic learning will occur if and only if the information depth (or the trail of the information) increases without bound as the network becomes larger.

Corollary 2. *Assume that the signal structure $(\mathbb{F}_0, \mathbb{F}_1)$ has unbounded private beliefs. Assume that the network topology is deterministic. Then, asymptotic learning occurs for all equilibria if the sequence of information depths $\{L(n)\}_{n \in \mathbb{N}}$ goes to infinity. If the sequence $\{L(n)\}_{n \in \mathbb{N}}$ does not go to infinity, then asymptotic learning does not occur in any equilibrium.*

Proof. See Appendix A. \parallel

4.3. No learning under bounded private beliefs

In the full observation network topology, bounded beliefs imply lack of asymptotic learning. One might thus expect a converse to Theorem 2, whereby asymptotic learning fails whenever signals are bounded. Under general network topologies, learning dynamics turn out to be more interesting and richer. The next theorem provides a partial converse to Theorem 2 and shows that for a wide range of deterministic and stochastic network topologies, bounded beliefs imply no asymptotic learning. However, somewhat surprisingly, Theorem 4 will show that the same is not true with more general stochastic network topologies.

Theorem 3. *Assume that the signal structure $(\mathbb{F}_0, \mathbb{F}_1)$ has bounded private beliefs. If the network topology $\{Q_n\}_{n \in \mathbb{N}}$ satisfies one of the following conditions,*

- (a) $B(n) = \{1, \dots, n-1\}$ for all n ,
- (b) $|B(n)| \leq 1$ for all n , or
- (c) there exists some constant M such that $|B(n)| \leq M$ for all n and

$$\lim_{n \rightarrow \infty} \max_{b \in B(n)} b = \infty \quad \text{with probability 1,}$$

then, asymptotic learning does not occur in any equilibrium $\sigma \in \Sigma$.

Proof. See Appendix B in Supplementary Material. \parallel

This theorem implies that in many common deterministic and stochastic network topologies, bounded private beliefs imply lack of asymptotic learning. Part (a) of this theorem is already proved by Smith and Sorensen (2000), we provide an alternative proof in Appendix B in Supplementary Material. Intuitively, in all cases, we can provide an upper bound on the amount of learning. We establish this upper bound by showing that either individuals start following potentially incorrect social beliefs after a certain stage of learning has been reached or they rely on their own private signal, taking the incorrect action with positive probability.

Part (c) of the theorem is the more substantive contribution. The idea behind the proof of Part (c) is as follows. We suppose, to obtain a contradiction, that asymptotic learning occurs. We then show that this implies that the social belief of agents who only observe neighbours selecting Action 1 converges to 1. The probability that the first K agents select Action 1 is positive (for any K). In case of such an event and K large, eventually all agents will have a social belief close to 1, which is proved using a union bound and the fact that each agent observes at most M actions. But since, with bounded private beliefs, an agent with social belief close to 1 will ignore her signal, eventually decisions will contain no new information, contradicting asymptotic learning.

The following corollary illustrates the implications of Theorem 3. It shows that, when private beliefs are bounded, there will be no asymptotic learning (in any equilibrium) in stochastic networks with random sampling.

Corollary 3. *Assume that the signal structure $(\mathbb{F}_0, \mathbb{F}_1)$ has bounded private beliefs. Assume that each agent n samples M agents uniformly and independently among $\{1, \dots, n-1\}$ for some $M \geq 1$. Then, asymptotic learning does not occur in any equilibrium $\sigma \in \Sigma$.*

Proof. See Appendix B in Supplementary Material. \parallel

4.4. Asymptotic learning with bounded private beliefs

While in the full observation network topology studied by Smith and Sorensen (2000) and in other instances investigated in the previous literature, bounded private beliefs preclude asymptotic learning, this is no longer true under general stochastic social networks. Characterizing the set of stochastic network topologies under which learning occurs for all distributions of private beliefs is the next major contribution of our paper. To present these results, we first introduce the notion of a *non-persuasive neighbourhood*.

Definition 7. *A finite set $B \subset \mathbb{N}$ is a non-persuasive neighbourhood in equilibrium $\sigma \in \Sigma$ if*

$$\mathbb{P}_\sigma(\theta = 1 \mid x_k = y_k \text{ for all } k \in B) \in (1 - \bar{\beta}, 1 - \underline{\beta})$$

for any set of values $y_k \in \{0, 1\}$ for each k . We denote the set of all non-persuasive neighbourhoods by \mathcal{U}_σ .

A neighbourhood B is non-persuasive in equilibrium $\sigma \in \Sigma$ if for any set of decisions that agent n observes, his behaviour may still depend on his private signal. That is, for any set of decisions the agent observes, there exists some private signal such that the agent chooses Action 0 and some private signal such that the agent chooses Action 1.

A non-persuasive neighbourhood is defined with respect to a particular equilibrium. However, we will provide below several different sufficient conditions for a collection of neighbourhoods to be non-persuasive in any equilibrium.

Our main theorem for learning with bounded beliefs provides a broad class of stochastic social networks where asymptotic learning takes place for any signal structure. It relies on the

society comprising two subsets, one learning from agents with non-persuasive neighbourhoods and the other learning from the former subset. This necessitates strengthening the expanding observations condition to ensure that the latter subset receives new information from the former subset.

Definition 8. *Given a subset $S \subseteq \mathbb{N}$, the network topology $\{\mathbb{Q}_n\}_{n \in \mathbb{N}}$ has expanding observations with respect to S if for all K ,*

$$\lim_{n \rightarrow \infty} \mathbb{Q}_n \left(\max_{b \in B(n) \cap S} b < K \right) = 0.$$

Theorem 4. *Let $(\mathbb{F}_0, \mathbb{F}_1)$ be an arbitrary signal structure and let $S \subseteq \mathbb{N}$. Let σ be the equilibrium where agents break ties in favour of Action 0.¹¹ Assume the network topology $\{\mathbb{Q}_n\}_{n \in \mathbb{N}}$ has expanding observations with respect to S and has a lower bound on the probability of observing the entire history of actions along S , i.e. there exists some $\underline{\varepsilon} > 0$ such that*

$$\mathbb{Q}_n(B(n) = \{1, \dots, n-1\}) \geq \underline{\varepsilon} \quad \text{for all } n \in S.$$

Assume further that for some positive integer M and non-persuasive neighbourhoods C_1, \dots, C_M , i.e. $C_i \in \mathcal{U}_\sigma$ for all $i = 1, \dots, M$, we have

$$\sum_{n \in S} \sum_{i=1}^M \mathbb{Q}_n(B(n) = C_i) = \infty.$$

Then, asymptotic learning occurs in equilibrium σ .

Proof. See Appendix A. \parallel

This is a rather surprising result, particularly in view of existing results in the literature, which generate herds and information cascades (and no learning) with bounded beliefs. This theorem indicates that learning dynamics become significantly richer when we consider general social networks.

Theorem 4 highlights a class of network structures that lead to learning. In these structures, there is an infinite subset S that forms the core of the network. Within this core S , there are two important subgroups of agents. There are infinitely many agents who act based (partially) on their signals since they have non-persuasive neighbourhoods. The existence of these agents does not preclude learning since the probability that an agent n has a non-persuasive neighbourhood can go to 0 as n goes to infinity. By acting based on their private signals, these agents play the essential function of bringing new information into the society. The second key subgroup within the core of agents in S is the set of all agents who observe the entire history of play. These agents play the role of collecting the information brought in by the agents who have non-persuasive neighbourhoods. From the martingale convergence theorem, the social belief of these agents converges with probability 1, and because there is sufficient information generated by the agents with non-persuasive neighbourhoods, the social belief must converge to the correct state of the world. The expanding observations with respect to S guarantees that information spreads to all other agents. In particular, agents who do not observe the entire history of actions

11. This assumption is not necessary to prove this result but assuming a particular equilibrium significantly simplifies the notation needed in the proof.

can simply copy the action of their highest neighbours within S . This strategy provides a lower bound on the pay-off to these agents and in fact guarantees asymptotic learning. Therefore, these agents also asymptotically learn the state in equilibrium.¹²

We conclude this subsection by providing sufficient conditions for a neighbourhood to be non-persuasive in any equilibrium. It is straightforward to see that $B = \emptyset$, i.e. the empty neighbourhood, is non-persuasive in any equilibrium. The following two propositions present two other classes of non-persuasive neighbourhoods.

Proposition 3. *Let $(\mathbb{G}_0, \mathbb{G}_1)$ be private belief distributions. Assume that the first K agents have empty neighbourhoods, i.e. $B(n) = \emptyset$ for all $n \leq K$, and K satisfies*

$$K < \min \left\{ \frac{\log\left(\frac{\bar{\beta}}{1-\bar{\beta}}\right)}{\log\left(\frac{\mathbb{G}_0(1/2)}{\mathbb{G}_1(1/2)}\right)}, \frac{\log\left(\frac{\underline{\beta}}{1-\underline{\beta}}\right)}{\log\left(\frac{1-\mathbb{G}_0(1/2)}{1-\mathbb{G}_1(1/2)}\right)} \right\}. \tag{4.5}$$

Then, any subset $B \subseteq \{1, 2, \dots, K\}$ is a non-persuasive neighbourhood.

Proof. See Appendix A. \parallel

To obtain the intuition for equation (4.5), consider the equivalent pair of inequalities

$$\left(1 + \frac{\mathbb{G}_0(1/2)^K}{\mathbb{G}_1(1/2)^K}\right)^{-1} > 1 - \bar{\beta} \quad \text{and} \quad \left(1 + \frac{(1 - \mathbb{G}_0(1/2))^K}{(1 - \mathbb{G}_1(1/2))^K}\right)^{-1} < 1 - \underline{\beta}. \tag{4.6}$$

The factor of $(1 + (\mathbb{G}_0(1/2)/\mathbb{G}_1(1/2))^K)^{-1}$ represents the probability that the state θ is equal to 1 conditional on K agents independently selecting Action 0. For such a neighbourhood of K agents acting independently to be non-persuasive, there must exist a signal strong enough in favour of State 1 such that an agent would select Action 1 after observing K independent agents choosing 0. From the equilibrium decision rule (cf. Proposition 2), it follows that this holds if $(1 + (\mathbb{G}_0(1/2)/\mathbb{G}_1(1/2))^K)^{-1} + \bar{\beta} > 1$. We can repeat the same argument for Action 1 to obtain the second inequality in equation (4.6).

The condition in equation (4.5) has a natural interpretation: for a neighbourhood B to be non-persuasive, any $|B|$ decisions have to be less informative than a single very informative signal. In the case where an agent $k \in B$ has an empty neighbourhood, the informativeness of her decision $x_k = 0$ is given by $\mathbb{G}_0(1/2)/\mathbb{G}_1(1/2)$ and can be interpreted as the informativeness of an ‘‘average’’ signal in favour of State 0 since any signal mildly in favour of State 0 would lead to Action 0. Therefore, the ratio between the informative value of an average signal and the informative value of an extreme signal determines the size $|B|$ of the largest non-persuasive neighbourhood.

12. It is important to emphasize the difference between this result and that in SgROI (2002), which shows that a social planner can ensure some degree of information aggregation by forcing a subsequence of agents to make decisions without observing past actions. With the same reasoning, one might conjecture that asymptotic learning may occur if a particular subsequence of agents, such as that indexed by prime numbers, has empty neighbourhoods. However, there will not be asymptotic learning in this deterministic topology since $\liminf_{n \rightarrow \infty} \mathbb{P}_\sigma(x_n = \theta) < 1$. For the result that there is asymptotic learning (i.e. $\liminf_{n \rightarrow \infty} \mathbb{P}_\sigma(x_n = \theta) = 1$) in Theorem 4, the feature that the network topology is stochastic is essential.

Proposition 4. Let $(\mathbb{G}_0, \mathbb{G}_1)$ be private belief distributions. Assume that the first K agents observe the full history of past actions, i.e. $B(n) = \{1, \dots, n-1\}$ for all $n \leq K$, and K satisfies

$$K < \min \left\{ \frac{\log\left(\frac{\bar{\beta}}{1-\bar{\beta}}\right)}{\log\left(\frac{\mathbb{G}_0(1/2)}{\mathbb{G}_1(1/2)}\right)}, \frac{\log\left(\frac{\beta}{1-\beta}\right)}{\log\left(\frac{1-\mathbb{G}_0(1/2)}{1-\mathbb{G}_1(1/2)}\right)} \right\}. \quad (4.7)$$

Assume also that the private belief distributions $(\mathbb{G}_0, \mathbb{G}_1)$ satisfy the following monotonicity conditions: the functions

$$\left(\frac{1}{q} - 1\right) \frac{\mathbb{G}_0(1-q)}{\mathbb{G}_1(1-q)} \quad \text{and} \quad \left(\frac{1}{q} - 1\right) \frac{1 - \mathbb{G}_0(1-q)}{1 - \mathbb{G}_1(1-q)} \quad (4.8)$$

are both non-increasing in q . Then, any subset $B \subseteq \{1, 2, \dots, K\}$ is a non-persuasive neighbourhood.

