PIER Working Paper 10-005

"Non-Bayesian Social Learning"
Second Version

by

Ali Jadbabaie, Alvaro Sandroni, and Alireza Tahbaz-Salehi

Non-Bayesian Social Learning*

Ali Jadbabaie[†], Alvaro Sandroni[‡], and Alireza Tahbaz-Salehi[§]

First version: June 2009
This version: February 2010

Abstract

We develop a dynamic model of opinion formation in social networks. Relevant information is spread throughout the network in such a way that no agent has enough data to learn a payoff-relevant parameter. Individuals engage in communication with their neighbors in order to learn from their experiences. However, instead of incorporating the views of their neighbors in a fully Bayesian manner, agents use a simple updating rule which linearly combines their personal experience and the views of their neighbors (even though the neighbors' views may be quite inaccurate). This non-Bayesian learning rule is motivated by the formidable complexity required to fully implement Bayesian updating in networks. We show that, under mild assumptions, repeated interactions lead agents to successfully aggregate information and to learn the true underlying state of the world. This result holds in spite of the apparent naïvité of agents' updating rule, the agents' need for information from sources (i.e., other agents) the existence of which they may not be aware of, the possibility that the most persuasive agents in the network are precisely those least informed and with worst prior views, and the assumption that no agent can tell whether their own views or their neighbors' views are more accurate.

Keywords: Social networks, learning, information aggregation.
JEL Classification: D83, L14.

---

*We thank seminar participants at MIT, University of Pennsylvania, and University of California, Santa Barbara for useful comments and suggestions. Jadbabaie and Tahbaz-Salehi gratefully acknowledge financial support from the Defense Advanced Research Projects Agency and the Office of Naval Research. Sandroni gratefully acknowledges the financial support of the National Science Foundation (awards SES – 0820472 and SES – 0922404).

[†]Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104. (e-mail: jadbabai@seas.upenn.edu)

[‡]Department of Managerial Economics and Decision Sciences, Kellogg School of Management, Northwestern University, Evanston, IL 60208. (e-mail: sandroni@kellogg.northwestern.edu)

[§]*Corresponding author*, Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA 02139. (e-mail: alirezat@mit.edu)

"Civilization advances by extending the number of important operations which we can perform without thinking about them."

- Alfred North Whitehead

1 Introduction

In everyday life, people form opinions over various economic, political, and social issues – such as how to educate their children or whether to vote for a certain candidate – which do not have an obvious solution. These issues allow for a great variety of opinions, because even if a satisfactory solution exists, it is not easily recognizable. In addition, relevant information for such problems are often not concentrated in any source or body of sufficient knowledge. Instead, the data is dispersed throughout a vast network, where each individual observes only a small fraction, consisting of his/her personal experience. This motivates an individual to engage in communication with others in order to learn from other people's experiences. For example, Hagerstrand (1969) and Rogers (1983) document such a phenomenon in the choice of new agricultural techniques by various farmers, while Kotler (1986) shows the importance of learning from others in the purchase of consumer products.

In many scenarios, however, the information available to an individual is not directly observable by others. At most, each individual only knows the opinions of few individuals (such as colleagues, family members, and maybe a few news organizations), will never know the opinions of everyone in the society, and might not even know the full personal experience of anyone but herself. This limited observability, coupled with the complex interactions of opinions arising from dispersed information over the network, makes it highly impractical for agents to incorporate other people's views in a Bayesian fashion.

The difficulties with Bayesian updating are further intensified if agents do not have complete information about the structure of the social network, which means that they would need to form and update opinions not only on the states of the world, but also on the network topology. This lack of information significantly complicates the required calculations, well beyond individuals' regular computational capabilities, even in networks of moderate size. Nevertheless, the complications with Bayesian learning persist even when individuals have complete information about the network structure, as they still need to perform deductions about the information of every other individual in the network, while only observing the evolution of opinions of their neighbors. Such fully Bayesian agents need to form opinions about who the source of what bit of information is, how the information is spread around the network, and how every agent's opinions affect everyone else's. The necessary information and the computational burden of these calculations are simply prohibitive for adopting Bayesian learning, even in relatively simple networks.[1]

[1] See e.g., Gale and Kariv (2003).

In this paper, we study the evolution of opinions in a society where agents, instead of using Bayesian updates, apply a simple learning rule to incorporate the views of individuals in their social clique. We assume that at every time period, each individual receives a private signal, and observes the opinions (i.e., the beliefs) held by her neighbors at the previous period. The individual updates her belief as a convex combination of the Bayesian posterior belief conditioned on her private signal and the opinions of her neighbors. The weight an individual assigns to the opinion of a neighbor represents the influence (or persuasion power) of that neighbor on her. At the end of the period, agents report their opinions truthfully to their neighbors. The influence that agents exert on one another can be large or small, and may depend on each pair of agents. Moreover, this persuasion power may be independent from the informativeness of their signals. In particular, more persuasive agents may not be better informed or hold more accurate views. In such cases, in initial periods, agents' views will move towards the views of the most persuasive agents and, hence, away from the data generating process.

We analyze the flow of opinions as new observations accumulate. First, we show that agents eventually make correct forecasts, provided that the social network is strongly connected; that is, there exists either a direct or an indirect information path between any two agents. By the means of an example we show that the assumption of strong connectivity cannot be disposed of. Hence, the seemingly naïve updating rule will eventually transform the existing data into a near perfect guide for the future even though the truth is not recognizable, agents do not know if their views are more or less accurate than the views of their neighbors, and the most persuasive agents may have the least accurate views.

We further show that in strongly connected networks, the non-Bayesian learning rule also enables agents to successfully aggregate disperse information. Each agent eventually learns the truth even though no agent and her neighbors, by themselves, may have enough information to infer the underlying parameter. Eventually, each agent learns as if she were completely informed of all observations of *all* agents *and* updated her beliefs according to Bayes' rule. This aggregation of information is achieved while agents avoid the computational complexity involved in Bayesian updating.

Our results also highlight the role of social networks in information processing and aggregation. An agent can learn from individuals whom she is not in direct contact with and may not even be aware of their existence altogether. In other words, the indirect communication path in the social network guarantees that she will eventually incorporate the information revealed to them through their signals correctly into her opinions. For example, assume that one agent receives signals that allow her to determine whether the state is a or b while another agent can distinguish between states b and c. These two agents do not know each other and do not communicate directly with one another, but they are indirectly connected through other agents. Our results establish that all agents' beliefs will eventually be as if they could all distinguish states a and c even though no agent, by herself or together with her neighbors, can do so.

Our basic learning results hold in a wide spectrum of networks and under conditions that are seemingly not conducive to learning. For example, assume that one agent receives uninformative signals and have strong persuasive powers over all agents including the only agent in the network that has informative signals (but who may not know that her signals are more informative than the signals of others). The agent with informative signals cannot directly influence the persuasive agents and only has a small, direct persuasive power over a few other agents. Even so, all agents' views will eventually be as if they were based on informative signals although most agents will have never seen these informative signals and will not know where they come from. These results stand in contrast with the results of Golub and Jackson (2010) who show, in a different model, that "influential" individuals – those who are connected to a large number of people – may make learning impossible.

Another distinctive feature of our results is the absence of absolute continuity of the true measure with respect to *all* prior beliefs, as a requirement for social learning. Even though it has been shown in different contexts that absolute continuity is a strong assumption,[2] it is essential for Bayesian learning and its absence can lead to incorrect forecasts.[3] Moreover, in a network in which all agents are Bayesians, the absence of absolute continuity of the true measure with respect to prior of an agent, not only prevents that agent from forecasting future correctly, but can also affect the learning of others. For instance, consider an agent located at the network's bottleneck, functioning as the only link connecting two components of the network. This agent may prevent agents on each side from obtaining valuable information from the other side. In contrast, individuals in our non-Bayesian model learn the underlying state of the world even if the true measure is not absolutely continuous with respect to all prior beliefs: as long as the social network is strongly connected and the true measure is absolutely continuous with respect to prior belief of a *single* agent, complete learning is achieved. So, even if there is only one agent (agent i) in the network whose prior belief assigns positive measure to the true parameter, no agent in the economy – with the exception of i's neighbors – is aware of her existence, and she is one of the least persuasive of all agents (even among her neighbors), all individuals will still learn the true parameter, although they may never know the ultimate source of learning.

The paper is organized as follows. The next section discusses the related literature. Section 3 contains our model. Our main results are presented in Section 4 and Section 5 concludes. All proofs can be found in the Appendix.

2 Related Literature

There exists a large body of works on learning over social networks, both boundedly and fully rational. The Bayesian social learning literature focuses on formulating the problem as a dynamic game with incomplete information and characterizing its

[2]For example, see Miller and Sanchirico (1997) and Nachbar (1997).

[3]Lehrer and Smorodinsky (1996) and Sandroni (1998) have shown that agents can correctly forecast future events under slightly weaker conditions on the prior beliefs. However, if the state space is finite, as in our model, these conditions coincide with absolute continuity.

equilibria. However, since characterizing the equilibria in complex networks is generally intractable, the literature studies relatively simple and stylized environments. More specifically, rather than considering repeated interactions over the network, it focuses on models where agents interact sequentially and communicate with their neighbors only *once*. Examples include Banerjee (1992), Bikchandani, Hirshleifer, and Welch (1992), Smith and Sørensen (2000), Banerjee and Fudenberg (2004), and more recently, Acemoglu, Dahleh, Lobel, and Ozdaglar (2008). In contrast, in our model, there are repeated social interactions and information exchange among individuals. Moreover, the network is quite flexible and can accommodate general structures.

Our work is also related to the social learning literature that focuses on non-Bayesian learning models, such as Ellison and Fudenberg (1993, 1995) and Bala and Goyal (1998, 2001), in which agents use simple rule-of-thumb methods to update their beliefs. In the same spirit are DeMarzo, Vayanos, and Zwiebel (2003), Golub and Jackson (2010), and Acemoglu, Ozdaglar, and Parandeh-Gheibi (2010), which are based on the opinion formation model of DeGroot (1974). In DeGroot-style models, each individual initially receives *one* signal about the state of the world and the focus is on conditions under which individuals in the connected components of the social network converge to similar opinions. Golub and Jackson further show that if the size of the network grows unboundedly, this asymptotic consensus opinion converges to the true state of the world, provided that there are not overly influential agents in the society.

A feature that distinguishes our model from the works that are based on DeGroot's model, such as Golub and Jackson (2010), is the existence of time dynamics. Whereas in DeGroot's model each agent has only a single observation, the individuals in our model receive information in small bits over time. The existence of time dynamics in our model can potentially lead to learning in finite networks, a feature absent in DeGroot-style models, where learning can only occur when the number of agents increases unboundedly.

As mentioned above, the crucial difference in results between the Golub and Jackson (2010) model and our model is the role played by the social network in successful information aggregation. Golub and Jackson show that presence of "influential" individuals – those who are connected to a large number of people – makes learning impossible. In contrast, in our environment, strong connectivity is the only requirement on the network for successful learning, and neither the network topology nor the influence level of different agents can prevent learning. In fact, social learning is achieved even if the most influential agents (both in terms of their persuasion power and in terms of their location in the network) are the ones with the least informative signals.

Finally, our work is also related to Epstein, Noor, and Sandroni (2008a), who provide choice-theoretic foundations for non-Bayesian opinion formation dynamics of a *single* agent. However, the focus of our analysis is on the process of information aggregation over a network comprising of many agents.

3 The Model

3.1 Agents and Observations

Let Θ denote a finite set of possible states of the world and let $\theta^* \in \Theta$ denote the true underlying state of the world. We consider a set $\mathcal{N} = \{1, 2, \ldots, n\}$ of agents interacting over a social network. Each agent i starts with a prior belief about the true state, denoted by $\mu_{i,0} \in \Delta\Theta$ which is a probability distribution over the set Θ. More generally, we denote the opinion of agent i at time period $t \in \{0, 1, 2 \ldots\}$ by $\mu_{i,t} \in \Delta\Theta$.

Conditional on the state of the world θ, at each time period $t \geq 1$, an observation profile $s_t = (s_t^1, \ldots, s_t^n) \in S_1 \times \cdots \times S_n \equiv S$ is generated by the likelihood function $\ell(s_t|\theta)$. We let $s_t^i \in S_i$ denote the signal privately observed by agent i at period t and S_i denote agent i's signal space, which we assume to be finite. The privately observed signals are independent over time, but might be correlated among agents at the same time period. We assume that $\ell(s|\theta) > 0$ for all $(s, \theta) \in S \times \Theta$ and use $\ell_i(\cdot|\theta)$ to denote the i-th marginal of $\ell(\cdot|\theta)$. We further assume that every agent i knows the conditional likelihood function $\ell_i(\cdot|\theta)$, known as her *signal structure*.

Note that we do not require the observations to be informative about the state. In fact, each agent may face an identification problem, in the sense that she might not be able to distinguish between two states. We say two states are *observationally equivalent from the point of view of an agent* if the conditional distributions of her signals coincide. More specifically, the elements of the set $\bar{\Theta}_i = \{\theta \in \Theta : \ell_i(s^i|\theta) = \ell_i(s^i|\theta^*) \text{ for all } s^i \in S_i\}$ are observationally equivalent to the true state θ^* from the point of view of agent i.

We also impose the mild technical restriction that the signal structure of each agent is such that there exists a signal which is most likely under the true state θ^* than any other state θ, unless θ is observationally equivalent to θ^*. More precisely:

Assumption $(\star)$**.** For any agent i, there exists a signal $\hat{s}^i \in S_i$ and a positive number δ_i such that

$$\frac{\ell_i(\hat{s}^i|\theta)}{\ell_i(\hat{s}^i|\theta^*)} \leq \delta_i < 1 \qquad \forall \theta \notin \bar{\Theta}_i.$$

The above assumption means that signal $\hat{s}^i$ is more likely to realize under the true state θ^* than any other state θ, unless θ is indistinguishable from θ^* by agent i.[4] Notice that signal $\hat{s}^i$ can have an arbitrarily small probability and δ_i can be arbitrarily close to one. Also notice that, Assumption $(\star)$ does not exclude the possibility of existence of other signals (with even higher probabilities than $\hat{s}^i$), that are much more probable under θ than θ^*.

Finally, for a fixed $\theta \in \Theta$, we define a probability triple $(\Omega, \mathcal{F}, \mathbb{P}^\theta)$, where Ω is the space containing sequences of realizations of the signals $s_t \in S$ over time, and $\mathbb{P}^\theta$ is the probability measure induced over sample paths in Ω. In other words, $\mathbb{P}^\theta = \otimes_{t=1}^{\infty} \ell(\cdot|\theta)$. We use $\mathbb{E}^\theta[\cdot]$ to denote the expectation operator associated with measure $\mathbb{P}^\theta$. Define $\mathcal{F}_{i,t}$

[4] We make this assumption for technical reasons. We conjecture that all our results still hold even in its absence.

as the σ-field generated by the past history of agent i's observations up to time period t, and let $\mathcal{F}_t$ be the smallest σ-field containing all $\mathcal{F}_{i,t}$ for $1 \leq i \leq n$.

3.2 Social Structure

When updating their opinions about the true state of the world, agents observe the opinions currently held by their neighbors. We capture the social interaction structure between agents by a directed graph $G = (V, E)$, where each vertex in V corresponds to an agent, and an edge connecting vertex i to vertex j, denoted by the ordered pair $(i, j) \in E$, captures the fact that agent j has access to the opinion held by agent i. Note that because of the way we have defined the social network, opinion of agent i might be accessible to agent j, but not the other way around.

For each agent i, define $\mathcal{N}_i = \{j \in V : (j, i) \in E\}$, called the set of *neighbors* of agent i. The elements of this set are agents whose opinions are available to agent i at each time period. We assume that individuals report their opinions truthfully to their neighbors.

A directed path in $G = (V, E)$ from vertex i to vertex j, is a sequence of vertices starting with i and ending with j such that each vertex is a neighbor of the next vertex in the sequence. We say the social network is *strongly connected*, if there exists a directed path from each vertex to any other vertex.