Proof. See Appendix A. \parallel

The monotonicity condition of equation (4.8) guarantees the following: let $0 \leq q' \leq q'' \leq 1$; then $\mathbb{P}_\sigma(\theta = 1 \mid q_n = q', x_n = 0) \leq \mathbb{P}_\sigma(\theta = 1 \mid q_n = q'', x_n = 0)$ and $\mathbb{P}_\sigma(\theta = 1 \mid q_n = q', x_n = 1) \leq \mathbb{P}_\sigma(\theta = 1 \mid q_n = q'', x_n = 1)$. That is, an agent n with a given social belief $q_n = q'$ selecting Action 0 (or Action 1) represents a stronger signal in favour of State $\theta = 0$ than an agent with social belief $q_n = q'' \geq q'$ choosing Action 0 (or, respectively, Action 1).

When the monotonicity condition of equation (4.8) holds, we obtain a partial ordering on the informativeness of sets of decisions. In particular, it determines that the most informative sequence of actions (x_1, x_2, \dots, x_K) of the first K agents in favour of State 0 is the sequence $(0, 0, \dots, 0)$. Without the monotonicity condition, this seemingly mild claim is not generally true. We can bound the informativeness of the sequence of decisions $(x_1, \dots, x_K) = (0, \dots, 0)$ using the fact (cf. Lemma A1(c) from Appendix A) that $\mathbb{G}_0(q_1)/\mathbb{G}_1(q_1) \geq \mathbb{G}_0(q_2)/\mathbb{G}_1(q_2) \geq \dots \geq \mathbb{G}_0(q_K)/\mathbb{G}_1(q_K)$, where q_k is the social belief of agent K . In this situation, all decisions are less informative than Agent 1's decision yielding

$$\frac{\mathbb{P}_\sigma(x_k = 0, k \leq K \mid \theta = 0)}{\mathbb{P}_\sigma(x_k = 0, k \leq K \mid \theta = 1)} \leq \left(\frac{\mathbb{G}_0(q_1)}{\mathbb{G}_1(q_1)}\right)^K = \left(\frac{\mathbb{G}_0(1/2)}{\mathbb{G}_1(1/2)}\right)^K.$$

Using the same argument as in Proposition 3, we obtain that the condition in equation (4.7) is sufficient to guarantee the neighbourhood B is unpersuasive.

5. CONCLUSION

In this paper, we studied the problem of sequential learning over a general social network. A large social learning literature, pioneered by Bikhchandani, Hirshleifer and Welch (1992), Banerjee (1992), and Smith and Sorensen (2000), has studied equilibria of sequential-move games, where each individual observes all past actions. The focus has been on whether equilibria lead to aggregation of information (and thus to asymptotic learning).

In many relevant situations, individuals obtain their information not by observing all past actions, but from their "social network". This raises the question of how the structures of social networks in which individuals are situated affects the equilibrium outcomes. To address these questions, we formulated a sequential learning model over a general social network.

In our model, each individual receives a signal about the underlying state of the world and observes the past actions of a stochastically generated neighbourhood of individuals. The stochastic

process generating the neighbourhoods defines the network topology. The signal structure determines the conditional distributions of the signals received by each individual as a function of the underlying state. The social network consists of the network topology and the signal structure. Each individual then chooses one of two possible actions depending on his posterior beliefs given his signal and the realized neighbourhood. We characterized pure-strategy (perfect Bayesian) equilibria for arbitrary stochastic and deterministic social networks and characterized the conditions under which there is asymptotic learning. Asymptotic learning corresponds to individual decisions converging (in probability) to the right action as the social network becomes large.

Two concepts turn out to be crucial in determining whether there will be asymptotic learning. The first is common with the previous literature: following [Smith and Sorensen \(2000\)](#), we say that private beliefs are bounded if the likelihood ratio implied by individual signals is bounded and there is a maximum amount of information that can be generated from these signals. Conversely, private beliefs are unbounded if the corresponding likelihood ratio is unbounded. The second important concept is that of expanding or non-expanding observations. A network topology has non-expanding observations if there exists infinitely many agents observing the actions of only a finite subset of (excessively influential) agents. Most network topologies feature expanding observations.

Non-expanding observations do not allow asymptotic learning since there exists infinitely many agents who do not receive sufficiently many observations to be able to aggregate information.

Our first main theorem, [Theorem 2](#), shows that expanding observations and unbounded private signals are sufficient to ensure asymptotic learning. Since expanding observations is a relatively mild restriction, to the extent that unbounded private beliefs constitute a good approximation to the informativeness of individual signals, this result implies that all equilibria featuring asymptotic learning applies in a wide variety of settings. Another implication is that asymptotic learning is possible even when there are “influential agents” or “information leaders”, *i.e.* individuals who are observed by many, most, or even all agents (while others may be observed not at all or much less frequently). It is only when individuals are excessively influential—loosely speaking when they act as the sole source of information for infinitely many agents—that asymptotic learning ceases to apply.

We also provided a partial converse to this result, showing that under the most common deterministic or stochastic network topologies, bounded beliefs imply no asymptotic learning. However, our second main theorem, [Theorem 4](#), characterizes the class of network structures where asymptotic learning occurs with bounded beliefs. This result shows the importance of general (stochastic) network topologies in the study of sequential learning since asymptotic learning with bounded beliefs is impossible in the deterministic network topologies studied in the previous literature.

It is also useful to acknowledge the limitations of our model. In particular, our results rely on the fact that each agent acts only once and the order of actions is exogenous. Therefore, the predictions of the model are most likely to apply in situations where agents act infrequently and have limited control over the timing of their actions. The analysis of Bayesian learning when a network of agents interacts repeatedly and learn from communication as well as observation is an interesting area for future research.

We believe that the framework developed here opens the way for a more general analysis of the impact of structure of social networks on learning dynamics in sequential settings. Among the questions that can be studied in future work using this framework are the following: (1) the effect of network structure on the speed (rate of convergence) of sequential learning, (2) equilibrium learning when there are heterogeneous preferences, (3) equilibrium learning when

the underlying state is changing dynamically, and (4) the influence of a subset of a social network (e.g. the media or interested parties) in influencing the views of the rest as a function of the network structure.

APPENDIX A: MAIN PROOFS

Proof of Proposition 2.

We prove that if

$$\mathbb{P}_\sigma(\theta = 1 | s_n) + \mathbb{P}_\sigma(\theta = 1 | B(n), x_k \text{ for all } k \in B(n)) > 1, \quad (\text{A1})$$

then $x_n = 1$. The proofs of the remaining statements follow the same line of argument. We first show that equation (A1) holds if and only if

$$\mathbb{P}_\sigma(\theta = 1 | I_n) > 1/2, \quad (\text{A2})$$

therefore implying that $x_n = 1$ by the equilibrium condition [cf. equation (3.2)]. By Bayes' rule, equation (A2) is equivalent to

$$\mathbb{P}_\sigma(\theta = 1 | I_n) = \frac{d\mathbb{P}_\sigma(I_n | \theta = 1)\mathbb{P}_\sigma(\theta = 1)}{\sum_{j=0}^1 d\mathbb{P}_\sigma(I_n | \theta = j)\mathbb{P}_\sigma(\theta = j)} = \frac{d\mathbb{P}_\sigma(I_n | \theta = 1)}{\sum_{j=0}^1 d\mathbb{P}_\sigma(I_n | \theta = j)} > 1/2,$$

where the second equality follows from the assumption that States 0 and 1 are equally likely. Hence, equation (A2) holds if and only if

$$d\mathbb{P}_\sigma(I_n | \theta = 1) > d\mathbb{P}_\sigma(I_n | \theta = 0).$$

Conditional on state θ , the private signals and the observed decisions are independent, i.e.

$$d\mathbb{P}_\sigma(I_n | \theta = j) = d\mathbb{P}_\sigma(s_n | \theta = j)\mathbb{P}_\sigma(B(n), x_k, k \in B(n) | \theta = j).$$

Combining the preceding two relations, it follows that equation (A2) is equivalent to

$$\frac{\mathbb{P}_\sigma(B(n), x_k, k \in B(n) | \theta = 1)}{\sum_{j=0}^1 \mathbb{P}_\sigma(B(n), x_k, k \in B(n) | \theta = j)} > \frac{d\mathbb{P}_\sigma(s_n | \theta = 0)}{\sum_{j=0}^1 d\mathbb{P}_\sigma(s_n | \theta = j)}.$$

Since both states are equally likely, this can be rewritten as follows:

$$\frac{\mathbb{P}_\sigma(B(n), x_k, k \in B(n) | \theta = 1)P_\sigma(\theta = 1)}{\sum_{j=0}^1 \mathbb{P}_\sigma(B(n), x_k, k \in B(n) | \theta = j)P_\sigma(\theta = j)} > \frac{d\mathbb{P}_\sigma(s_n | \theta = 0)P_\sigma(\theta = 0)}{\sum_{j=0}^1 d\mathbb{P}_\sigma(s_n | \theta = j)P_\sigma(\theta = j)}.$$

Applying Bayes' rule on both sides of the inequality above, we see that the preceding relation is identical to

$$\mathbb{P}_\sigma(\theta = 1 | B(n), x_k, k \in B(n)) > \mathbb{P}_\sigma(\theta = 0 | s_n) = 1 - \mathbb{P}_\sigma(\theta = 1 | s_n),$$

completing the proof. \parallel

Proof of Theorem 1.

Suppose that the network has non-expanding observations. This implies that there exists some $K \in \mathbb{N}$, $\varepsilon > 0$, and a subsequence of agents \mathcal{N} such that for all $n \in \mathcal{N}$,

$$\mathbb{Q}_n \left(\max_{b \in B(n)} b < K \right) \geq \varepsilon. \quad (\text{A3})$$

For any such agent $n \in \mathcal{N}$, we have

$$\begin{aligned} \mathbb{P}_\sigma(x_n = \theta) &= \mathbb{P}_\sigma \left(x_n = \theta \mid \max_{b \in B(n)} b < K \right) \mathbb{Q}_n \left(\max_{b \in B(n)} b < K \right) \\ &\quad + \mathbb{P}_\sigma \left(x_n = \theta \mid \max_{b \in B(n)} b \geq K \right) \mathbb{Q}_n \left(\max_{b \in B(n)} b \geq K \right) \\ &\leq \mathbb{P}_\sigma \left(x_n = \theta \mid \max_{b \in B(n)} b < K \right) \mathbb{Q}_n \left(\max_{b \in B(n)} b < K \right) + \mathbb{Q}_n \left(\max_{b \in B(n)} b \geq K \right) \\ &= 1 + \mathbb{Q}_n \left(\max_{b \in B(n)} b < K \right) \left(-1 + \mathbb{P}_\sigma \left(x_n = \theta \mid \max_{b \in B(n)} b < K \right) \right) \\ &\leq 1 - \varepsilon + \varepsilon \mathbb{P}_\sigma \left(x_n = \theta \mid \max_{b \in B(n)} b < K \right), \end{aligned} \quad (\text{A4})$$

where the second inequality follows from equation (A3).

Given some equilibrium $\sigma \in \Sigma$ and agent n , we define z_n as the decision that maximizes the conditional probability of making a correct decision given the private signals and neighbourhoods of the first $K - 1$ agents and agent n , *i.e.*

$$z_n = \operatorname{argmax}_{y \in \{0,1\}} \mathbb{P}_{y, \sigma - n}(y = \theta \mid s_i, B(i), \text{ for } i = 1, \dots, K - 1, n). \tag{A5}$$

We denote a particular realization of private signal s_i by s_i and a realization of neighbourhood $B(i)$ by $\mathfrak{B}(i)$ for all i . Given the equilibrium σ and the realization s_1, \dots, s_{K-1} and $\mathfrak{B}(1), \dots, \mathfrak{B}(K - 1)$, all decisions x_1, \dots, x_{K-1} are recursively defined [*i.e.* they are non-stochastic; see the definition of the information set in equation (3.1)]. Therefore, for any $\mathfrak{B}(n)$ that satisfies $\max_{b \in \mathfrak{B}(n)} b < K$, the decision x_n is also defined. By the definition of z_n [cf. equation (A5)], this implies that

$$\begin{aligned} \mathbb{P}_\sigma(x_n = \theta \mid s_i = s_i, B(i) = \mathfrak{B}(i), \text{ for } i = 1, \dots, K - 1, n) \\ \leq \mathbb{P}_\sigma(z_n = \theta \mid s_i = s_i, B(i) = \mathfrak{B}(i), \text{ for } i = 1, \dots, K - 1, n). \end{aligned}$$

By integrating over all possible $s_1, \dots, s_{K-1}, s_n, \mathfrak{B}(1), \dots, \mathfrak{B}(K - 1)$, this yields

$$\mathbb{P}_\sigma(x_n = \theta \mid B(n) = \mathfrak{B}(n)) \leq \mathbb{P}_\sigma(z_n = \theta \mid B(n) = \mathfrak{B}(n)),$$

for any $\mathfrak{B}(n)$ that satisfies $\max_{b \in \mathfrak{B}(n)} b < K$. By integrating over all $\mathfrak{B}(n)$ that satisfy this condition, we obtain

$$\mathbb{P}_\sigma\left(x_n = \theta \mid \max_{b \in B(n)} b < K\right) \leq \mathbb{P}_\sigma\left(z_n = \theta \mid \max_{b \in B(n)} b < K\right). \tag{A6}$$

Moreover, since the sequence of neighbourhoods $\{B(i)\}_{i \in \mathbb{N}}$ is independent of θ and the sequence of private signals $\{s_i\}_{i \in \mathbb{N}}$, it follows from equation (A5) that the decision z_n is given by

$$z_n = \operatorname{argmax}_{y \in \{0,1\}} \mathbb{P}_\sigma(y = \theta \mid s_1, \dots, s_{K-1}, s_n). \tag{A7}$$

Therefore, z_n is also independent of the sequence of neighbourhoods $\{B(i)\}_{i \in \mathbb{N}}$ and we have

$$\mathbb{P}_\sigma(z_n = \theta \mid \max_{b \in B(n)} b < K) = \mathbb{P}_\sigma(z_n = \theta).$$

Since the private signals have the same distribution, it follows from equation (A7) that for any $n, m \geq K$, the random variables z_n and z_m have identical probability distributions. Hence, for any $n \geq K$, equation (A7) implies that

$$\mathbb{P}_\sigma(z_n = \theta) = \mathbb{P}_\sigma(z_K = \theta).$$

Combining the preceding two relations with equation (A6), we have for any $n \geq K$,

$$\mathbb{P}_\sigma\left(x_n = \theta \mid \max_{b \in B(n)} b < K\right) \leq \mathbb{P}_\sigma\left(z_n = \theta \mid \max_{b \in B(n)} b < K\right) = \mathbb{P}_\sigma(z_n = \theta) = \mathbb{P}_\sigma(z_K = \theta).$$

Substituting this relation in equation (A4), we obtain for any $n \in \mathcal{N}, n \geq K$,

$$\mathbb{P}_\sigma(x_n = \theta) \leq 1 - \varepsilon + \varepsilon \mathbb{P}_\sigma(z_K = \theta).$$

Therefore,

$$\liminf_{n \rightarrow \infty} \mathbb{P}_\sigma(x_n = \theta) \leq 1 - \varepsilon + \varepsilon \mathbb{P}_\sigma(z_K = \theta). \tag{A8}$$

We finally show that in view of the assumption that \mathbb{F}_0 and \mathbb{F}_1 are absolutely continuous with respect to each other [which implies $\mathbb{P}_\sigma(x_1 = \theta) < 1$], we have $\mathbb{P}_\sigma(z_K = \theta) < 1$ for any given K . If $\mathbb{P}_\sigma(x_1 = \theta) < 1$ holds, then we have either $\mathbb{P}_\sigma(x_1 = \theta \mid \theta = 1) < 1$ or $\mathbb{P}_\sigma(x_1 = \theta \mid \theta = 0) < 1$. Assume without loss of generality that we have

$$\mathbb{P}_\sigma(x_1 = \theta \mid \theta = 1) < 1. \tag{A9}$$

Let $\bar{\mathcal{S}}_\sigma$ denote the set of all private signals such that if $s_1 \in \bar{\mathcal{S}}_\sigma$, then $x_1 = 0$ in equilibrium σ . Since the first agent's decision is a function of s_1 , then equation (A9) is equivalent to

$$\mathbb{P}_\sigma(s_1 \in \bar{\mathcal{S}}_\sigma \mid \theta = 1) > 0.$$

Since the private signals are conditionally independent given θ , this implies that

$$\mathbb{P}_\sigma(s_i \in \bar{S}_\sigma \text{ for all } i \leq K \mid \theta = 1) = \mathbb{P}_\sigma(s_1 \in \bar{S}_\sigma \mid \theta = 1)^K > 0. \quad (\text{A10})$$

We next show that if $s_i \in \bar{S}_\sigma$ for all $i \leq K$, then $z_K = 0$. Using Bayes' rule, we have

$$\begin{aligned} \mathbb{P}_\sigma(\theta = 0 \mid s_i \in \bar{S}_\sigma \text{ for all } i \leq K) &= \left[1 + \frac{\mathbb{P}_\sigma(s_i \in \bar{S}_\sigma \text{ for all } i \leq K \mid \theta = 1)}{\mathbb{P}_\sigma(s_i \in \bar{S}_\sigma \text{ for all } i \leq K \mid \theta = 0)} \right]^{-1} \\ &= \left[1 + \frac{\prod_{i=1}^K \mathbb{P}_\sigma(s_i \in \bar{S}_\sigma \mid \theta = 1)}{\prod_{i=1}^K \mathbb{P}_\sigma(s_i \in \bar{S}_\sigma \mid \theta = 0)} \right]^{-1} \\ &= \left[1 + \left(\frac{\mathbb{P}_\sigma(s_1 \in \bar{S}_\sigma \mid \theta = 1)}{\mathbb{P}_\sigma(s_1 \in \bar{S}_\sigma \mid \theta = 0)} \right)^K \right]^{-1}, \end{aligned} \quad (\text{A11})$$

where the second equality follows from the conditional independence of the private signals and the third equality holds since private signals are identically distributed. Applying Bayes' rule on the second term in parentheses in equation (A11), this implies that

$$\mathbb{P}_\sigma(\theta = 0 \mid s_i \in \bar{S}_\sigma \text{ for all } i \leq K) = \left[1 + \left(\frac{1}{\mathbb{P}_\sigma(\theta = 0 \mid s_1 \in \bar{S}_\sigma)} - 1 \right)^K \right]^{-1}. \quad (\text{A12})$$

Since $s_1 \in \bar{S}_\sigma$ induces $x_1 = 0$, we have $\mathbb{P}_\sigma(\theta = 0 \mid s_1 \in \bar{S}_\sigma) \geq 1/2$. Because the function on the R.H.S. of equation (A12) is non-decreasing in $\mathbb{P}_\sigma(\theta = 0 \mid s_1 \in \bar{S}_\sigma)$ for any value in $[1/2, 1]$, we obtain

$$\mathbb{P}_\sigma(\theta = 0 \mid s_i \in \bar{S}_\sigma \text{ for all } i \leq K) \geq \frac{1}{2}.$$

By the definition of z_K , this implies that if $s_i \in \bar{S}_\sigma$ for all $i \leq K$, then $z_K = 0$ (we can let z_n be equal to 0 whenever both states are equally likely given the private signals). Combined with the fact that the event $\{s_i \in \bar{S}_\sigma \text{ for all } i \leq K\}$ has positive conditional probability given $\theta = 1$ under measure \mathbb{P}_σ [cf. equation (A10)], this implies that $\mathbb{P}_\sigma(z_K = \theta) < 1$. Substituting this relation in equation (A8), we have

$$\liminf_{n \rightarrow \infty} \mathbb{P}_\sigma(x_n = \theta) < 1,$$

thus showing that asymptotic learning does not occur. \parallel

Proof of Theorem 2.

The proof follows by combining several lemmas, which we next present.

As a first step, we characterize certain important properties of private belief distributions.

Lemma A1. *For any private belief distributions $(\mathbb{G}_0, \mathbb{G}_1)$, the following relations hold.*

(a) *For all $r \in (0, 1)$, we have*

$$\frac{d\mathbb{G}_0}{d\mathbb{G}_1}(r) = \frac{1-r}{r}.$$

(b) *We have*

$$\begin{aligned} \mathbb{G}_0(r) &\geq \left(\frac{1-r}{r} \right) \mathbb{G}_1(r) + \frac{r-z}{2} \mathbb{G}_1(z) \text{ for all } 0 < z < r < 1, \\ 1 - \mathbb{G}_1(r) &\geq (1 - \mathbb{G}_0(r)) \left(\frac{r}{1-r} \right) + \frac{w-r}{2} (1 - \mathbb{G}_1(w)) \text{ for all } 0 < r < w < 1. \end{aligned}$$

(c) *The ratio $\mathbb{G}_0(r)/\mathbb{G}_1(r)$ is non-increasing in r and $\mathbb{G}_0(r)/\mathbb{G}_1(r) > 1$ for all $r \in (\underline{\beta}, \bar{\beta})$.*

Proof. See Appendix B. \parallel

We next show that the *ex ante* probability of an agent making the correct decision (and thus his expected pay-off) is no less than the probability of any of the agents in his realized neighbourhood making the correct decision.

Lemma A2 (Information Monotonicity). Let $\sigma \in \Sigma$ be an equilibrium. For any agent n and neighborhood \mathfrak{B} , we have

$$\mathbb{P}_\sigma(x_n = \theta \mid B(n) = \mathfrak{B}) \geq \max_{b \in \mathfrak{B}} \mathbb{P}_\sigma(x_b = \theta).$$

Proof. See Appendix B in Supplementary Material. \parallel

Let us next focus on a specific network topology where each agent observes the decision of a single agent. This will enable us to establish a preliminary version of the *strong improvement principle*, which provides a lower bound on the increase in the *ex ante* probability that an individual will make a correct decision over his neighbour's probability. We will then generalize this result to arbitrary social networks.

For each n and strategy profile σ , let us define Y_n^σ and N_n^σ as the probabilities of agent n making the correct decision conditional on state θ . More formally, these are defined as

$$Y_n^\sigma = \mathbb{P}_\sigma(x_n = 1 \mid \theta = 1), \quad N_n^\sigma = \mathbb{P}_\sigma(x_n = 0 \mid \theta = 0).$$

The unconditional probability of a correct decision is then

$$\frac{1}{2}(Y_n^\sigma + N_n^\sigma) = \mathbb{P}_\sigma(x_n = \theta). \tag{A13}$$

We also define the thresholds L_n^σ and U_n^σ in terms of these probabilities:

$$L_n^\sigma = \frac{1 - N_n^\sigma}{1 - N_n^\sigma + Y_n^\sigma}, \quad U_n^\sigma = \frac{N_n^\sigma}{N_n^\sigma + 1 - Y_n^\sigma}. \tag{A14}$$

The next lemma shows that the equilibrium decisions are fully characterized in terms of these thresholds.

Lemma A3. Let $B(n) = \{b\}$ for some agent n . Let $\sigma \in \Sigma$ be an equilibrium, and let L_b^σ and U_b^σ be given by equation (A14). Then, agent n 's decision x_n in equilibrium σ satisfies

$$x_n = \begin{cases} 0, & \text{if } p_n < L_b^\sigma, \\ x_b, & \text{if } p_n \in (L_b^\sigma, U_b^\sigma), \\ 1, & \text{if } p_n > U_b^\sigma. \end{cases}$$

The proof is omitted since it is an immediate application of Proposition 2 [use Bayes' rule to determine $\mathbb{P}_\sigma(\theta = 1 \mid x_b = j)$ for each $j \in \{0, 1\}$].

Note that the sequence $\{(U_n, L_n)\}$ only depends on $\{(Y_n, N_n)\}$ and is thus deterministic. This reflects the fact that each individual recognizes the amount of information that will be contained in the action of the previous agent, which determines his own decision thresholds. Individual actions are still stochastic since they are determined by whether the individual's private belief is below L_b , above U_b , or in between (see Figure 2).

Using the structure of the equilibrium decision rule, the next lemma provides an expression for the probability of agent n making the correct decision conditional on his observing agent $b < n$, in terms of the private belief distributions and the thresholds L_b^σ and U_b^σ .

Lemma A4. Let $B(n) = \{b\}$ for some agent n . Let $\sigma \in \Sigma$ be an equilibrium, and let L_b^σ and U_b^σ be given by equation (A14). Then,

$$\begin{aligned} \mathbb{P}_\sigma(x_n = \theta \mid B(n) = \{b\}) &= \frac{1}{2}[\mathbb{G}_0(L_b^\sigma) + (\mathbb{G}_0(U_b^\sigma) - \mathbb{G}_0(L_b^\sigma))N_b^\sigma + (1 - \mathbb{G}_1(U_b^\sigma)) + (\mathbb{G}_1(U_b^\sigma) - \mathbb{G}_1(L_b^\sigma))Y_b^\sigma]. \end{aligned}$$

Proof. By definition, agent n receives the same expected utility from all his possible equilibrium choices. We can thus compute the expected utility by supposing that the agent will choose $x_n = 0$ when indifferent. Then, the expected utility of agent n (the probability of the correct decision) can be written as follows:

$$\begin{aligned} \mathbb{P}_\sigma(x_n = \theta \mid B(n) = \{b\}) &= \mathbb{P}_\sigma(p_n \leq L_b^\sigma \mid \theta = 0)\mathbb{P}(\theta = 0) + \mathbb{P}_\sigma(p_n \in (L_b^\sigma, U_b^\sigma], x_b = 0 \mid \theta = 0)\mathbb{P}(\theta = 0) \\ &\quad + \mathbb{P}_\sigma(p_n > U_b^\sigma \mid \theta = 1)\mathbb{P}(\theta = 1) + \mathbb{P}_\sigma(p_n \in (L_b^\sigma, U_b^\sigma], x_b = 1 \mid \theta = 1)\mathbb{P}(\theta = 1). \end{aligned}$$




FIGURE 2

The equilibrium decision rule when observing a single agent illustrated on the private belief space

The result then follows using the fact that p_n and x_b are conditionally independent given θ and the notation for the private belief distributions [cf. equation (3.4)]. \parallel

Using the previous lemma, we next provide a lower bound on the amount of improvement in the *ex ante* probability of making the correct decision between an agent and his neighbour.