3.3 Belief Updates

Before the beginning of each period, agents observe the opinions of their neighbors. At the beginning of period t, signal profile $s_t = (s_t^1, \ldots, s_t^n)$ is realized according to the probability law $\ell(\cdot|\theta^*)$, and signal s_t^i is privately observed by agent i. Following the realization of the private signals, each agent computes her Bayesian posterior belief conditional on the signal observed, and then sets her final belief to be a linear combination of the Bayesian posterior and the opinions of her neighbors, observed right before the beginning of the period. At the end of the period, agents report their opinions to their neighbors. More precisely, if we denote the belief that agent i assigns to state $\theta \in \Theta$ at time period t by $\mu_{i,t}(\theta)$, then

$$\mu_{i,t+1} = a_{ii}\,\mathrm{BU}(\mu_{i,t}; s_{t+1}^i) + \sum_{j \in \mathcal{N}_i} a_{ij}\mu_{j,t}, \tag{1}$$

where $a_{ij} \in \mathbb{R}_+$ captures the weight that agent i assigns to the opinion of agent j in her neighborhood, $\mathrm{BU}(\mu_{i,t}; s_{t+1}^i)(\cdot)$ is the Bayesian update of $\mu_{i,t}$ when signal s_{t+1}^i is observed, and a_{ii} is the weight that the agent assigns to her Bayesian posterior conditional on her private signal, which we refer to as the measure of *self-reliance* of agent i.[5] Note

[5]One can generalize this belief update model and assume that agent i's belief update also depends on his own beliefs at the previous time period, $\mu_{i,t}$. Such an assumption is equivalent to adding a prior-bias to the model, as stated in Epstein, Noor, and Sandroni (2008b). Since this added generality does not change the results or the economic intuitions, we assume that agents have no prior bias.

that weights a_{ij} must satisfy $\sum_{j\in\mathcal{N}_i\cup\{i\}} a_{ij} = 1$, in order for the period $t+1$ beliefs to form a well-defined probability distribution.

The beliefs of agent i on Θ at time t induce time $(t+1)$-beliefs on S_i given by

$$m_{i,t}(s^i_{t+1}) = \int_\Theta \ell_i(s^i_{t+1}|\theta)d\mu_{i,t}(\theta), \tag{2}$$

where we refer to probability measure $m_{i,t}(\cdot)$ as agent i's forecasts. Therefore, the law of motion for the beliefs about the parameters can be written as

$$\mu_{i,t+1}(\theta) = a_{ii}\mu_{i,t}(\theta)\frac{\ell_i(s^i_{t+1}|\theta)}{m_{i,t}(s^i_{t+1})} + \sum_{j\in\mathcal{N}_i} a_{ij}\mu_{j,t}(\theta), \tag{3}$$

for all $\theta \in \Theta$. Note that the dynamics of belief update in our model is *local*, in the sense that each individual only uses the beliefs of her immediate neighbors to form her opinions, ignores the structure of the network, and does not make any inferences about the beliefs of other individuals in the society. The above dynamics for opinion formation, compared to the Bayesian case, places a significantly smaller computational burden on the individuals. Moreover, individuals do not need to keep track of the identities of their neighbors and the exact information provided by them. They only need to know the "average belief" held in their neighborhood, given by the term $\sum_{j\in\mathcal{N}_i} a_{ij}\mu_{j,t}(\cdot)$. In the special case that the signals observed by an agent are uninformative (or equivalently, there are no signals) after time $t = 0$, equation (3) reduces to the belief update model of DeGroot (1974), used by Golub and Jackson (2010).

When analyzing the asymptotic behavior of the beliefs, sometimes it is more convenient to use a matrix notation. Define A to be a real $n \times n$ matrix which captures the social interaction of the agents as well as the weight that each agent assigns to her neighbors. More specifically, we let the ij element of the matrix A be a_{ij} when agent j is a neighbor of agent i, and zero otherwise. Thus, equation (3) can be rewritten as

$$\mu_{t+1}(\theta) = A\mu_t(\theta) + \mathrm{diag}\left(a_{11}[\frac{\ell_1(s^1_{t+1}|\theta)}{m_{1,t}(s^1_{t+1})} - 1], \ldots, a_{nn}[\frac{\ell_n(s^n_{t+1}|\theta)}{m_{n,t}(s^n_{t+1})} - 1]\right)\mu_t(\theta) \tag{4}$$

where $\mu_t(\cdot) = [\mu_{1,t}, \ldots, \mu_{n,t}]'(\cdot)$, and diag of a vector is a diagonal matrix which has the entries of the vector as its diagonal. In the special case that A is the identity matrix, our model reduces to the standard Bayesian case, in which the society consists of n Bayesian agents who do not have access to the beliefs of other members of the society, and only observe their own private signals.

4 Social Learning

Given the model described above, we are interested in the evolution of opinions in the network, and whether this evolution can lead to learning in the long run. Learning may either signify learning the true parameter or learning to forecast future outcomes.

These two notions of learning are distinct and might not occur simultaneously. We start this section by specifying what we exactly mean by either type of learning.

Suppose that $\theta^* \in \Theta$ is the true state of the world and thus, the measure $\mathbb{P}^* = \otimes_{t=1}^{\infty} \ell(\cdot|\theta^*)$ is the probability law generating the signals $\{s_t\}_{t=1}^{\infty}$.

Definition 1. The forecasts of agent i are *eventually correct* on a path $\{s_t\}_{t=1}^{\infty}$ if, along that path,

$$m_{i,t}(\cdot) \to \ell_i(\cdot|\theta^*) \quad \text{as} \quad t \to \infty.$$

This notion of learning, called *weak merging* of opinions, captures the ability of agents to correctly forecast events in near future (see Kalai and Lehrer (1994)). It is well-known, that repeated applications of the Bayes' rule leads to eventually correct forecasts with probability 1 under the truth, given suitable conditions, the key condition being absolute continuity of the true measure with respect to initial beliefs.[6] In the presence of absolute continuity, the mere repetition of Bayes' rule eventually transforms the historical record into a near perfect guide for the future. However, predicting events in near future accurately is not the same as learning the underlying state of the world. In fact, depending on the signal structure of each agent, there might be an "identification problem" which can potentially prevent the agent from learning the true parameter θ^*. The other type of learning that we are concerned with, precisely captures this notion:

Definition 2. Agent $i \in \mathcal{N}$ *asymptotically learns* the true parameter θ^* on a path $\{s_t\}_{t=1}^{\infty}$ if, along that path,

$$\mu_{i,t}(\theta^*) \to 1 \quad \text{as} \quad t \to \infty.$$

Asymptotic learning occurs when agent assigns probability one to the true parameter. As mentioned earlier, making correct forecasts about future events does not necessarily guarantee learning the true state. In general, the converse is not true either. However, it is straightforward to show that in the absence of time correlations, as in our model, asymptotically learning θ^* implies eventually correct forecasts.[7]

4.1 Correct Forecasts in Strongly Connected Societies

We now turn to the main question of this paper: under what circumstances does learning occur over the social network?

Our first result shows that under very mild assumptions, in spite of local interactions, limited observability, and the non-Bayesian belief update, agents will eventually make correct forecasts. The proof is provided in Appendix B.

[6] Lehrer and Smorodinsky (1996) show that an assumption weaker than absolute continuity, known as *accommodation*, is sufficient for weak merging of the opinion.

[7] See Lehrer and Smorodinsky (1996), for an example of the case that learning the true parameter does not guarantee weak merging.

Proposition 1. *Suppose that the social network is strongly connected, all agents have strictly positive self-reliances, and there exists an agent with positive prior belief on the true parameter θ^*. Then, the forecasts of all agents are eventually correct with $\mathbb{P}^*$– probability one.*

This proposition states that, when agents use non-Bayesian update (3) to form and update their opinions, they will eventually make accurate predictions. Note that as long as the social network remains strongly connected, neither the topology of the network nor the influence levels of different individuals prevent agents from making correct forecasts.

Another substantive feature of Proposition 1 is the absence of absolute continuity of the true measure with respect to the prior beliefs of all agents in the society, as a requirement for eventually correct forecasts: as long as some agent assigns a positive prior belief to the true parameter θ^*, all agents will eventually make accurate predictions. In fact, all forecasts are eventually correct even if the only agent for whom absolute continuity holds is located at the fringe of the society, has very small persuasive power over her neighbors, and almost everyone in the network is unaware of her existence.

Besides the existence of an agent with a positive prior belief on the true state, the above proposition requires the existence of positive self-reliances to guarantee correct forecasts. This requirement is quite intuitive: it prohibits agents from completely discarding information provided to them through their observations. Clearly, if all agents discard their private signals, no new information is incorporated into their opinions, and (3) simply turns into a diffusion of prior beliefs.

The final requirement for accurate predictions is strong connectivity of the social network. The following example illustrates that this assumption cannot be disposed of.

Example 1. Consider a society consisting of two agents, $\mathcal{N} = \{1, 2\}$, and assume that $\Theta = \{\theta_1, \theta_2\}$ with the true state being $\theta^* = \theta_1$. Both agents have non-degenerate prior beliefs over Θ. Assume that signals observed by the agents are conditionally independent, and belong to the set $S_1 = S_2 = \{H, T\}$. We further assume that Agent 2's signals are non-informative, while Agent 1's observations are perfectly informative about the state; that is, $\ell_1(H|\theta_1) = \ell_1(T|\theta_2) = 1$, and $\ell_2(s|\theta_1) = \ell_2(s|\theta_2)$ for $s \in \{H, T\}$. As for the social structure, we assume that agent 1 has access to the opinion of Agent 2, while Agent 2 cannot observe the opinion of Agent 1. Clearly, the social network is not strongly connected. We let the social interaction matrix be

$$A = \begin{bmatrix} 1-\alpha & \alpha \\ 0 & 1 \end{bmatrix},$$

where $\alpha \in (0, 1)$ is the weight that Agent 1 assigns to the opinion of Agent 2, when updating her beliefs using equation (3). Since the private signals observed by the latter are non-informative, her beliefs, at all times, remain equal to her prior. Clearly, she makes correct forecasts at all times. Agent 1's forecasts, on the other hand, will always remain incorrect. Notice that since her signals are perfectly informative, Agent 1's forecasts are eventually correct if and only if she eventually assigns probability 1 to the true state, θ_1. However, the belief she assigns to θ_2 follows the law of motion

$$\mu_{1,t+1}(\theta_2) = (1-\alpha)\mu_{1,t}(\theta_2)\frac{\ell_1(s^1_{t+1}|\theta_2)}{m_{1,t}(s^1_{t+1})} + \alpha\mu_{2,t}(\theta_2)$$

which cannot converge to zero, as $\mu_{2,t}(\theta_2) = \mu_{2,0}(\theta_2)$ is strictly positive.

The intuition for failure of learning in this example is simple. Notice that given the same observations, the two agents make different interpretations about the state, even if they have equal prior beliefs. Moreover, Agent 1 follows the beliefs of the less informed Agent 2 but is unable of influencing her back. This one-way persuasion and non-identical interpretations of signals (due to non-identical signal structures) result in incorrect forecasts on the part of Agent 1.

4.2 Social Agreement

The key implication of Proposition 1 is that as long as the social network is strongly connected, the forecasts of all agents will eventually be correct. Our next result establishes that not only agents make accurate predictions about their private observations, but will also hold asymptotically equal beliefs about the underlying state.

Proposition 2. *Suppose that the social network is strongly connected, all agents have strictly positive self-reliances, and there exists an agent with positive prior belief on the true parameter θ^*. Moreover, suppose that Assumption $(\star)$ holds. Then, the beliefs of all agents converge with $\mathbb{P}^*$-probability 1. Moreover, all agents have asymptotically equal beliefs $\mathbb{P}^*$-almost surely. That is, with $\mathbb{P}^*$-probability 1, $\lim_{t\to\infty}\mu_{i,t}(\theta)$ exists for all $i \in \mathcal{N}$ and all $\theta \in \Theta$, and its value does not depend on i.*

The above proposition states that when the social network is strongly connected, the opinions do not fluctuate forever and reach some limit asymptotically. Moreover, social interactions among agents and the opinion formation process described in equation (3) transform initial heterogeneity and diversity of opinions into homogeneity and agreement. Note that these results are achieved regardless of how the social network is structured (aside from strong connectivity), the influence level of individuals on one another, and their signal structures. On the other hand, Proposition 2 also highlights the role played by the social network in generating asymptotic agreement. Clearly, due to nonidentical signal structures and potential identification problems, agents would not have reached the same opinions had they been updating their beliefs in isolation.

Proposition 2 relies heavily on Proposition 1. The intuition is as follows: an individual's forecasts are eventually correct on a sample path only if her opinions converge asymptotically on that path. Once the convergence of beliefs to some limit is established, the fact that each agent's belief lies in the convex hull of her neighbors' opinions implies that all opinions must coincide asymptotically.

4.3 Social Learning

Proposition 2 indicates that in strongly connected social networks, all individuals will eventually hold similar opinions. On the other hand, given the fact that their forecasts are eventually correct (Proposition 1), their asymptotic opinions cannot be arbitrary. The following theorem, which is our main result, establishes that strong connectivity of the social network not only leads to agreement, but also guarantees information aggregation over the network, in the sense that all individuals learn the true state.

Theorem 3. *Suppose that:*

(a) The social network is strongly connected.

(b) All agents have strictly positive self-reliances.

(c) There exists an agent with positive prior belief on the true parameter θ^*.

(d) Assumption $(\star)$ holds.

(e) There is no state $\theta \neq \theta^*$ that is observationally equivalent to θ^* from the point of view of all agents in the network.

Then, all agents in the social network learn the true state of the world $\mathbb{P}^*$– almost surely; that is, $\mu_{i,t}(\theta^*) \longrightarrow 1$ with $\mathbb{P}^*$– probability one for all $i \in \mathcal{N}$, as $t \to \infty$.

The above theorem states that under fairly mild assumptions on the social network's topology and the individuals' signal structures, all agents will eventually learn the true underlying state of the world. Notice that agents only interact with their neighbors and perform no deductions beyond their immediate neighbors. Yet, the non-Bayesian updating rule eventually enables them to obtain relevant information from others, without exactly knowing where it comes from. In fact, they can be completely oblivious to important features of the social network – such as the number of individuals in the society, the topology of the network, other people's signal structures, the existence of some agent who considers the truth plausible, or the influence level of any agent in the network – and still learn the true parameter. Moreover, all these results are achieved with a significantly smaller computational burden than what is required for Bayesian learning.

The other significant feature of our result in Theorem 3 is the fact that neither network's topology, the signal structures, nor the influence levels of different agents prevent learning. For instance, even if the agents with the least informative signals are the most persuasive ones and are located at the bottlenecks of the network, everyone will eventually learn the true state. It is important to emphasize once again that social learning is achieved despite the fact that the truth is not recognizable to any individual, and she would not have learned it by herself in isolation.

As mentioned earlier, the assumptions of Theorem 3 are quite mild. The strong connectivity assumption simply creates the possibility of information flow between any pair of agents in the social network. The assumption on positive self-reliances guarantees that agents do not discard the information provided to them through their private observations. The third assumption states that it is sufficient to have only one agent in the society who assigns a positive prior belief to the truth, even if that agent is at the fringe of the society, has a very small influence on her neighbors, and almost no one is aware of her existence. Hence, the ultimate source of learning may remain unknown to almost everyone in the society. Clearly, if the prior beliefs of all agents assigned to the truth is equal to zero, then they will never learn. Also note that, as mentioned earlier, assumption (c) is weaker than what is required for learning in a network in which all agents are Bayesians.

Assumption ($\star$) states that for any agent i, there exists a signal which is most likely under the true state θ^* than any other state θ, unless θ is observationally equivalent to θ^* from the point of view of agent i.[8] Finally, the last assumption indicates that the collection of observations of all agents is informative enough about the true state; that is, $\bar{\Theta}_1 \cap \cdots \cap \bar{\Theta}_n = \{\theta^*\}$.[9] This assumption guarantees that it is possible to learn the truth if one has access to the observations of all agents. Notice that in the absence of assumption (e) even highly sophisticated, Bayesian agents with access to all relevant information (such as the topology of the network and the signal structures), would not be able to completely learn the state, due to the presence of an identification problem.

The next examples show the power of Theorem 3.

Example 2. Consider the collection of agents $\mathcal{N} = \{1, 2, \ldots, 7\}$ who are located in a social network as depicted in Figure 1: at every time period, agent $i \leq 6$ can observe the opinion of agent $i+1$ and agent 7 has access to the opinion held by agent 1. Clearly, this is a strongly connected social network.

Assume that the set of possible states of the world is given by $\Theta = \{\theta^*, \theta_1, \theta_2, \ldots, \theta_7\}$, where θ^* is the true underlying state of the world. We also assume that the signals observed by the agents belong to the set $S_i = \{H, T\}$ for all i, are conditionally independent, and have conditional distributions given by

$$\ell_i(H|\theta) = \begin{cases} \frac{i}{i+1} & \text{if } \theta = \theta_i \\ \\ \frac{1}{(i+1)^2} & \text{otherwise} \end{cases}$$

for all $i \in \mathcal{N}$.