Lemma A5. *Let $B(n) = \{b\}$ for some agent n . Let $\sigma \in \Sigma$ be an equilibrium and let L_b^σ and U_b^σ be given by equation (A14). Then,*

$$\mathbb{P}_\sigma(x_n = \theta \mid B(n) = \{b\}) \geq \mathbb{P}_\sigma(x_b = \theta) + \frac{(1 - N_b^\sigma)L_b^\sigma}{8} \mathbb{G}_1\left(\frac{L_b^\sigma}{2}\right) + \frac{(1 - Y_b^\sigma)(1 - U_b^\sigma)}{8} \left[1 - \mathbb{G}_0\left(\frac{1 + U_b^\sigma}{2}\right)\right].$$

Proof. In Lemma A1(b), let $r = L_b^\sigma$, $z = L_b^\sigma/2$, so that we obtain

$$(1 - N_b^\sigma)\mathbb{G}_0(L_b^\sigma) \geq Y_b^\sigma \mathbb{G}_1(L_b^\sigma) + \frac{(1 - N_b^\sigma)L_b^\sigma}{4} \mathbb{G}_1\left(\frac{L_b^\sigma}{2}\right).$$

Next, again using Lemma A1(b) and letting $r = U_b^\sigma$ and $w = (1 + U_b^\sigma)/2$, we have

$$(1 - Y_b^\sigma)[1 - \mathbb{G}_1(U_b^\sigma)] \geq N_b^\sigma [1 - \mathbb{G}_0(\frac{1}{2})] + \frac{(1 - Y_b^\sigma)(1 - U_b^\sigma)}{4} \left[1 - \mathbb{G}_0\left(\frac{1 + U_b^\sigma}{2}\right)\right].$$

Combining the preceding two relations with Lemma A4 and using the fact that $Y_b^\sigma + N_b^\sigma = 2P_\sigma(x_b = \theta)$ [cf. equation (A13)], the desired result follows. \parallel

The next lemma establishes that the lower bound on the amount of improvement in the *ex ante* probability is uniformly bounded away from 0 for unbounded private beliefs and when $\mathbb{P}_\sigma(x_b = \theta) < 1$, *i.e.* when asymptotic learning is not achieved.

Lemma A6. *Let $B(n) = \{b\}$ for some n . Let $\sigma \in \Sigma$ be an equilibrium and denote $\alpha = \mathbb{P}_\sigma(x_b = \theta)$. Then,*

$$\mathbb{P}_\sigma(x_n = \theta \mid B(n) = \{b\}) \geq \alpha + \frac{(1 - \alpha)^2}{8} \min\left\{\mathbb{G}_1\left(\frac{1 - \alpha}{2}\right), 1 - \mathbb{G}_0\left(\frac{1 + \alpha}{2}\right)\right\}.$$

Proof. We consider two cases separately.

Case 1: $N_b^\sigma \leq \alpha$. From the definition of L_b^σ and the fact that $Y_b^\sigma = 2\alpha - N_b^\sigma$ [cf. equation (A13)], we have

$$L_b^\sigma = \frac{1 - N_b^\sigma}{1 - 2N_b^\sigma + 2\alpha}.$$

Since σ is an equilibrium, we have $\alpha \geq 1/2$, and thus the R.H.S. of the preceding inequality is a non-increasing function of N_b^σ . Since $N_b^\sigma \leq \alpha$, this relation therefore implies that $L_b^\sigma \geq 1 - \alpha$. Combining the relations $1 - N_b^\sigma \geq 1 - \alpha$ and $L_b^\sigma \geq 1 - \alpha$, we obtain

$$\frac{(1 - N_b^\sigma)L_b^\sigma}{8} \mathbb{G}_1\left(\frac{L_b^\sigma}{2}\right) \geq \frac{(1 - \alpha)^2}{8} \mathbb{G}_1\left(\frac{1 - \alpha}{2}\right). \tag{A15}$$

Case 2: $N_b^\sigma \geq \alpha$. Since $Y_b^\sigma + N_b^\sigma = 2\alpha$, this implies that $Y_b^\sigma \leq \alpha$. Using the definition of U_b^σ and a similar argument as the one above, we obtain

$$\frac{(1 - Y_b^\sigma)(1 - U_b^\sigma)}{8} \left[1 - \mathbb{G}_0\left(\frac{1 + U_b^\sigma}{2}\right)\right] \geq \frac{(1 - \alpha)^2}{8} \left[1 - \mathbb{G}_0\left(\frac{1 + \alpha}{2}\right)\right]. \tag{A16}$$

Combining equations (A15) and (A16), we obtain

$$\begin{aligned} &\frac{(1 - N_b^\sigma)L_b^\sigma}{8} \mathbb{G}_1\left(\frac{L_b^\sigma}{2}\right) + \frac{(1 - Y_b^\sigma)(1 - U_b^\sigma)}{8} \left[1 - \mathbb{G}_0\left(\frac{1 + U_b^\sigma}{2}\right)\right] \\ &\geq \frac{(1 - \alpha)^2}{8} \min\left\{\mathbb{G}_1\left(\frac{1 - \alpha}{2}\right), 1 - \mathbb{G}_0\left(\frac{1 + \alpha}{2}\right)\right\}, \end{aligned}$$

where we also used the fact that each term on the L.H.S. of the preceding inequality is non-negative. Substituting this into Lemma A5, the desired result follows. \parallel

The preceding lemma characterizes the improvements in the probability of making the correct decision between an agent and his neighbour. To study the limiting behaviour of these improvements, we introduce the function $\bar{\mathcal{Z}} : [1/2, 1] \rightarrow [1/2, 1]$ defined by

$$\bar{\mathcal{Z}}(\alpha) = \alpha + \frac{(1 - \alpha)^2}{8} \min\left\{\mathbb{G}_1\left(\frac{1 - \alpha}{2}\right), 1 - \mathbb{G}_0\left(\frac{1 + \alpha}{2}\right)\right\}. \tag{A17}$$

Lemma A6 establishes that for n , which has $B(n) = \{b\}$, we have

$$\mathbb{P}_\sigma(x_n = \theta | B(n) = \{b\}) \geq \bar{\mathcal{Z}}(\mathbb{P}_\sigma(x_b = \theta)), \tag{A18}$$

i.e. the function $\bar{\mathcal{Z}}$ acts as an improvement function for the evolution of the probability of making the correct decision. The function $\bar{\mathcal{Z}}(\cdot)$ has several important properties, which are formally stated in the next lemma.

Lemma A7. *The function $\bar{\mathcal{Z}} : [1/2, 1] \rightarrow [1/2, 1]$ given in equation (A17) satisfy the following properties:*

(a) *The function $\bar{\mathcal{Z}}$ has no upwards jumps. That is, for any $\alpha \in [1/2, 1]$,*

$$\bar{\mathcal{Z}}(\alpha) = \lim_{r \uparrow \alpha} \bar{\mathcal{Z}}(r) \geq \lim_{r \downarrow \alpha} \bar{\mathcal{Z}}(r).$$

(b) *For any $\alpha \in [1/2, 1]$, $\bar{\mathcal{Z}}(\alpha) \geq \alpha$.*

(c) *If the private beliefs are unbounded, then for any $\alpha \in [1/2, 1)$, $\bar{\mathcal{Z}}(\alpha) > \alpha$.*

Proof. Since \mathbb{G}_0 and \mathbb{G}_1 are cumulative distribution functions, they cannot have downward jumps, i.e. for each $j \in \{0, 1\}$, $\lim_{r \uparrow \alpha} \mathbb{G}_j(r) \leq \lim_{r \downarrow \alpha} \mathbb{G}_j(r)$ for any $\alpha \in [1/2, 1]$, establishing Part (a). Part (b) follows from the fact that cumulative distribution functions take values in $[0, 1]$. For Part (c), suppose that for some $\alpha \in [1/2, 1)$, $\bar{\mathcal{Z}}(\alpha) = \alpha$. This implies that

$$\min\left\{\mathbb{G}_1\left(\frac{1 - \alpha}{2}\right), 1 - \mathbb{G}_0\left(\frac{1 + \alpha}{2}\right)\right\} = 0. \tag{A19}$$

However, from the assumption on the private beliefs, we have that for all $\alpha \in (0, 1)$ and any $j \in \{0, 1\}$, $\mathbb{G}_j(\alpha) \in (0, 1)$, contradicting equation (A19). \parallel

The properties of the $\bar{\mathcal{Z}}$ function will be used in the analysis of asymptotic learning in general networks in the last step of the proof of Theorem 2. The analysis of asymptotic learning requires the relevant improvement function to be both continuous and monotone. However, $\bar{\mathcal{Z}}$ does not necessarily satisfy these properties. We next construct a

related function $\mathcal{Z} : [1/2, 1] \rightarrow [1/2, 1]$ that satisfies these properties and can be used as the improvement function in the asymptotic analysis. Let \mathcal{Z} be defined as follows:

$$\mathcal{Z}(\alpha) = \frac{1}{2} \left(\alpha + \sup_{r \in [1/2, \alpha]} \bar{\mathcal{Z}}(r) \right). \quad (\text{A20})$$

This function shares the same “improvement” properties as $\bar{\mathcal{Z}}$ but is also non-decreasing and continuous. The properties of the function $\mathcal{Z}(\cdot)$ are stated in the following lemma.

Lemma A8. *The function $\mathcal{Z} : [1/2, 1] \rightarrow [1/2, 1]$ given in equation (A20) satisfies the following properties:*

- (a) For any $\alpha \in [1/2, 1]$, $\mathcal{Z}(\alpha) \geq \alpha$.
- (b) If the private beliefs are unbounded, then for any $\alpha \in [1/2, 1)$, $\mathcal{Z}(\alpha) > \alpha$.
- (c) The function \mathcal{Z} is increasing and continuous.

Proof. Parts (a) and (b) follow immediately from Lemma A7, Parts (b) and (c), respectively. The function $\sup_{r \in [1/2, \alpha]} \bar{\mathcal{Z}}(r)$ is non-decreasing and the function α is increasing, therefore the average of these two functions, which is \mathcal{Z} , is an increasing function, establishing the first part of Part (c).

We finally show that \mathcal{Z} is a continuous function. We first show $\mathcal{Z}(\alpha)$ is continuous for all $\alpha \in [1/2, 1)$. To obtain a contradiction, assume that \mathcal{Z} is discontinuous at some $\alpha^* \in [1/2, 1)$. This implies that $\sup_{r \in [1/2, \alpha]} \bar{\mathcal{Z}}(r)$ is discontinuous at α^* . Since $\sup_{r \in [1/2, \alpha]} \bar{\mathcal{Z}}(r)$ is a non-decreasing function, we have

$$\lim_{\alpha \downarrow \alpha^*} \sup_{r \in [1/2, \alpha]} \bar{\mathcal{Z}}(r) > \sup_{r \in [1/2, \alpha^*]} \bar{\mathcal{Z}}(r),$$

from which it follows that there exists some $\varepsilon > 0$ such that for any $\delta > 0$

$$\sup_{r \in [1/2, \alpha^* + \delta]} \bar{\mathcal{Z}}(r) > \bar{\mathcal{Z}}(\alpha) + \varepsilon \quad \text{for all } \alpha \in [1/2, \alpha^*].$$

This contradicts the fact that the function $\bar{\mathcal{Z}}$ does not have an upward jump [cf. Lemma A7(a)] and establishes the continuity of $\mathcal{Z}(\alpha)$ for all $\alpha \in [1/2, 1)$. The continuity of the function $\mathcal{Z}(\alpha)$ at $\alpha = 1$ follows from Part (a). \parallel

The next lemma shows that the function \mathcal{Z} is also a (strong) improvement function for the evolution of the probability of making the correct decision.

Lemma A9 (Strong Improvement Principle). *Let $B(n) = \{b\}$ for some n . Let $\sigma \in \Sigma$ be an equilibrium. Then, we have*

$$\mathbb{P}_\sigma(x_n = \theta \mid B(n) = \{b\}) \geq \mathcal{Z}(\mathbb{P}_\sigma(x_b = \theta)). \quad (\text{A21})$$

Proof. Let α denote $\mathbb{P}_\sigma(x_b = \theta)$. If $\mathcal{Z}(\alpha) = \alpha$, then the result is immediate. Suppose next that $\mathcal{Z}(\alpha) > \alpha$. This implies that $\mathcal{Z}(\alpha) < \sup_{r \in [1/2, \alpha]} \bar{\mathcal{Z}}(r)$. Therefore, there exists some $\bar{\alpha} \in [1/2, \alpha]$ such that

$$\bar{\mathcal{Z}}(\bar{\alpha}) > \mathcal{Z}(\alpha). \quad (\text{A22})$$

We next show that $\mathbb{P}_\sigma(x_n = \theta \mid B(n) = b) \geq \bar{\mathcal{Z}}(\bar{\alpha})$. Agent n can always (privately) make the information from his observation of x_b coarser (*i.e.* not observe x_b according to some probability). Let the observation thus generated by agent n be denoted by \tilde{x}_b , and suppose that it is given by

$$\tilde{x}_b = \begin{cases} x_b, & \text{with probability } (2\bar{\alpha} - 1)/(2\alpha - 1), \\ 0, & \text{with probability } (\alpha - \bar{\alpha})/(2\alpha - 1), \\ 1, & \text{with probability } (\alpha - \bar{\alpha})/(2\alpha - 1), \end{cases}$$

where the realizations of \tilde{x}_b are independent from agent n 's information set. Next observe that $\mathbb{P}_\sigma(\tilde{x}_b = \theta) = \bar{\alpha}$. Then, Lemma A6 implies that $\mathbb{P}_\sigma(x_n = \theta \mid B(n) = b) \geq \bar{\mathcal{Z}}(\bar{\alpha})$. Since $\bar{\mathcal{Z}}(\bar{\alpha}) > \mathcal{Z}(\alpha)$ [cf. equation (A22)], the desired result follows. \parallel

We next generalize the results presented so far to an arbitrary network topology. We first present an information monotonicity result with the amount of improvement given by the improvement function \mathcal{Z} defined in equation (A21).