Notice that Assumption ($\star$) is satisfied. The signal structures are such that each agent suffers from some identification problem; i.e., the information in the observations of any agent is not sufficient for learning the true state of the world in isolation. More precisely, $\bar{\Theta}_i = \Theta / \{\theta_i\}$ for all i, which means that from the point of view of agent i, all states except for θ_i are observationally equivalent to the true state θ^*. Nevertheless, for any given state $\theta \neq \theta^*$, there exists an agent whose signals are informative enough to

[8]We conjecture that all our results hold even in the absence of Assumption ($\star$).

[9]This is a stronger restriction than requiring $\ell(\cdot|\theta) \neq \ell(\cdot|\theta^*)$ for all $\theta \neq \theta^*$. For more on this, see Example 4 in Appendix A.

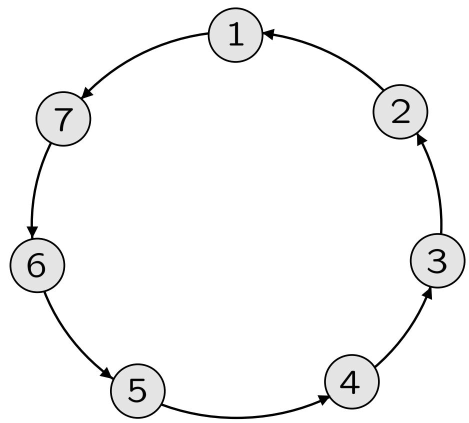

Figure 1: The figure illustrates a strongly connected social network of 7 agents, which is of the form of a directed cycle.

distinguish the two; that is, $\cap_{i=1}^{7} \bar{\Theta}_i = \{\theta^*\}$. Therefore, Theorem 3 implies that as long as one agent assigns a positive prior belief on the true state θ^* and all agents have strictly positive self-reliances when applying (3), then $\mu_{i,t}(\theta^*) \to 1$, as $t \to \infty$ for all agents i, with $\mathbb{P}^*$–probability one. In other words, all agents will asymptotically learn the true underlying state of the world. Clearly, if agents discard the information provided to them by their neighbors, they have no means of learning the true state.

Example 3. Consider a collection of agents who are connected to one another according to the social network depicted in Figure 2. The values on the edges depict the persuasion power of different agents on each other, where $\epsilon > 0$ is some arbitrarily small number. As the figure suggests, Agent M is the most influential agent in the network, both in terms of persuasion power and connectivity: she can highly influence almost everyone in the society, while being only marginally influenced by the public opinion herself. One can think of M representing a far reaching news media.

Even though highly influential, agent M is not well-informed about the true underlying state of the world $\theta^* \in \Theta$. More specifically, we assume that her signals are completely non-informative and that she does not consider θ^* a possible candidate for the truth, i.e., she assigns a zero prior belief to that state. In fact, we assume that agent A – who is neither highly persuasive nor can broadcast her opinions beyond her immediate neighbors – is the only agent in the society who assigns some positive prior belief to θ^*. In addition, we assume that agent S is the only agent in the social network with access to informative signals, enabling her to distinguish different states from one another.

Since the social network is strongly connected, as long as Assumption ($\star$) is satisfied for the signal structure of Agent S, Theorem 3 implies that all agents will asymptotically learn the truth. This is despite the fact that in initial periods, due to the high persuasion power of agent M and her far reach, the views of all agents (including agents A and S) will move towards the initial views of agent M. However, such effects are only transient and will not last forever. As time progresses, due to the possibility of reciprocal

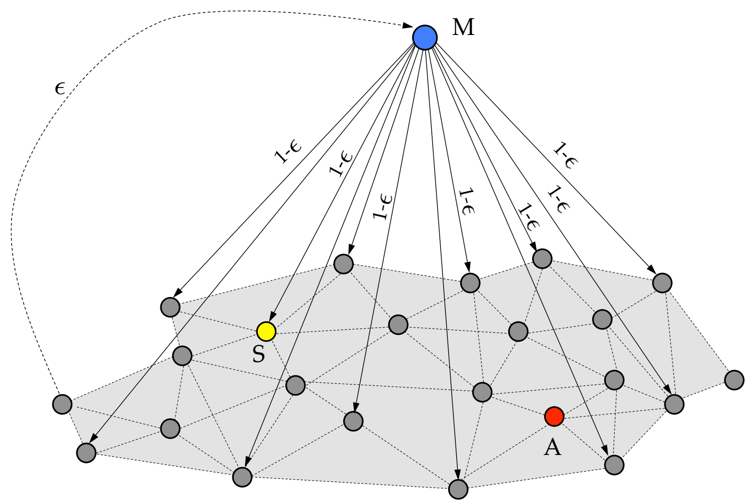

Figure 2: "The power of truth": Agents A and S, with absolutely continuous priors and informative signals respectively, eventually lead every other agent to learn the truth, even though agent M (with access to neither good priors nor informative signals) is much more persuasive than any other agent.

persuasion in the network (although highly asymmetric), the views of agents A and S about the true parameter are spread throughout the network. Since such views are consistent with the personal experience of *all* agents, they are eventually consolidated all across the social network. Thus, in the tension between high persuasion power and global reach of M versus the grain of truth of the beliefs of the "obscure" agents A and S, eventually, agents A and S prevail. These results hold even though at no point in time the truth is recognizable to any of the agents, including agents A and S themselves.

5 Conclusions

In this paper, we study a model of dynamic opinion formation in social networks. Agents fail to incorporate the views of their neighbors in a fully Bayesian manner, and instead, use a local updating rule. More specifically, at every time period, the belief of each individual is a convex combination of her Bayesian posterior belief and her neighbors' expressed beliefs. Agents eventually make correct forecasts, as long as the social network is strongly connected. In addition, agents successfully aggregate all information over the entire social network: they eventually learn the true underlying state of the world as if they were completely informed of *all* signals *and* updated their beliefs according to Bayes' rule. Furthermore, in contrast to standard Bayesian learning results, absolute continuity of the true measure with respect to all prior beliefs is not a necessary condition for social learning. As long as some individual places strictly positive prior probability on the true parameter, social learning is achieved.

The aggregation of information is achieved even if individuals are unaware of important features of the environment. In particular, agents do not need to have any information (or form beliefs) about the structure of the social network nor the views or characteristics of most agents, as they only update their opinions *locally* and do not make any deductions beyond their immediate neighbors. Moreover, the individuals do not need to know the signal structure of any other agent in the network, besides their own. The simplicity of the local update rule guarantees that individuals eventually achieve full learning, while at the same time, avoiding highly complex computations that are essential for full Bayesian learning over the network.

Appendix A

A.1 An Example of Incomplete Learning

Example 4. Consider a strongly connected social network consisting of two individuals $\mathcal{N} = \{1, 2\}$. Assume that $\Theta = \{\theta_1, \theta_2\}$, and $S_1 = S_2 = \{H, T\}$. Also assume that the distribution function describing the random private observations of the agents conditional on the underlying state of the world is given by the following tables:

$\ell(s_1 s_2|\theta_1)$:

	H	T
H	1/2	0
T	0	1/2

$\ell(s_1 s_2|\theta_2)$:

	H	T
H	0	1/2
T	1/2	0

In other words, under state θ_1, the private observations of the two agents are perfectly correlated, while when the underlying state of the world is θ_2, their observations are perfectly negatively correlated. Notice that even though the joint distributions of the signals generated by θ_1 and θ_2 are different, we have $\ell_i(H|\theta_1) = \ell_i(H|\theta_2) = \frac{1}{2}$ for $i = 1, 2$; i.e., the local signal structure of each agent is the same under either state. As a result, despite the fact that agents will eventually agree on their opinions and make correct forecasts, learning as defined in Definition 2 does not occur, because θ_1 and θ_2 are observationally equivalent from the point of view of both agents – which amounts to failure of assumption (e) of Theorem 3. However, notice that, due to this identification problem, even highly sophisticated, Bayesian agents would not be able to completely learn the state either.

Appendix B: Proofs

B.1 Two Auxiliary Lemmas

Before presenting the proofs of the results in the paper, we state and prove two lemmas, both of which are consequences of the martingale convergence theorem.

Lemma 1. *Let A denote the matrix of social interactions. The sequence $\sum_{i=1}^{n} v_i \mu_{i,t}(\theta^*)$ converges $\mathbb{P}^*$–almost surely as $t \to \infty$, where v is any non-negative left eigenvector of A corresponding to its unit eigenvalue.*

Proof: First, note that since A is stochastic,[10] it always has at least one eigenvalue equal to 1. Moreover, there exists a non-negative left eigenvector corresponding to this eigenvalue.[11] We denote such a vector by v.