Even though a full characterization of equilibrium decisions in general network topologies is a non-tractable problem, it is possible to establish an analogue of Lemma A9, i.e. a *generalized strong improvement principle*, which provides a lower bound on the amount of increase in the probabilities of making the correct decision.

Lemma A10 (Generalized Strong Improvement Principle). *For any $n \in \mathbb{N}$, any set $\mathfrak{B} \subseteq \{1, \dots, n-1\}$ and any equilibrium $\sigma \in \Sigma$, we have*

$$\mathbb{P}_\sigma(x_n = \theta \mid B(n) = \mathfrak{B}) \geq \mathcal{Z} \left(\max_{b \in \mathfrak{B}} \mathbb{P}_\sigma(x_b = \theta) \right).$$

Proof. Given an equilibrium $\sigma \in \Sigma$ and agent n , let h_σ be a function that maps any subset of $\{1, \dots, n-1\}$ to an element of $\{1, \dots, n-1\}$ such that for any $\mathfrak{B} \subseteq \{1, \dots, n-1\}$, we have

$$h_\sigma(\mathfrak{B}) \in \operatorname{argmax}_{b \in \mathfrak{B}} \mathbb{P}_\sigma(x_b = \theta). \tag{A23}$$

We define w_n as the decision that maximizes the conditional probability of making the correct decision given the private signal s_n and the decision of the agent $h_\sigma(B(n))$, i.e.

$$w_n \in \operatorname{argmax}_{y \in \{0,1\}} \mathbb{P}_\sigma(y = \theta \mid s_n, x_{h_\sigma(B(n))}).$$

The equilibrium decision x_n of agent n satisfies

$$\mathbb{P}_\sigma(x_n = \theta \mid s_n, B(n), x_k, k \in B(n)) \geq \mathbb{P}_\sigma(w_n = \theta \mid s_n, B(n), x_k, k \in B(n)),$$

[cf. the characterization of the equilibrium decision rule in equation (3.2)]. Integrating over all possible private signals and decisions of neighbours, we obtain for any $\mathfrak{B} \subseteq \{1, \dots, n-1\}$,

$$\mathbb{P}_\sigma(x_n = \theta \mid B(n) = \mathfrak{B}) \geq \mathbb{P}_\sigma(w_n = \theta \mid B(n) = \mathfrak{B}). \tag{A24}$$

Because w_n is an optimal choice given a single observation, equation (A21) holds and yields

$$\mathbb{P}_\sigma(w_n = \theta \mid B(n) = \mathfrak{B}) \geq \mathcal{Z}(\mathbb{P}_\sigma(x_{h_\sigma(\mathfrak{B})} = \theta)). \tag{A25}$$

Combining equations (A23), (A24), and (A25), we obtain the desired result. \parallel

This lemma is the key step in our proof. It shows that, under unbounded private beliefs, there are improvements in pay-offs (probabilities of making correct decisions) that are bounded away from zero. We will next use this generalized strong improvement principle to prove Theorem 2. The proof involves showing that under the expanding observations and the unbounded private beliefs assumptions, the amount of improvement in the probabilities of making the correct decision given by \mathcal{Z} accumulates until asymptotic learning is reached.

Proof of Theorem 2. The proof consists of two parts. In the first part of the proof, we construct two sequences $\{a_k\}$ and $\{\phi_k\}$ such that for all $k \geq 0$, there holds

$$\mathbb{P}_\sigma(x_n = \theta) \geq \phi_k \text{ for all } n \geq a_k. \tag{A26}$$

The second part of the proof shows that ϕ_k converges to 1, thus establishing the result.

Given some integer $K > 0$ and scalar $\varepsilon > 0$, let $N(K, \varepsilon) > 0$ be an integer such that for all $n \geq N(K, \varepsilon)$,

$$\mathbb{Q}_n \left(\max_{b \in B(n)} b < K \right) < \varepsilon,$$

(such an integer exists in view of the fact that, by hypothesis, the network topology features expanding observations). We let $a_1 = 1$ and $\phi_1 = 1/2$ and define the sequences $\{a_k\}$ and ϕ_k recursively by

$$a_{k+1} = N \left(a_k, \frac{1}{2} \left[1 - \frac{\phi_k}{\mathcal{Z}(\phi_k)} \right] \right), \quad \phi_{k+1} = \frac{\phi_k + \mathcal{Z}(\phi_k)}{2}.$$

Using the fact that the range of the function \mathcal{Z} is $[1/2, 1]$, it can be seen that $\phi_k \in [1/2, 1]$ for all k , therefore the preceding sequences are well-defined.

We use induction on the index k to prove relation (A26). Since σ is an equilibrium, we have

$$\mathbb{P}_\sigma(x_n = \theta) \geq \frac{1}{2} \quad \text{for all } n \geq 1,$$

which together with $\alpha_1 = 1$ and $\phi_1 = 1/2$ shows relation (A26) for $k = 1$. Assume that the relation (A26) holds for an arbitrary k , *i.e.*

$$\mathbb{P}_\sigma(x_j = \theta) \geq \phi_k \quad \text{for all } j \geq \alpha_k. \quad (\text{A27})$$

Consider some agent n with $n \geq \alpha_{k+1}$. By integrating the relation from Lemma A10 over all possible neighbourhoods $B(n)$, we obtain

$$\mathbb{P}_\sigma(x_n = \theta) \geq \mathbb{E}_{B(n)} \left[\mathcal{Z} \left(\max_{b \in B(n)} \mathbb{P}_\sigma(x_b = \theta) \right) \right],$$

where $\mathbb{E}_{B(n)}$ denotes the expectation with respect to the neighbourhood $B(n)$ (*i.e.* the weighted sum over all possible neighbourhoods $B(n)$). We can rewrite the preceding as follows:

$$\begin{aligned} \mathbb{P}_\sigma(x_n = \theta) &\geq \mathbb{E}_{B(n)} \left[\mathcal{Z} \left(\max_{b \in B(n)} \mathbb{P}_\sigma(x_b = \theta) \right) \mid \max_{b \in B(n)} b \geq \alpha_k \right] \mathbb{Q}_n \left(\max_{b \in B(n)} b \geq \alpha_k \right) \\ &\quad + \mathbb{E}_{B(n)} \left[\mathcal{Z} \left(\max_{b \in B(n)} \mathbb{P}_\sigma(x_b = \theta) \right) \mid \max_{b \in B(n)} b < \alpha_k \right] \mathbb{Q}_n \left(\max_{b \in B(n)} b < \alpha_k \right). \end{aligned}$$

Since the terms on the R.H.S. of the preceding relation are non-negative, this implies that

$$\mathbb{P}_\sigma(x_n = \theta) \geq \mathbb{E}_{B(n)} \left[\mathcal{Z} \left(\max_{b \in B(n)} \mathbb{P}_\sigma(x_b = \theta) \right) \mid \max_{b \in B(n)} b \geq \alpha_k \right] \mathbb{Q}_n \left(\max_{b \in B(n)} b \geq \alpha_k \right).$$

Since $\max_{b \in B(n)} b \geq \alpha_k$, equation (A27) implies that

$$\max_{b \in B(n)} \mathbb{P}_\sigma(x_b = \theta) \geq \phi_k.$$

Since the function \mathcal{Z} is non-decreasing [cf. Lemma A8(c)], combining the preceding two relations, we obtain

$$\begin{aligned} \mathbb{P}_\sigma(x_n = \theta) &\geq \mathbb{E}_{B(n)} \left[\mathcal{Z}(\phi_k) \mid \max_{b \in B(n)} b \geq \alpha_k \right] \mathbb{Q}_n \left(\max_{b \in B(n)} b \geq \alpha_k \right) \\ &= \mathcal{Z}(\phi_k) \mathbb{Q}_n \left(\max_{b \in B(n)} b \geq \alpha_k \right), \end{aligned}$$

where the equality follows since the sequence $\{\phi_k\}$ is deterministic. Using the definition of α_k , this implies that

$$\mathbb{P}_\sigma(x_n = \theta) \geq \mathcal{Z}(\phi_k) \frac{1}{2} \left[1 + \frac{\phi_k}{\mathcal{Z}(\phi_k)} \right] = \phi_{k+1},$$

thus completing the induction.

We finally prove that $\phi_k \rightarrow 1$ as $k \rightarrow \infty$. Since $\mathcal{Z}(a) \geq a$ for all $a \in [1/2, 1]$ [cf. Lemma A8(a)], it follows from the definition of ϕ_k that $\{\phi_k\}_{k \in \mathbb{N}}$ is a non-decreasing sequence. It is also bounded and therefore it converges to some ϕ^* . Taking the limit in the definition of ϕ_k , we obtain

$$2\phi^* = 2 \lim_{k \rightarrow \infty} \phi_k = \lim_{k \rightarrow \infty} [\phi_k + \mathcal{Z}(\phi_k)] = \phi^* + \mathcal{Z}(\phi^*),$$

where the third equality follows since \mathcal{Z} is a continuous function [cf. Lemma A8(c)]. This shows that $\phi^* = \mathcal{Z}(\phi^*)$, *i.e.* ϕ^* is a fixed point of \mathcal{Z} . Since the private beliefs are unbounded, the unique fixed point of \mathcal{Z} is 1, showing that $\phi_k \rightarrow 1$ as $k \rightarrow \infty$ and completing the proof. \parallel

Proofs of Corollaries 1 and 2.

Proof of Corollary 1. We first show that if $c \geq 1$, then the network topology has non-expanding observations. To show this, we set $K = 1$ in Definition 5 and show that the probability of infinitely many agents having empty neighbourhoods is uniformly bounded away from 0. We first consider the case $c > 1$. Then, the probability that the neighbourhood of agent $n + 1$ is the empty set is given by

$$\mathbb{Q}_{n+1}(B(n+1) = \emptyset) = \left(1 - \frac{a}{n^c} \right)^n,$$

which converges to 1 as n goes to infinity. If $c = 1$, then

$$\lim_{n \rightarrow \infty} \mathbb{Q}_{n+1}(B(n+1) = \emptyset) = \lim_{n \rightarrow \infty} \left(1 - \frac{a}{n}\right)^n = e^{-a}.$$

Therefore, for infinitely many agents, $\mathbb{Q}_{n+1}(B(n+1) = \emptyset) \geq e^{-a}/2$. The preceding show that the network topology has non-expanding observations for $c \geq 1$, hence the result follows from Theorem 1.

We next assume that $c < 1$. For any K and all $n \geq K$, we have

$$\mathbb{Q}_{n+1} \left(\max_{b \in B(n+1)} b \leq K \right) = \left(1 - \frac{a}{n^c}\right)^{n-K},$$

which converges to 0 as n goes to infinity. Hence, the network topology is expanding in observations and the result follows from Theorem 2. \parallel

Proof of Corollary 2. We show that $\{L(n)\}_{n \in \mathbb{N}}$ goes to infinity if and only if the deterministic sequence $\{\max_{b \in B(n)} b\}_{n \in \mathbb{N}}$ goes to infinity. Suppose first $\{L(n)\}_{n \in \mathbb{N}}$ diverges to infinity. Then, for every K , there exists N such that for all $n \geq N$, $L(n) \geq K$. Note that

$$L(n) \leq 1 + \max_{b \in B(n)} b$$

because the longest information path must be a subset of the sequence $(1, 2, \dots, \max_{b \in B(n)} b, n)$. So, for $n \geq N$, if $L(n) \geq K$, then $\max_{b \in B(n)} b > K$, thus proving the first part of the lemma. Suppose next that $\{\max_{b \in B(n)} b\}_{n \in \mathbb{N}}$ goes to infinity as n goes to infinity. We show by induction that for each $d \in \mathbb{N}$, there exists some integer c_d such that $L(n) \geq d$ for all $n \geq c_d$. Since $L(n) \geq 1$ for all n , then $c_1 = 1$. Assume such c_d exists for some d . Then, we show that such a c_{d+1} also exists. Since $\{\max_{b \in B(n)} b\}_{n \in \mathbb{N}}$ goes to infinity, there exists some N_d such that for all $n \geq N_d$,

$$\max_{b \in B(n)} b \geq c_d.$$

Now, for any $n \geq N_d$, there exists a path with size d up to some $k \geq c_d$ and then another observation from k to n , therefore $L(n) \geq d + 1$. Hence, $c_{d+1} = N_d$. \parallel

Proof of Theorem 4.

This proof consists of two parts. In the first part, we show that agents who observe the full history of actions are able to learn the state of the world. This part of the proof is centred on the martingale convergence of beliefs for agents who observe the entire history. The second part of the proof uses the improvement principle to extend the asymptotic learning result to all agents. Note that the conditions of the theorem do not hold if the set S is finite. We, therefore, assume that S has infinite cardinality without loss of generality.

Part 1. For each n , let $x^n = (x_1, \dots, x_n)$ represent the sequence of decisions up to and including x_n . Let $q^*(x^n)$ denote the ‘‘social belief’’ when x^n is observed under equilibrium σ , i.e.

$$q^*(x^n) = \mathbb{P}_\sigma(\theta = 1 | x^n).$$

The social belief $q^*(x^n)$ is a martingale and, by the martingale convergence theorem, converges with probability 1 to some random variable \hat{q} . Conditional on $\theta = 1$, the likelihood ratio

$$\frac{1 - q^*(x^n)}{q^*(x^n)} = \frac{\mathbb{P}_\sigma(\theta = 0 | x^n)}{\mathbb{P}_\sigma(\theta = 1 | x^n)}$$

is also a martingale [see Doob, 1953, equation (7.12)]. Therefore, conditional on $\theta = 1$, the ratio $(1 - q^*(x^n))/q^*(x^n)$ converges with probability 1 to some random variable $(1 - \hat{q}_1)/\hat{q}_1$. In particular, we have

$$\mathbb{E}_\sigma \left[\frac{1 - \hat{q}_1}{\hat{q}_1} \right] < \infty,$$

(see Breiman (1968), Theorem 5.14], and therefore $\hat{q}_1 > 0$ with probability 1. Similarly, $q^*(x^n)/(1 - q^*(x^n))$ is a martingale conditional on $\theta = 0$ and converges with probability 1 to some random variable $\hat{q}_0/(1 - \hat{q}_0)$, where $\hat{q}_0 < 1$ with probability 1. Therefore,

$$\mathbb{P}_\sigma(\hat{q} > 0 | \theta = 1) = 1 \text{ and } \mathbb{P}_\sigma(\hat{q} < 1 | \theta = 0) = 1. \tag{A28}$$

The key element of Part 1 of the proof is to show that the support of \hat{q} is contained in the set $\{0, 1\}$. This fact combined with equation (A28) guarantees that $\hat{q} = \theta$ (i.e. the agents that observe the entire history eventually know what the state of the world θ is).

To show that the support of \hat{q} is contained in $\{0, 1\}$, we study the evolution dynamics of $q^*(x^n)$. Suppose $x_{n+1} = 0$. Using Bayes' rule twice, we have

$$\begin{aligned} q^*((x^n, 0)) &= \frac{\mathbb{P}_\sigma(x_{n+1} = 0, x^n | \theta = 1)}{\sum_{j=0}^1 \mathbb{P}_\sigma(x_{n+1} = 0, x^n | \theta = j)} = \left[1 + \frac{\mathbb{P}_\sigma(x_{n+1} = 0, x^n | \theta = 0)}{\mathbb{P}_\sigma(x_{n+1} = 0, x^n | \theta = 1)} \right]^{-1} \\ &= \left[1 + \left(\frac{1}{q^*(x^n)} - 1 \right) \frac{\mathbb{P}_\sigma(x_{n+1} = 0 | \theta = 0, x^n)}{\mathbb{P}_\sigma(x_{n+1} = 0 | \theta = 1, x^n)} \right]^{-1}. \end{aligned}$$

To simplify notation, let

$$f_n(x^n) = \frac{\mathbb{P}_\sigma(x_{n+1} = 0 | \theta = 0, x^n)}{\mathbb{P}_\sigma(x_{n+1} = 0 | \theta = 1, x^n)}, \quad (\text{A29})$$

so that

$$q^*((x^n, 0)) = \left[1 + \left(\frac{1}{q^*(x^n)} - 1 \right) f_n(x^n) \right]^{-1}. \quad (\text{A30})$$

We next show that if $q^*(x^n) \geq 1/2$ for some $n+1 \in S$, we have

$$f_n(x^n) \geq 1 + \delta \quad \text{for some } \delta > 0, \quad (\text{A31})$$

which will allow us to establish a bound on the difference between $q^*(x^n)$ and $q^*((x^n, 0))$. By conditioning on the neighbourhood $B(n+1)$, we have for any $j \in \{0, 1\}$,

$$\mathbb{P}_\sigma(x_{n+1} = 0 | \theta = j, x^n) = \sum_{\bar{B}} \mathbb{Q}_{n+1}(B(n+1) = \bar{B}) \mathbb{P}_\sigma(x_{n+1} = 0 | \theta = j, x^n, B(n+1) = \bar{B}). \quad (\text{A32})$$

Define the auxiliary measure $\tilde{\mathbb{Q}}_{n+1}$ over the subsets of $\{1, \dots, n\}$ to have the following value for any $\bar{B} \subseteq \{1, \dots, n\}$,

$$\tilde{\mathbb{Q}}_{n+1}(\bar{B}) = \begin{cases} \mathbb{Q}_{n+1}(B(n+1) = \bar{B}), & \text{if } \bar{B} \neq \{1, \dots, n\}; \\ \mathbb{Q}_{n+1}(B(n+1) = \bar{B}) - \varepsilon, & \text{if } \bar{B} = \{1, \dots, n\}, \end{cases}$$

where $\varepsilon > 0$ is given by the lower bound on the probability of observing the history of actions for $n+1 \in S$. Note that by the definition of ε , $\tilde{\mathbb{Q}}_{n+1}$ is a non-negative measure. Rewriting equation (A32) in terms of $\tilde{\mathbb{Q}}_{n+1}$, we obtain

$$\begin{aligned} \mathbb{P}_\sigma(x_{n+1} = 0 | \theta = j, x^n) &= \varepsilon \mathbb{P}_\sigma(x_{n+1} = 0 | \theta = j, x^n, B(n+1) = \{1, \dots, n\}) \\ &\quad + \sum_{\bar{B}} \tilde{\mathbb{Q}}_{n+1}(\bar{B}) \mathbb{P}_\sigma(x_{n+1} = 0 | \theta = j, x^n, B(n+1) = \bar{B}). \end{aligned} \quad (\text{A33})$$

Let $\tilde{q}(\bar{B}, x^n)$ denote the "social belief" when the subset \bar{B} of actions of x^n is observed under equilibrium σ , *i.e.*

$$\tilde{q}(\bar{B}, x^n) = \mathbb{P}_\sigma(\theta = 1 | x_k \text{ for all } k \in \bar{B}).$$

Using the assumption that agents break ties in favour of Action 0, Proposition 2 implies that for any $\bar{B} \subseteq \{1, \dots, n\}$,

$$\mathbb{P}_\sigma(x_{n+1} = 0 | \theta = j, x^n, B(n+1) = \bar{B}) = \mathbb{G}_j(1 - \tilde{q}(\bar{B}, x^n)).$$

Hence, equation (A33) can be rewritten as follows:

$$\mathbb{P}_\sigma(x_{n+1} = 0 | \theta = j, x^n) = \varepsilon \mathbb{G}_j(1 - q^*(x^n)) + \sum_{\bar{B}} \tilde{\mathbb{Q}}_{n+1}(\bar{B}) \mathbb{G}_j(1 - \tilde{q}(\bar{B}, x^n)).$$

Returning to equation (A29), we have that

$$f_n(x^n) = \frac{\varepsilon \mathbb{G}_0(1 - q^*(x^n)) + \sum_{\bar{B}} \tilde{\mathbb{Q}}_{n+1}(\bar{B}) \mathbb{G}_0(1 - \tilde{q}(\bar{B}, x^n))}{\varepsilon \mathbb{G}_1(1 - q^*(x^n)) + \sum_{\bar{B}} \tilde{\mathbb{Q}}_{n+1}(\bar{B}) \mathbb{G}_1(1 - \tilde{q}(\bar{B}, x^n))}.$$

From Lemma A1(c), we know that $\mathbb{G}_0(r) \geq \mathbb{G}_1(r)$ for any r and, therefore,

$$f_n(x^n) \geq \frac{\varepsilon \mathbb{G}_0(1 - q^*(x^n)) + \sum_{\bar{B}} \tilde{\mathbb{Q}}_{n+1}(\bar{B}) \mathbb{G}_1(1 - \tilde{q}(\bar{B}, x^n))}{\varepsilon \mathbb{G}_1(1 - q^*(x^n)) + \sum_{\bar{B}} \tilde{\mathbb{Q}}_{n+1}(\bar{B}) \mathbb{G}_1(1 - \tilde{q}(\bar{B}, x^n))}.$$

We also obtain from Lemma A1(c) that

$$\mathbb{G}_0(1 - q^*(x^n)) \geq \frac{\mathbb{G}_0(1/2)}{\mathbb{G}_1(1/2)} \mathbb{G}_1(1 - q^*(x^n))$$

since $q^*(x^n)$ is assumed to be greater than or equal to $1/2$. Therefore,

$$f_n(x^n) \geq \frac{\frac{\mathbb{G}_0(1/2)}{\mathbb{G}_1(1/2)} \mathbb{G}_1(1 - q^*(x^n)) + \sum_{\bar{B}} \tilde{Q}_{n+1}(\bar{B}) \mathbb{G}_1(1 - \tilde{q}(\bar{B}, x^n))}{\varepsilon \mathbb{G}_1(1 - q^*(x^n)) + \sum_{\bar{B}} \tilde{Q}_{n+1}(\bar{B}) \mathbb{G}_1(1 - \tilde{q}(\bar{B}, x^n))}.$$

Note that since $\frac{\mathbb{G}_0(1/2)}{\mathbb{G}_1(1/2)} > 1$, the R.H.S. of the equation above is increasing in $\mathbb{G}_1(1 - q^*(x^n))$ and decreasing in $\sum_{\bar{B}} \tilde{Q}_{n+1}(\bar{B}) \mathbb{G}_1(1 - \tilde{q}(\bar{B}, x^n))$. Since $\mathbb{G}_1(1 - q^*(x^n)) \leq \mathbb{G}_1(1/2)$ and $\sum_{\bar{B}} \tilde{Q}_{n+1}(\bar{B}) \mathbb{G}_1(1 - \tilde{q}(\bar{B}, x^n)) \geq \sum_{\bar{B}} \tilde{Q}_{n+1}(\bar{B}) = 1 - \varepsilon$, we obtain

$$f_n(x^n) \geq \frac{\varepsilon \mathbb{G}_0(1/2) + 1 - \varepsilon}{\varepsilon \mathbb{G}_1(1/2) + 1 - \varepsilon} = 1 + \delta \quad \text{for some } \delta > 0 \text{ if } n + 1 \in S,$$

thus proving equation (A31).

We now show that the support of the limiting social belief \hat{q} does not include the interval $(1/2, 1)$. Combining equation (A31) with equation (A30) yields

$$q^*((x^n, 0)) \leq \left[1 + \left(\frac{1}{q^*(x^n)} - 1 \right) (1 + \delta) \right]^{-1} \quad \text{if } q^*(x^n) \geq 1/2 \text{ and } n + 1 \in S. \tag{A34}$$

Suppose now $q^*(x^n) \in [1/2, 1 - \varepsilon]$ for some $\varepsilon > 0$. We show that there exists some constant $K(\delta, \varepsilon) > 0$ such that

$$q^*(x^n) - q^*((x^n, 0)) \geq K(\delta, \varepsilon) \text{ if } n + 1 \in S. \tag{A35}$$

Define $g : [1/2, 1 - \varepsilon] \rightarrow [0, 1]$ as

$$g(q) = q - \left[1 + \left(\frac{1}{q} - 1 \right) (1 + \delta) \right]^{-1}.$$

It can be seen that $g(q)$ is a concave function over $q \in [1/2, 1 - \varepsilon]$. Let $K(\delta, \varepsilon)$ be

$$K(\delta, \varepsilon) = \inf_{q \in [1/2, 1 - \varepsilon]} g(q) = \min\{g(1/2), g(1 - \varepsilon)\} = \min \left\{ \frac{\delta}{2(2 + \delta)}, \frac{\varepsilon \delta (1 - \varepsilon)}{1 + \varepsilon \delta} \right\} > 0.$$

From equation (A34), it follows that

$$q^*(x^n) - q^*((x^n, 0)) \geq q^*(x^n) - \left[1 + \left(\frac{1}{q^*(x^n)} - 1 \right) (1 + \delta) \right]^{-1} \geq g(q^*(x^n)) \geq K(\delta, \varepsilon),$$

thus proving equation (A35).