Evaluate equation (4) at the true parameter θ^* and multiply both sides by v' from left

$$v'\mu_{t+1}(\theta^*) = v'A\mu_t(\theta^*) + \sum_{i=1}^{n} v_i \mu_{i,t}(\theta^*) a_{ii} [\frac{\ell_i(s_{t+1}^i|\theta^*)}{m_{i,t}(s_{t+1}^i)} - 1].$$

Thus,

$$\mathbb{E}^* \left[\sum_{i=1}^{n} v_i \mu_{i,t+1}(\theta^*) | \mathcal{F}_t \right] = \sum_{i=1}^{n} v_i \mu_{i,t}(\theta^*) + \sum_{i=1}^{n} v_i a_{ii} \mu_{i,t}(\theta^*) \mathbb{E}^* \left[\frac{\ell_i(s_{t+1}^i|\theta^*)}{m_{i,t}(s_{t+1}^i)} - 1 | \mathcal{F}_t \right], \quad (5)$$

where $\mathbb{E}^*$ denotes the expectation operator associated with measure $\mathbb{P}^*$. Since $f(x) = 1/x$ is a convex function, Jensen's inequality implies that

$$\mathbb{E}^* \left[\frac{\ell_i(s_{t+1}^i|\theta^*)}{m_{i,t}(s_{t+1}^i)} | \mathcal{F}_t \right] \geq \left(\mathbb{E}^* \left[\frac{m_{i,t}(s_{t+1}^i)}{\ell_i(s_{t+1}^i|\theta^*)} | \mathcal{F}_t \right] \right)^{-1} = 1,$$

and therefore,

$$\mathbb{E}^* \left[\sum_{i=1}^{n} v_i \mu_{i,t+1}(\theta^*) | \mathcal{F}_t \right] \geq \sum_{i=1}^{n} v_i \mu_{i,t}(\theta^*).$$

The last inequality is due to the fact that v is element-wise non-negative. As a result, $\sum_{i=1}^{n} v_i \mu_{i,t}(\theta^*)$ is a submartingale with respect to the filtration $\mathcal{F}_t$, which is also bounded above by $\|v\|_1$. Hence, it converges $\mathbb{P}^*$-almost surely.

Lemma 2. *Suppose that there exists an agent i such that $\mu_{i,0}(\theta^*) > 0$. Also suppose that the social network is strongly connected. Then, the sequence $\sum_{i=1}^{n} v_i \log \mu_{i,t}(\theta^*)$ converges $\mathbb{P}^*$-almost surely as $t \to \infty$, where v is any non-negative left eigenvector of A corresponding to its unit eigenvalue.*

[10] A matrix is said to be stochastic if it is entry-wise non-negative and all its row sums are equal to one.

[11] This is a consequence of the Perron-Frobenius theorem. For more on the properties of non-negative and stochastic matrices, see Berman and Plemmons (1979).

Proof: Similar to the proof of the previous lemma, we show that $\sum_{i=1}^n v_i \log \mu_{i,t}(\theta^*)$ is a bounded submartingale and invoke the martingale convergence theorem to obtain almost sure convergence.

By evaluating the law of motion at θ^*, taking $\log$ from both sides, and using the fact that the row sums of A are equal to one, we obtain

$$\log \mu_{i,t+1}(\theta^*) \geq a_{ii} \log \mu_{i,t}(\theta^*) + a_{ii} \log \left(\frac{\ell_i(s_{t+1}^i|\theta^*)}{m_{i,t}(s_{t+1}^i)} \right) + \sum_{j \in \mathcal{N}_i} a_{ij} \log \mu_{j,t}(\theta^*),$$

where we have used the concavity of the logarithm function. Note that since the social network is strongly connected, the existence of one agent with a positive prior on θ^* guarantees that after at most n periods all agents assign a strictly positive probability to the true parameter, which means that $\log \mu_{i,t}(\theta^*)$ is well-defined for large enough t and all i.

Our next step is to show that $\mathbb{E}^* \left[\log \frac{\ell_i(s_{t+1}^i|\theta^*)}{m_{i,t}(s_{t+1}^i)} |\mathcal{F}_t \right] \geq 0$. To obtain this,

$$\begin{aligned} \mathbb{E}^* \left[\log \frac{\ell_i(s_{t+1}^i|\theta^*)}{m_{i,t}(s_{t+1}^i)} |\mathcal{F}_t \right] &= -\mathbb{E}^* \left[\log \frac{m_{i,t}(s_{t+1}^i)}{\ell_i(s_{t+1}^i|\theta^*)} |\mathcal{F}_t \right] \\ &\geq -\log \left(\mathbb{E}^* \left[\frac{m_{i,t}(s_{t+1}^i)}{\ell_i(s_{t+1}^i|\theta^*)} |\mathcal{F}_t \right] \right) \\ &= 0. \end{aligned}$$

Thus,

$$\mathbb{E}^* \left[\log \mu_{i,t+1}(\theta^*) | \mathcal{F}_t \right] \geq a_{ii} \log \mu_{i,t}(\theta^*) + \sum_{j \in \mathcal{N}_i} a_{ij} \log \mu_{j,t}(\theta^*),$$

which can be rewritten in matrix form as $\mathbb{E}^* \left[\log \mu_{t+1}(\theta^*) | \mathcal{F}_t \right] \geq A \log \mu_t(\theta^*)$, where by the logarithm of a vector, we mean its entry-wise logarithm. Multiplying both sides by A's non-negative left eigenvector v' leads to

$$\mathbb{E}^* \left[\sum_{i=1}^n v_i \log \mu_{i,t+1}(\theta^*) | \mathcal{F}_t \right] \geq \sum_{i=1}^n v_i \log \mu_{i,t}(\theta^*).$$

Thus, the non-positive sequence $\sum_{i=1}^n v_i \log \mu_{i,t}(\theta^*)$ is a submartingale with respect to filtration $\mathcal{F}_t$, and therefore, converges with $\mathbb{P}^*$-probability one. ∎

With these lemmas in hand, we can prove Proposition 1.

B.2 Proof of Proposition 1

First, note that since the social network is strongly connected, the social interaction matrix A is an irreducible stochastic matrix, and therefore its left eigenvector corresponding to the unit eigenvalue is strictly positive.[12]

[12] An $n \times n$ matrix A is said to be *reducible*, if for some permutation matrix P, the matrix $P'AP$ is block upper triangular. If a square matrix is not reducible, it is said to be *irreducible*. For more on this, see e.g., Berman and Plemmons (1979).

According to Lemma 1, $\sum_{i=1}^n v_i \mu_{i,t}(\theta^*)$ converges with $\mathbb{P}^*$-probability one, where v is the positive left eigenvector of A corresponding to its unit eigenvalue. Therefore, equation (5) implies that

$$\sum_{i=1}^n v_i a_{ii} \mu_{i,t}(\theta^*) \left(\mathbb{E}^* \left[\frac{\ell_i(s^i_{t+1}|\theta^*)}{m_{i,t}(s^i_{t+1})} |\mathcal{F}_t \right] - 1 \right) \longrightarrow 0 \quad \mathbb{P}^* - \text{a.s.}$$

Since the term $v_i a_{ii} \mu_{i,t}(\theta^*) \mathbb{E}^* \left[\ell_i(s^i_{t+1}|\theta^*)/m_{i,t}(s^i_{t+1}) - 1 |\mathcal{F}_t \right]$ is non-negative for all i, each such term converges to zero with $\mathbb{P}^*$-probability one. Moreover, the assumptions that all diagonal entries of A are strictly positive and that of its irreducibility (which means that v is entry-wise positive) lead to

$$\mu_{i,t}(\theta^*) \left(\mathbb{E}^* \left[\frac{\ell_i(s^i_{t+1}|\theta^*)}{m_{i,t}(s^i_{t+1})} |\mathcal{F}_t \right] - 1 \right) \longrightarrow 0 \quad \text{for all } i \quad \mathbb{P}^* - \text{a.s.} \tag{6}$$

Furthermore, Lemma 2 guarantees that $\sum_{i=1}^n v_i \log \mu_{i,t}(\theta^*)$ converges almost surely, implying that $\mu_{i,t}(\theta^*)$ is uniformly bounded away from zero for all i with probability one. Note that, once again we are using the fact that v is a strictly positive vector. Hence, $\mathbb{E}^* \left[\frac{\ell_i(s^i_{t+1}|\theta^*)}{m_{i,t}(s^i_{t+1})} |\mathcal{F}_t \right] \to 1$ almost surely. Thus,

$$\begin{aligned}
\mathbb{E}^* \left[\frac{\ell_i(s^i_{t+1}|\theta^*)}{m_{i,t}(s^i_{t+1})} |\mathcal{F}_t \right] - 1 &= \sum_{s \in S_i} \ell_i(s|\theta^*) \left(\frac{\ell_i(s|\theta^*)}{m_{i,t}(s)} - 1 \right) \\
&= \sum_{s \in S_i} \left(\ell_i(s|\theta^*) \frac{\ell_i(s|\theta^*) - m_{i,t}(s)}{m_{i,t}(s)} + m_{i,t}(s) - \ell_i(s|\theta^*) \right) \\
&= \sum_{s \in S_i} \frac{\left[\ell_i(s|\theta^*) - m_{i,t}(s)\right]^2}{m_{i,t}(s)} \longrightarrow 0 \qquad \mathbb{P}^* - \text{a.s.},
\end{aligned}$$

where the second equality is due to the fact that both $\ell_i(\cdot|\theta^*)$ and $m_{i,t}(\cdot)$ are measures on S_i, and therefore, $\sum_{s \in S_i} \ell_i(s|\theta^*) = \sum_{s \in S_i} m_{i,t}(s) = 1$.

In the last expression, the term in the braces and the denominator are always non-negative and therefore,

$$m_{i,t}(s) \longrightarrow \ell_i(s|\theta^*) \qquad \mathbb{P}^* - \text{a.s.}$$

for all $s \in S_i$ and all $i \in \mathcal{N}$. ■

B.3 Proof of Proposition 2

We prove this proposition in two steps. First, we focus on the set of states that are observationally equivalent to the true state from the point of view of all agents in the society. On the second part, we focus on the compliment set.

Suppose that a state $\theta \in \Theta$ is observationally equivalent to θ^* from the point of view of all individuals; that is $\theta \in \bar{\Theta}_1 \cap \cdots \cap \bar{\Theta}_n$.[13] For any such state, Proposition 1 guarantees

[13]Recall that $\bar{\Theta}_i \subseteq \Theta$ is the defined as $\bar{\Theta}_i = \{\theta \in \Theta : \ell_i(s^i|\theta) = \ell_i(s^i|\theta^*) \text{ for all} s^i \in S_i\}$.

that $m_{i,t}(\cdot) \to \ell_i(\cdot|\theta^*) = \ell_i(\cdot|\theta)$. Therefore, by equation (4), $\mu_{t+1}(\theta) - A\mu_t(\theta) \to 0$ with $\mathbb{P}^*$-probability one. That is, on almost all sample paths and for any $\epsilon > 0$, there exists a large enough time T such that for all $t \geq T$,

$$|\mu_{i,t+1}(\theta) - \sum_{k=1}^{n} a_{ik}\mu_{k,t}(\theta)| < \frac{\epsilon}{2} \qquad \forall i \in \mathcal{N} \tag{7}$$

Therefore, given any two agents i and j,

$$|\,(\mu_{i,t+1}(\theta) - \mu_{j,t+1}(\theta)) - \sum_{k=1}^{n} \mu_{k,t}(\theta)(a_{ik} - a_{jk})| < \epsilon, \tag{8}$$

and hence,

$$|\mu_{i,t+1}(\theta) - \mu_{j,t+1}(\theta)| < \epsilon + |\sum_{k=1}^{n} \mu_{k,t}(\theta)(a_{ik} - a_{jk})|. \tag{9}$$

Since A is a stochastic matrix, $\sum_{k=1}^{n}(a_{ik} - a_{jk}) = 0$. Therefore, we can use Paz's inequality to find an upper bound for the right hand side of the above inequality:[14]

$$|\mu_{i,t+1}(\theta) - \mu_{j,t+1}(\theta)| < \epsilon + \frac{1}{2}\max_{p,q}|\mu_{p,t}(\theta) - \mu_{q,t}(\theta)| \sum_{k=1}^{n} |a_{ik} - a_{jk}|. \tag{10}$$

Thus,

$$\max_{i,j}|\mu_{i,t+1}(\theta) - \mu_{j,t+1}(\theta)| < \epsilon + \tau(A)\max_{i,j}|\mu_{i,t}(\theta) - \mu_{j,t}(\theta)|, \tag{11}$$

where $\tau(A) = \frac{1}{2}\max_{i,j}\sum_{k=1}^{n}|a_{ik} - a_{jk}|$ is known as the coefficient of ergodicity of matrix A. It is an easy exercise to show that the coefficient of ergodicity of any stochastic matrix lies in the interval $[0, 1]$. This along with the fact that $\epsilon > 0$ is arbitrary imply that $\max_{i,j}|\mu_{i,t}(\theta) - \mu_{j,t}(\theta)|$ is a non-increasing sequence and hence, converges. We now prove that the limit is in fact zero.

Since the social network is strongly connected, A is an irreducible matrix. Therefore, as shown by Seneta (1981), there exists a positive integer r such that $\tau(A^r) < 1$. Moreover, eventually correct forecasts imply that $\mu_{t+r}(\theta) - A^r\mu_t(\theta) \to 0$, as $t \to \infty$. Following the same steps in as equations (7)-(11) for matrix A^r leads to

$$\max_{i,j}|\mu_{i,t+r}(\theta) - \mu_{j,t+r}(\theta)| < \epsilon + \tau(A^r)\max_{i,j}|\mu_{i,t}(\theta) - \mu_{j,t}(\theta)|$$

for an arbitrary $\epsilon > 0$ and a large enough t. Thus, for any positive integer p, we have

$$\begin{aligned}
\max_{i,j}|\mu_{i,pr}(\theta) - \mu_{j,pr}(\theta)| &< \sum_{k=0}^{p-1}[\tau(A^r)]^k\epsilon + [\tau(A^r)]^p\max_{i,j}|\mu_{i,0}(\theta) - \mu_{j,0}(\theta)| \\
&= \frac{1 - [\tau(A^r)]^p}{1 - \tau(A^r)}\epsilon + [\tau(A^r)]^p\max_{i,j}|\mu_{i,0}(\theta) - \mu_{j,0}(\theta)| \\
&\to \frac{\epsilon}{1 - \tau(A^r)} \quad \text{as} \quad p \to \infty.
\end{aligned}$$

[14] Paz's inequality states that if d is a vector with an entry-wise sum of zero, then for any arbitrary vector z of the same size, $|d'z| \leq \frac{1}{2}\|d\|_1 \max_{i,j}|z_i - z_j|$. This inequality can be found in the book of Paz (1971) or Kirkland, Neumann, and Shader (1998).

Since $\epsilon > 0$ is arbitrary, the right hand side can be made arbitrarily small, and as a result, $\max_{i,j} |\mu_{i,t}(\theta) - \mu_{j,t}(\theta)|$ must converge to zero; i.e., in any strongly connected network, $\mu_{i,t}(\theta) - \mu_{j,t}(\theta) \to 0$ for all $i, j \in \mathcal{N}$, and all $\theta \in \bar{\Theta}_1 \cap \cdots \cap \bar{\Theta}_n$.

In order to complete the proof of the first part, all we need to show is that $\lim_{t\to\infty} \mu_{i,t}(\theta)$ exists. Since $|\mu_{i,t}(\theta) - \mu_{j,t}(\theta)| \to 0$ as $t \to \infty$, for any $\delta > 0$, there exists a large enough T such that for any $t \geq T$, we have $-\delta < \mu_{i,t}(\theta) - \mu_{j,t}(\theta) < \delta$ uniformly for all i and j. Thus, for any finite positive integer p,

$$\mu_{j,t}(\theta) - \delta < \sum_{i=1}^{n} a_{ji}^{(p)} \mu_{i,t}(\theta) < \mu_{j,t}(\theta) + \delta,$$

where $a_{ij}^{(p)}$ is the (i, j)-entry of stochastic matrix A^p. Note that the term $\sum_{i=1}^{n} a_{ji}^{(p)} \mu_{i,t}(\theta)$ can be made arbitrarily close to $\mu_{j,t+p}(\theta)$ for large enough t, implying that $-\delta < \mu_{j,t+p}(\theta) - \mu_{j,t}(\theta) < \delta$. Therefore, $\{\mu_{j,t}(\theta)\}_{t=1}^{\infty}$ is a Cauchy sequence for all j and hence, converges.