Recall that $q^*(x^n)$ converges to \hat{q} with probability 1. We show that for any $\varepsilon > 0$, the support of \hat{q} does not contain $(1/2, 1 - \varepsilon)$. Assume, to arrive at a contradiction, that it does. Consider a sample path that converges to a value in the interval $(1/2, 1 - \varepsilon)$. For this sample path, there exists some N such that for all $n \geq N$, $q^*(x^n) \in [1/2, 1 - \varepsilon]$. By the Borel–Cantelli lemma, there are infinitely many agents n within S that observe a non-persuasive neighbourhood C_i for some i because the neighbourhoods are independent and $\sum_{n \in S} \sum_{i=1}^M r_i(n) = \infty$. Since these are non-persuasive neighbourhoods, an infinite subset will choose Action 0. Therefore, for infinitely many n that satisfy $n \geq N$ and $n \in S$, by equation (A35),

$$q^*(x^{n+1}) \leq q^*(x^n) - K(\delta, \varepsilon).$$

But this implies the sequence $q^*(x^n)$ is not Cauchy and contradicts the fact that $q^*(x^n)$ is a convergent sequence. Hence, we conclude that the support of \hat{q} does not contain $(1/2, 1 - \varepsilon)$. Since this argument holds for any $\varepsilon > 0$, the support of \hat{q} cannot contain $(1/2, 1)$. A similar argument leads to the conclusion that the support of \hat{q} does not include $(0, 1/2]$. Therefore, the support of \hat{q} is a subset of the set $\{0, 1\}$. By equation (A28), this implies that $\hat{q} = \theta$ with probability 1. Combining the result above with Proposition 2,

$$\begin{aligned} \lim_{n \rightarrow \infty} \mathbb{P}_\sigma(x_n = \theta | B(n) = \{1, \dots, n - 1\}) &= \lim_{n \rightarrow \infty} \mathbb{P}_\sigma(\lfloor p_n + q_n \rfloor = \theta | B(n) = \{1, \dots, n - 1\}) \\ &= \lim_{n \rightarrow \infty} \mathbb{P}_\sigma(\lfloor p_n + q^*(x^{n-1}) \rfloor = \theta) \\ &= \lim_{n \rightarrow \infty} \mathbb{P}_\sigma(\lfloor p_n + \theta \rfloor = \theta) = 1, \end{aligned}$$

where p_n and q_n are, respectively, the private and social beliefs of agent n . We thus conclude that the subsequence of agents that observe the full history of actions eventually learn the state of the world θ .

Part 2. In this part of the proof, we show that asymptotic learning occurs. We first show that agents in S asymptotically learn the state θ . In the previous part, we obtained that $\mathbb{P}_\sigma(x_n = \theta | B(n) = \{1, \dots, n-1\})$ converges to 1 as n goes to infinity. Define $\{\psi_n\}_{n \in S}$ to be a non-decreasing sequence of numbers such that $\lim_{n \rightarrow \infty} \psi_n = 1$ and for all $n \in S$,

$$\mathbb{P}_\sigma(x_n = \theta | B(n) = \{1, \dots, n-1\}) \geq \psi_n.$$

Define also a sequence of random variables $\{D_n\}_{n \in S}$ such that

$$D_n = \begin{cases} 0, & \text{with probability } 1, \text{ if } B(n) \neq \{1, 2, \dots, n-1\}; \\ 0, & \text{with probability } 1 - \varepsilon, \text{ if } B(n) = \{1, 2, \dots, n-1\}; \\ 1, & \text{with probability } \varepsilon, \text{ if } B(n) = \{1, 2, \dots, n-1\}, \end{cases}$$

where ε is given by the lower bound on the probability of observing the history of actions for S . Let D_n be independent of all other events conditionally on $B(n)$. Conditioning the event $x_n = \theta$ on D_n , we obtain that for any $n \in S$,

$$\begin{aligned} \mathbb{P}_\sigma(x_n = \theta) &= \mathbb{P}_\sigma(x_n = \theta | D_n = 0)(1 - \varepsilon) + \mathbb{P}_\sigma(x_n = \theta | D_n = 1)\varepsilon \\ &\geq \mathbb{P}_\sigma(x_n = \theta | D_n = 0)(1 - \varepsilon) + \psi_n \varepsilon. \end{aligned} \quad (\text{A36})$$

Note that expanding observations with respect to S hold conditionally on $D_n = 0$ since for any K ,

$$0 = \lim_{n \rightarrow \infty} \mathbb{Q}_n \left(\max_{b \in B(n) \cap S} b < K \right) \geq \lim_{n \rightarrow \infty} \mathbb{Q}_n \left(\max_{b \in B(n) \cap S} b < K \mid D_n = 0 \right) (1 - \varepsilon) \geq 0.$$

We can thus define a function $N(K, \delta)$ such that for all $K \in \mathbb{N}$, $\delta > 0$, and $n \geq N(K, \delta)$,

$$\mathbb{Q}_n \left(\max_{b \in B(n) \cap S} b < K \mid D_n = 0 \right) < \delta. \quad (\text{A37})$$

We now show inductively that agents in S asymptotically learn the state θ . We construct two sequences $\{\phi_K\}_{K \in \mathbb{N}}$ and $\{\alpha_K\}_{K \in \mathbb{N}}$ such that for all $n \in S$ and all $K \in \mathbb{N}$,

$$\mathbb{P}_\sigma(x_n = \theta) \geq \phi_K \text{ for all } n \geq \alpha_K. \quad (\text{A38})$$

Let $\alpha_1 = 1$ and $\phi_1 = 1/2$ and the equation above hold for $K = 1$. Assume equation (A38) holds for some K , and let

$$\phi_{K+1} = \phi_K + \frac{\varepsilon}{2}(1 - \phi_K). \quad (\text{A39})$$

We now show that there exists some α_{K+1} such that for all $n \in S$ and $n \geq \alpha_{K+1}$, $\mathbb{P}_\sigma(x_n = \theta) \geq \phi_{K+1}$. From equation (A36), we have that for any $n \in S$,

$$\begin{aligned} \mathbb{P}_\sigma(x_n = \theta) &\geq \mathbb{P}_\sigma(x_n = \theta | D_n = 0)(1 - \varepsilon) + \psi_n \varepsilon \\ &\geq \mathbb{E}_\sigma \left[\max_{b \in B(n) \cap S} \mathbb{P}_\sigma(x_b = \theta) | D_n = 0 \right] (1 - \varepsilon) + \psi_n \varepsilon, \end{aligned}$$

where the second inequality follows from the fact that agent n can select a better action (*ex ante*) than anyone in his neighbourhood (cf. Lemma A2). From equations (A37) and (A38), we obtain that for any $\delta > 0$, $n \in S$, and $n \geq N(\alpha_K, \delta)$,

$$\mathbb{P}_\sigma(x_n = \theta) \geq \phi_K (1 - \delta)(1 - \varepsilon) + \psi_n \varepsilon.$$

By selecting a n large (in order to have ψ_n close to 1) and a small δ , we can have $\phi_K (1 - \delta)(1 - \varepsilon) + \psi_n \varepsilon$ be arbitrarily close to $\phi_K + \varepsilon(1 - \phi_K)$. Therefore, there exists some α_{K+1} that satisfies equation (A39). Therefore, equation (A38) holds. To complete the argument that agents in S asymptotically learn the state note that $\lim_{K \rightarrow \infty} \phi_K = 1$ from the recursion in equation (A39) and

$$\lim_{n \rightarrow \infty, n \in S} \mathbb{P}_\sigma(x_n = \theta) \geq \lim_{K \rightarrow \infty} \phi_K = 1.$$

We now show that all agents asymptotically learn the state θ . From Lemma A2, we obtain that

$$\begin{aligned} \mathbb{P}_\sigma(x_n = \theta) &\geq \mathbb{E}_\sigma \left[\max_{b \in B(n) \cap S} \mathbb{P}_\sigma(x_b = \theta) \right] \\ &\geq \mathbb{E}_\sigma \left[\max_{b \in B(n) \cap S} \mathbb{P}_\sigma(x_b = \theta) \mid \max_{b \in B(n) \cap S} b \geq K \right] \mathbb{Q}_n \left(\max_{b \in B(n) \cap S} b \geq K \right), \end{aligned}$$

for any $K \in \mathbb{N}$. Since agents in S learn the state, for any K ,

$$\lim_{n \rightarrow \infty} \mathbb{E}_\sigma \left[\max_{b \in B(n) \cap S} \mathbb{P}_\sigma(x_b = \theta) \mid \max_{b \in B(n) \cap S} b \geq K \right] = 1.$$

From the expanding observations condition with respect to S we obtain that for all K ,

$$\lim_{n \rightarrow \infty} \mathbb{Q}_n \left(\max_{b \in B(n) \cap S} b \geq K \right) = 1.$$

Therefore, $\lim_{n \rightarrow \infty} \mathbb{P}_\sigma(x_n = \theta) = 1$, completing the proof. \parallel

Proofs of Propositions 3 and 4.

Proof of Proposition 3. Using Bayes' rule and the assumption of equal priors on the state θ , we have that the social belief given by observing neighbourhood B , with actions $x_k = y_k$ for each $k \in B$ in equilibrium σ is given by

$$\mathbb{P}_\sigma(\theta = 1 \mid x_k = y_k \text{ for all } k \in B) = \left(1 + \frac{\mathbb{P}_\sigma(x_k = y_k \text{ for all } k \in B \mid \theta = 0)}{\mathbb{P}_\sigma(x_k = y_k \text{ for all } k \in B \mid \theta = 1)} \right)^{-1}. \tag{A40}$$

Since the neighbourhoods of the agents in B are all empty, their actions are independent conditional on the state θ . Hence,

$$\mathbb{P}_\sigma(\theta = 1 \mid x_k = y_k \text{ for all } k \in B) = \left(1 + \frac{\prod_{k \in B} \mathbb{P}_\sigma(x_k = y_k \mid \theta = 0)}{\prod_{k \in B} \mathbb{P}_\sigma(x_k = y_k \mid \theta = 1)} \right)^{-1}.$$

Given the assumption that all agents break ties in favour of Action 0, we can obtain that for each agent n with an empty neighbourhood, and each $j \in \{0, 1\}$, $\mathbb{P}_\sigma(x_1 = 0 \mid \theta = j) = \mathbb{G}_j(1/2)$ and $\mathbb{P}_\sigma(x_1 = 1 \mid \theta = j) = 1 - \mathbb{G}_j(1/2)$. Therefore, the social belief of equation (A40) can be rewritten as follows:

$$\mathbb{P}_\sigma(\theta = 1 \mid x_k = y_k \text{ for all } k \in B) = \left(1 + \frac{(1 - \mathbb{G}_0(1/2))^{\sum_{k \in B} y_k} \mathbb{G}_0(1/2)^{|B| - \sum_{k \in B} y_k}}{(1 - \mathbb{G}_1(1/2))^{\sum_{k \in B} y_k} \mathbb{G}_1(1/2)^{|B| - \sum_{k \in B} y_k}} \right)^{-1}. \tag{A41}$$

From Lemma A1(c), we have that

$$\frac{\mathbb{G}_0(1/2)}{\mathbb{G}_1(1/2)} > 1 > \frac{1 - \mathbb{G}_0(1/2)}{1 - \mathbb{G}_1(1/2)},$$

which combined with equation (A41) implies

$$\begin{aligned} &\min_{B \subseteq \{1, \dots, K\}, \{y_k\}_{k \leq K} \in \{0, 1\}^K} \mathbb{P}_\sigma(\theta = 1 \mid x_k = y_k \text{ for all } k \in B) \\ &= \mathbb{P}_\sigma(\theta = 1 \mid x_k = 0 \text{ for all } k \in \{1, \dots, K\}) = \left(1 + \frac{\mathbb{G}_0(1/2)^K}{\mathbb{G}_1(1/2)^K} \right)^{-1} \text{ and} \\ &\max_{B \subseteq \{1, \dots, K\}, \{y_k\}_{k \leq K} \in \{0, 1\}^K} \mathbb{P}_\sigma(\theta = 1 \mid x_k = y_k \text{ for all } k \in B) \\ &= \mathbb{P}_\sigma(\theta = 1 \mid x_k = 1 \text{ for all } k \in \{1, \dots, K\}) = \left(1 + \frac{(1 - \mathbb{G}_0(1/2))^K}{(1 - \mathbb{G}_1(1/2))^K} \right)^{-1}. \end{aligned}$$

Therefore, any such neighbourhood $B \subseteq \{1, \dots, K\}$ is non-persuasive if it satisfies the conditions

$$\left(1 + \frac{\mathbb{G}_0(1/2)^K}{\mathbb{G}_1(1/2)^K} \right)^{-1} > 1 - \bar{\beta} \text{ and } \left(1 + \frac{(1 - \mathbb{G}_0(1/2))^K}{(1 - \mathbb{G}_1(1/2))^K} \right)^{-1} < 1 - \underline{\beta}, \tag{A42}$$

which together yield equation (4.5). \parallel

Proof of Proposition 4. We start by showing inductively that all networks of the form $\{1, \dots, n\}$ for any $n \leq K$ are non-persuasive. In particular, we first show that for all $n \leq K$, and any set of values $y_k \in \{0, 1\}$ for each $k \leq n$,

$$\mathbb{P}_\sigma(\theta = 1 \mid x_k = y_k \text{ for all } k \leq n) \geq \left(1 + \left(\frac{\mathbb{G}_0(1/2)}{\mathbb{G}_1(1/2)}\right)^n\right)^{-1}. \quad (\text{A43})$$

The upper bound on $\mathbb{P}_\sigma(\theta = 1 \mid x_k = y_k \text{ for all } k \leq n)$ is symmetric and we do not prove it to avoid repetition.