We now consider the case that $\theta \notin \bar{\Theta}_i$ for some agent i. We prove that the belief assigned to such state θ by agent i converges to zero. To show this, we pick a sample path over which $m_{i,t}(\cdot) \to \ell_i(\cdot|\theta^*)$ as $t \to \infty$. Note that Proposition 1 guarantees that such paths have $\mathbb{P}^*$-measure one. On such a sample path, for any $\epsilon > 0$, there exists a large enough time T, such that for all $t \geq T$,

$$\left| \sum_{\theta \in \Theta} \mu_{i,t}(\theta) \frac{\ell_i(s^i|\theta)}{\ell_i(s^i|\theta^*)} - 1 \right| < \epsilon \qquad \forall s^i \in S_i.$$

Therefore,

$$\left| \sum_{\theta \notin \bar{\Theta}_i} \mu_{i,t}(\theta) \frac{\ell_i(\hat{s}^i|\theta)}{\ell_i(\hat{s}^i|\theta^*)} + \sum_{\theta \in \bar{\Theta}_i} \mu_{i,t}(\theta) - 1 \right| < \epsilon,$$

where $\hat{s}^i$ is the signal that satisfies the inequality of Assumption $(\star)$. As a consequence, on any sample path that agent i's forecasts are eventually correct, and for large enough t,

$$0 \leq (1 - \delta_i) \sum_{\theta \notin \bar{\Theta}_i} \mu_{i,t}(\theta) < \epsilon.$$

Since, $\epsilon > 0$ is arbitrary, it must be the case that $\mu_{i,t}(\theta) \to 0$ as $t \to \infty$ for any $\theta \notin \bar{\Theta}_i$. Therefore, with $\mathbb{P}^*$-probability one, agent i assigns an asymptotic belief of zero to any state θ that from her point of view is not observationally equivalent to θ^*.

Now consider the belief update rule for agent i given by equation (3), evaluated at some state $\theta \notin \bar{\Theta}_i$:

$$\mu_{i,t+1}(\theta) = a_{ii} \mu_{i,t}(\theta) \frac{\ell_i(s_{t+1}^i|\theta)}{m_{i,t}(s_{t+1}^i)} + \sum_{j \in \mathcal{N}_i} a_{ij} \mu_{j,t}(\theta).$$

We have already shown that $\mu_{i,t}(\theta) \to 0$, $\mathbb{P}^*$-almost surely. However, this is possible only if $\sum_{j \in \mathcal{N}_i} a_{ij} \mu_{j,t}(\theta)$ converges to zero as well, which implies that $\mu_{j,t}(\theta) \to 0$ with

$\mathbb{P}^*$-probability one for all $j \in \mathcal{N}_i$. Note that this happens even if θ is observationally equivalent to θ^* from the point of view of agent j; that is, even if $\theta \in \bar{\Theta}_j$. As a result, all neighbors of agent i will assign an asymptotic belief of zero to parameter θ regardless of their signal structure. We can extend the same argument to the neighbors of neighbors of agent i, and by induction – since the social network is strongly connected – to all agents in the network. Thus, with $\mathbb{P}^*$-probability one,

$$\mu_{i,t}(\theta) \to 0 \qquad \forall i \in \mathcal{N} \quad , \quad \forall \theta \notin \bar{\Theta}_1 \cap \cdots \cap \bar{\Theta}_n.$$

This completes the proof. ■

B.4 Proof of Theorem 3

The proof of the Theorem 3 follows from the proof of Proposition 2, stated in subsection B.3. Recall, in the course of the proof of Proposition 2, we proved that

$$\mu_{i,t}(\theta) \to 0 \qquad \forall i \in \mathcal{N} \quad , \quad \forall \theta \notin \bar{\Theta}_1 \cap \cdots \cap \bar{\Theta}_n,$$

implying that all agents assign an asymptotic belief of zero on states that are not observationally equivalent to θ^* from the point of view of all individuals in the society. Therefore, statement (e) in the assumptions of Theorem 3 implies that $\mu_{i,t}(\theta) \to 0$ for all $\theta \neq \theta^*$, with $\mathbb{P}^*$– probability one, guaranteeing complete learning by all agents. ■

References

Acemoglu, Daron, Munther Dahleh, Ilan Lobel, and Asuman Ozdaglar (2008), "Bayesian learning in social networks." LIDS Working Paper ♯2780, Massachusetts Institute of Technology.

Acemoglu, Daron, Asuman Ozdaglar, and Ali Parandeh-Gheibi (2010), "Spread of (mis)information in social networks." *Games and Economic Behavior*. (Forthcoming).

Bala, Venkatesh and Sanjeev Goyal (1998), "Learning from neighbors." *The Review of Economic Studies*, 65, 595–621.

Bala, Venkatesh and Sanjeev Goyal (2001), "Conformism and diversity under social learning." *Economic Theory*, 17, 101–120.

Banerjee, Abhijit (1992), "A simple model of herd behavior." *The Quarterly Journal of Economics*, 107, 797–817.

Banerjee, Abhijit and Drew Fudenberg (2004), "Word-of-mouth learning." *Games and Economic Behavior*, 46, 1–22.

Berman, Abraham and Robert J. Plemmons (1979), *Nonnegative matrices in the mathematical sciences*. Academic Press, New York.

Bikchandani, Sushil, David Hirshleifer, and Ivo Welch (1992), "A theory of fads, fashion, custom, and cultural change as information cascades." *Journal of Political Economy*, 100, 992–1026.

DeGroot, Morris H. (1974), "Reaching a consensus." *Journal of American Statistical Association*, 69, 118–121.

DeMarzo, Peter M., Dimitri Vayanos, and Jeffery Zwiebel (2003), "Persuasion bias, social influence, and unidimensional opinions." *The Quarterly Journal of Economics*, 118, 909–968.

Ellison, Glenn and Drew Fudenberg (1993), "Rules of thumb for social learning." *The Journal of Political Economy*, 101, 612–643.

Ellison, Glenn and Drew Fudenberg (1995), "Word-of-mouth communication and social learning." *The Quarterly Journal of Economics*, 110, 93–126.

Epstein, Larry G., Jawwad Noor, and Alvaro Sandroni (2008a), "Non-Bayesian updating: a theoretical framework." *Theoretical Economics*, 3, 193–229.

Epstein, Larry G., Jawwad Noor, and Alvaro Sandroni (2008b), "Supplementary appendix for 'Non-Bayesian updating: a theoretical framework'." Working Paper 08-017, Penn Institute for Economic Research, University of Pennsylvania.

Gale, Douglas and Shachar Kariv (2003), "Bayesian learning in social networks." *Games and Economic Behavior*, 45, 329–346.

Golub, Benjamin and Mathew O. Jackson (2010), "Naïve learning in social networks and the wisdom of crowds." *American Economic Journal: Microeconomics*. (Forthcoming).

Hagerstrand, Torsten (1969), *Innovation Diffusion as a Spatial Process*. University of Chicago Press, Chicago, IL.

Kalai, Ehud and Ehud Lehrer (1994), "Weak and strong merging of opinions." *Journal of Mathematical Economics*, 23, 73–86.

Kirkland, Stephen J., Michael Neumann, and Bryan L. Shader (1998), "Applications of Paz's inequality to perturbation bounds for Markov chains." *Linear Algebra and its Applications*, 268, 183–196.

Kotler, Philip (1986), *The Principles of Marketing*, Third edition. Prentice-Hall, New York, NY.

Lehrer, Ehud and Rann Smorodinsky (1996), "Merging and learning." *Statistics, Probability and Game Theory*, 30, 147–168.

Miller, Ronald I. and Chris William Sanchirico (1997), "Almost everybody disagrees almost all the time: The genericity of weakly merging nowhere." Columbia Economics Department, Discussion Paper Series No. 9697-25.

Nachbar, John H. (1997), "Prediction, optimization, and learning in repeated games." *Econometrica*, 65, 275–309.

Paz, Azaria (1971), *Introduction to Probabilistic Automata*. Academic Press, New York.

Rogers, Everett M. (1983), *Diffusion of Innovations*, Third edition. Free Press, New York.

Sandroni, Alvaro (1998), "Necessary and sufficient conditions for convergence to Nash equilibrium: The almost absolute continuity hypothesis." *Games and Economic Behavior*, 22, 121–147.

Seneta, Eugene (1981), *Non-negative Matrices and Markov Chains*, Second edition. Springer, New York.

Smith, Lones and Peter Norman Sørensen (2000), "Pathological outcomes of observational learning." *Econometrica*, 68, 371–398.

Whitehead, Alfred North (1911), *An Introduction to Mathematics*. H. Holt and Company, New York, NY.