Note that equation (A43) is trivially true for $n = 0$, i.e. $\mathbb{P}_\sigma(\theta = 1) = 1/2$. Assume that equation (A43) is true for $n - 1$ as an inductive hypothesis. We now show it is true for n . Let the social belief given the neighbourhood $\{1, \dots, n\}$ and decisions $y^n = (y_1, \dots, y_n)$ be given by $q^*(y^n) = \mathbb{P}_\sigma(\theta = 1 \mid x_k = y_k \text{ for all } k \leq n)$. Then, using Bayes' rule, we obtain that

$$\frac{1}{q^*(y^n)} - 1 = \left(\frac{1}{q^*(y^{n-1})} - 1\right) \frac{\mathbb{P}_\sigma(x_n = y_n \mid \theta = 0, x_k = y_k \text{ for all } k \leq n-1)}{\mathbb{P}_\sigma(x_n = y_n \mid \theta = 1, x_k = y_k \text{ for all } k \leq n-1)}.$$

Since agent n observes the full history of past actions and breaks ties in favor of action 0, we have

$$\frac{1}{q^*(y^n)} - 1 = \begin{cases} \left(\frac{1}{q^*(y^{n-1})} - 1\right) \frac{1 - \mathbb{G}_0(1 - q^*(y^{n-1}))}{1 - \mathbb{G}_1(1 - q^*(y^{n-1}))}, & \text{if } y_n = 1; \\ \left(\frac{1}{q^*(y^{n-1})} - 1\right) \frac{\mathbb{G}_0(1 - q^*(y^{n-1}))}{\mathbb{G}_1(1 - q^*(y^{n-1}))}, & \text{if } y_n = 0. \end{cases}$$

If $y_n = 1$, then $q^*(y^n) \geq q^*(y^{n-1})$ because

$$\frac{1 - \mathbb{G}_0(1 - r)}{1 - \mathbb{G}_1(1 - r)} \leq 1 \quad \text{for any } r \geq \underline{\beta},$$

by Lemma A1(c). Using the inductive hypothesis,

$$q^*(y^n) \geq q^*(y^{n-1}) \geq \left(1 + \left(\frac{\mathbb{G}_0(1/2)}{\mathbb{G}_1(1/2)}\right)^{n-1}\right)^{-1} \geq \left(1 + \left(\frac{\mathbb{G}_0(1/2)}{\mathbb{G}_1(1/2)}\right)^n\right)^{-1},$$

thus proving equation (A43) for $y_n = 1$. If $y_n = 0$, then we need to consider two separate cases. Suppose first $q^*(y^{n-1}) < 1/2$. Then, by Lemma A1(c), we have that

$$\frac{\mathbb{G}_0(1 - q^*(y^{n-1}))}{\mathbb{G}_1(1 - q^*(y^{n-1}))} \leq \frac{\mathbb{G}_0(1/2)}{\mathbb{G}_1(1/2)}.$$

Combining the relation above with the inductive hypothesis, we obtain

$$\left(\frac{1}{q^*(y^{n-1})} - 1\right) \frac{\mathbb{G}_0(1 - q^*(y^{n-1}))}{\mathbb{G}_1(1 - q^*(y^{n-1}))} \leq \left(\frac{1}{q^*(y^{n-1})} - 1\right) \frac{\mathbb{G}_0(1/2)}{\mathbb{G}_1(1/2)} \leq \left(\frac{\mathbb{G}_0(1/2)}{\mathbb{G}_1(1/2)}\right)^n,$$

thus proving that equation (A43) holds in this case as well. Finally, consider the case where $q^*(y^{n-1}) \geq 1/2$. Here, we use the monotonicity assumption from equation (4.8) to show that

$$\frac{1}{q^*(y^n)} - 1 = \left(\frac{1}{q^*(y^{n-1})} - 1\right) \frac{\mathbb{G}_0(1 - q^*(y^{n-1}))}{\mathbb{G}_1(1 - q^*(y^{n-1}))} \leq \left(\frac{1}{1/2} - 1\right) \frac{\mathbb{G}_0(1/2)}{\mathbb{G}_1(1/2)} \leq \left(\frac{\mathbb{G}_0(1/2)}{\mathbb{G}_1(1/2)}\right)^n,$$

thus completing the proof of equation (A43). We thus conclude that any neighbourhood $\{1, \dots, n\}$ is non-persuasive if $n \leq K$ and K satisfies

$$\left(1 + \left(\frac{\mathbb{G}_0(1/2)}{\mathbb{G}_1(1/2)}\right)^K\right)^{-1} > 1 - \bar{\beta},$$

as well as the symmetric (upper bound)

$$\left(1 + \left(\frac{1 - \mathbb{G}_0(1/2)}{1 - \mathbb{G}_1(1/2)}\right)^K\right)^{-1} < 1 - \underline{\beta}.$$

Both conditions on K combined yield the condition of equation (4.7).

We have proved that the set $\{1, \dots, K\}$ is non-persuasive given the conditions of the proposition. We now show that any subset of $B \subseteq \{1, \dots, K\}$ is also non-persuasive. To prove this result, we bound by the worst possible event in terms of the actions of agents not in B with respect to the social belief, *i.e.*

$$\begin{aligned} & \mathbb{P}_\sigma(\theta = 1 \mid x_k = y_k \text{ for all } k \in B) \\ &= \sum_{z^n \in \{0,1\}^{\{1,\dots,K\} \setminus B}} \mathbb{P}_\sigma(\theta = 1 \mid x_k = y_k \text{ for } k \in B, x_k = z_k \text{ for } k \in \{1, \dots, K\} \setminus B) \\ & \quad \cdot \mathbb{P}_\sigma(x_k = z_k \text{ for all } k \in \{1, \dots, K\} \setminus B \mid x_k = y_k \text{ for all } k \in B) \\ &\geq \min_{z^n \in \{0,1\}^{\{1,\dots,K\} \setminus B}} \mathbb{P}_\sigma(\theta = 1 \mid x_k = y_k \text{ for } k \in B, x_k = z_k \text{ for } k \in \{1, \dots, K\} \setminus B). \end{aligned}$$

Use a symmetric bound to show the maximum value of the social belief. The proof is, therefore, complete since we have previously proven that the neighbourhood $\{1, \dots, K\}$ is non-persuasive. \square

Acknowledgments. We thank three anonymous referees and Andrea Prat, Lones Smith, and Peter Sorensen for useful comments and suggestions. We gratefully acknowledge financial support from the Air Force Office of Scientific Research grant (FA9550-09-1-0420) and the National Science Foundation grants (NSF SES-0729361 and EFRI-0735956).

SUPPLEMENTARY DATA

Supplementary Material is available at *Review of Economic Studies* online.

REFERENCES

BALA, V. and GOYAL, S. (1998), “Learning from Neighbours”, *The Review of Economic Studies*, **65**, 595–621.
 BALA, V. and GOYAL, S. (2001), “Conformism and Diversity under Social Learning”, *Economic Theory*, **17**, 101–120.
 BANERJEE, A. (1992), “A Simple Model of Herd Behavior”, *The Quarterly Journal of Economics*, **107**, 797–817.
 BANERJEE, A. and FUDENBERG, D. (2004), “Word-of-mouth Learning”, *Games and Economic Behavior*, **46**, 1–22.
 BESLEY, T. and CASE, A. (1994), “Diffusion as a Learning Process: Evidence from HYV Cotton” (Working Paper No. 228, Princeton University, Woodrow Wilson School of Public and International Affairs, Research Program in Development Studies).
 BIKHCHANDANI, S., HIRSHLEIFER, D. and WELCH, I. (1992), “A Theory of Fads, Fashion, Custom, and Cultural Change as Information Cascades”, *The Journal of Political Economy*, **100**, 992–1026.
 BIKHCHANDANI, S., HIRSHLEIFER, D. and WELCH, I. (1998), “Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades”, *The Journal of Economic Perspectives*, **12**, 151–170.
 BREIMAN, L. (1968), *Probability* (Reading, MA: Addison-Wellesley).
 CALLANDER, S. and HORNER, J. (2009), “The Wisdom of the Minority”, *Journal of Economic Theory*, **144**, 1421–1439.
 CELEN, B. and KARIV, S. (2004), “Observational Learning under Imperfect Information”, *Games and Economic Behavior*, **47**, 72–86.
 CHAMLEY, C. and GALE, D. (1994), “Information Revelation and Strategic Delay in a Model of Investment”, *Econometrica*, **62**, 1065–1086.
 CONDORCET, N. C. (1785), *Essai sur l’Application de l’Analyse à la Probabilité des Décisions Rendues à la Pluralité des Voix* (Paris, France: Imprimerie Royale).
 COVER, T. (1969), “Hypothesis Testing with Finite Statistics”, *The Annals of Mathematical Statistics*, **40**, 828–835.
 DEMARZO, P. M., VAYANOS, D. and ZWIEBEL, J. (2003), “Persuasion Bias, Social Influence, and Unidimensional Opinions”, *The Quarterly Journal of Economics*, **118**, 909–968.
 DOOB, J. L. (1953), *Stochastic Processes* (New York: John Wiley and Sons Inc.).
 ELLISON, G. and FUDENBERG, D. (1993), “Rules of Thumb for Social Learning”, *The Journal of Political Economy*, **101**, 612–643.
 ELLISON, G. and FUDENBERG, D. (1995), “Word-of-mouth Communication and Social Learning”, *The Quarterly Journal of Economics*, **110**, 93–126.
 FOSTER, A. and ROSENZWEIG, M. (1995), “Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture”, *The Journal of Political Economy*, **103**, 1176–1209.
 GALE, D. and KARIV, S. (2003), “Bayesian Learning in Social Networks”, *Games and Economic Behavior*, **45**, 329–346.
 GOLUB, B. and JACKSON, M. O. (2010), “Naive Learning in Social Networks: Convergence, Influence, and the Wisdom of Crowds”, *American Economic Journal: Microeconomics*, **2**, 112–149.
 GRANOVETTER, M. (1973), “The Strength of Weak Ties”, *The American Journal of Sociology*, **78**, 1360–1380.

- IOANNIDES, M. and LOURY, L. (2004), "Job Information Networks, Neighborhood Effects, and Inequality", *The Journal of Economic Literature*, **42**, 1056–1093.
- JACKSON, M. O. (2006), "The Economics of Social Networks", in Blundell, R., Newey, W. and Persson, T. (eds) *Chapter 1 in Volume I of Advances in Economics and Econometrics, Theory and Applications: Ninth World Congress of the Econometric Society* (Cambridge: Cambridge University Press).
- JACKSON, M. O. (2007), "The Study of Social Networks in Economics", in Rauch, J. E. (ed) *The Missing Links: Formation and Decay of Economic Networks* (New York: Russell Sage Foundation).
- JACKSON, M. O. and ROGERS, B. (2007), "Meeting Strangers and Friends of Friends: How Random are Socially Generated Networks?" *American Economic Review*, **97**, 890–915.
- LEE, I. (1993), "On the Convergence of Informational Cascades", *The Journal of Economic Theory*, **61**, 395–411.
- LORENZ, J., MARCINISZYN, M., and STEGER, A. (2007), "Observational Learning in Random Networks", *Lecture Notes in Computer Science*, Vol. 4539 (Springer Berlin Heidelberg), 574–588.
- MONTGOMERY, J. D. (1991), "Social Networks and Labor-Market Outcomes: Toward an Economic Analysis", *The American Economic Review*, **81**, 1408–1418.
- MUNSHI, K. (2003), "Networks in the Modern Economy: Mexican Migrants in the U.S. Labor Market", *The Quarterly Journal of Economics*, **118**, 549–597.
- MUNSHI, K. (2004), "Social Learning in a Heterogeneous Population: Technology Diffusion in the Indian Green Revolution", *The Journal of Development Economics*, **73**, 185–215.
- PAPASTAVROU, J. D. and M. ATHANS (1992), "Distributed Detection by a Large Team of Sensors in Tandem", *IEEE Transactions on Aerospace and Electronic Systems*, **28**, 639–653.
- SGROI, D. (2002), "Optimizing Information in the Herd: Guinea Pigs, Profits, and Welfare", *Games and Economic Behavior*, **39**, 137–166.
- SMITH, L. and SORENSEN, P. (2000), "Pathological Outcomes of Observational Learning", *Econometrica*, **68**, 371–398.
- SMITH, L. and SORENSEN, P. (2008), "Rational Social Learning with Random Sampling" (Working Paper: <http://lonessmith.com/sites/default/files/rational.pdf>).
- TAY, W. P., TSITSIKLIS, J. N., and WIN, M. Z. (2008), "On the Sub-Exponential Decay of Detection Error Probabilities in Long Tandems" *IEEE Transactions on Information Theory*, **54**, 4767–4771.
- UDRY, C. and CONLEY, T. (2001), "Social Learning through Networks: The Adoption of New Agricultural Technologies in Ghana", *American Journal of Agricultural Economics*, **83**, 668–673.
- VIVES, X. (1997), "Learning from Others: A Welfare Analysis", *Games and Economic Behavior*, **20**, 177–200.
- WELCH, I. (1992), "Sequential Sales, Learning and Cascades", *The Journal of Finance*, **47**, 695–732.