to appear in 2012 IEEE Conference on Decision and Control

Network Topology and Efficiency of Observational Social Learning

Pooya Molavi and Ali Jadbabaie

Abstract— **This paper explores the relationship between the topology of a network of agents and how efficiently they can learn a common unknown parameter. Agents repeatedly make private observations which are possibly informative about the unknown parameter; they also communicate their beliefs over the set of conceivable parameter values to their neighbors. It has been shown that for agents to learn the realized state, it is sufficient that they incorporate in their beliefs their private observations in a Bayesian way and the beliefs of their neighbors using a fixed linear rule. In this paper we establish upper and lower bounds on the rate by which agents performing such an update learn the realized state, and show that the bounds can be tight. These bounds enable us to compare efficiency of different networks in aggregating dispersed information. Our analysis yields an important insight: while for agents in large "balanced" social networks learning is much slower compared to that of a central observer, unbalanced networks could result in near efficient learning.**

I. Introduction

Social networks help shape the public opinion about different social, economic, and political issues by enabling individuals to gain information from the experiences of the others. In particular they influence individuals' opinions about hypotheses that are testable through observation, such as climate change or the risks associated with certain medical procedures. In light of this, it is important to understand when social networks could lead agents to hold accurate beliefs, whether the information of individual observations is aggregated efficiently, and how this depends on the network structure. It has been shown in [1] that under mild assumptions agents will eventually hold correct opinions. This paper builds upon [1] to provide answers to the last two of the questions posed above. We characterize the rate of learning (i.e., the rate by which agents' opinions become correct) in terms of information content of agents' observations and the structure of the social network. We then use the result to compare networks in terms of efficiency of information aggregation over them.

We base our analysis on the model of social learning proposed in [1]. Agents in a society desire to learn some unknown state which is drawn by nature from a finite set. They interact in a social structure that is represented by a weighted and possibly directed network. Agents have beliefs about the realized state, and use the information available to them to repeatedly update their beliefs. In every time period each agent privately observes a signal. It is assumed that agents cannot communicate their signals to each other; however, they observe the opinions held by their neighbors in the previous period.[1] Instead of processing all new information in a Bayesian way, agents use a simple rule to update their beliefs: Each agent first forms the Bayesian posterior given only her private signal, as an intermediate step. She then updates her belief to the convex combination of her Bayesian posterior and the beliefs of her neighbors, where the weights in the convex combination correspond to the trust she has in each of her neighbors. It is assumed that weights are fixed and independent of agents' observations.

The model provides a tractable framework to study the evolution of beliefs held by agents who interact in a social setting in addition to repeatedly making private observations. Repeated Bayesian updating over social networks is known to be computationally intractable—except for certain special cases—even when agents have knowledge of the network structure and other agents' observation models. This is since a fully Bayesian agent needs to form and update beliefs over the information of all other agents in the society, while only observing the beliefs of her neighbors. Incomplete information about the network structure and other agents' signal distributions intensifies these complications. The naïveté in the way agents incorporate the beliefs of their neighbors into their future beliefs in this model makes it tractable. It also serves to capture the assumption that agents are unaware of, and cannot learn the origin or quality of their neighbors' information. Nonetheless, it is shown in [1] that—for generic prior beliefs—the outcome of this model asymptotically agrees with that of a model with fully Bayesian agents.

The main contributions of this paper are as follows.

First, we find lower and upper bounds on the rate of learning, and show that the bounds can be tight. The rate of learning depends on the topology of the network and information content of individual agents' observations. This result signifies that even though network connectivity and global identifiability of the states are sufficient for agents to learn the realized state, the rate by which they do so is highly dependent on the network structure and informed agents' position in it. Namely, learning is faster whenever the most centrally located agents are also the ones who make the most informative observations.

Second, we show that the maximum efficiency by which agents can aggregate information is different in different networks. Social learning is inefficient in a large network

This research was supported by ONR MURI N000140810747 on Next Generation Network Science, AFOSR MURI CHASE, and AFOSR Complex Networks Program.

The authors are with the Department of Electrical and Systems Engineering and General Robotics, Automation, Sensing and Perception (GRASP) Laboratory, University of Pennsylvania, Philadelphia, PA 19104-6228, USA {pooya, jadbabaie}@seas.upenn.edu

[1] A rationale for this assumption is that agents' signals could belong to different spaces, rendering it difficult for an agent to interpret the signals of others, whereas all agents share the same belief space.

unless some agents are disproportionally central in it. Such agents correspond to "opinion leaders" whose opinions are observed by a non-vanishing fraction of all agents. Existence of opinion leaders is not sufficient, however, for efficient learning; rather, the opinion leaders also need to be the agents who make the most informative observations. Failure of this condition could provide an explanation for the widely observed phenomenon of individuals holding incorrect beliefs about factual issues, notwithstanding overwhelming evidence to the contrary.

This paper is related to a growing body of literature on learning over social networks, especially those with non-Bayesian updating rules. In this spirit are the works by DeGroot [2], DeMarzo *et al.* [3], Acemoglu *et al.* [4], Golub and Jackson [5], [6], and Jadbabaie *et al.* [1].

II. Model

A. Environment

There are a finite number of agents and possibilities for the state of the world. Let $\mathcal{N} = \{1, 2, \ldots, n\}$ be the set of agents, and let Θ be the finite set of possible states.

There are countably many time periods which are indexed by the non-negative integers. At time zero one of the states is realized. Each agent i has a prior belief about the realized state, denoted by $\mu_{i,0}(\cdot)$, which is a probability distribution over Θ. More generally, we denote the belief of agent i at time t by $\mu_{i,t}(\cdot)$. In every period $t \geqslant 1$, a signal profile $\omega_t = (\omega_{1,t}, \ldots, \omega_{n,t}) \in S_1 \times \cdots \times S_n = S$ is realized, where S_i is the finite set of possible signals for agent i. Conditional on the realized state being θ, the signal profiles are independent and identically distributed according to the likelihood function $\ell(\cdot|\theta)$. We assume that $\ell(s|\theta) > 0$ for all $s \in S$ and $\theta \in \Theta$.

The interactions between agents are captured by a directed graph $G = (\mathcal{N}, E)$. Agent j can observe the belief of agent i if there exists a directed edge from i to j, that is, if $(i, j) \in E$. Let $\mathcal{N}_i = \{j \in \mathcal{N} : (j, i) \in E\}$ be the set of neighbors of agent i whose beliefs she can observe.

B. Belief Update

Agents have two sources of information. At every time period $t \geqslant 1$, agent i privately observes $\omega_{i,t}$ in addition to the beliefs held by her neighbors in the previous time period. She then updates her belief using the information available to her. We assume that agents incorporate their private observations into their beliefs in a Bayesian way; however, they are naïve with respect to the other information available to them; namely,

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

where $\mathrm{BU}(\mu_{i,t}; \omega_{i,t+1})$ is the Bayesian update of $\mu_{i,t}$ given $\omega_{i,t+1}$, and a_{ij} are constants. Each agent updates her belief to a convex combination of her own Bayesian posterior, given only her private signal and neglecting the social network, and her neighbors' beliefs in the previous period. a_{ij} is the weight that agent i assigns to the opinion of agent j, and a_{ii}, called the *self-reliance* of agent i, is the weight she assigns to her Bayesian posterior conditional on her private signal. We assume that $a_{ij} > 0$ for all i and $j \in \mathcal{N}_i$ and that $\sum_{j \in \mathcal{N}_i \cup \{i\}} a_{ij} = 1$, in order for agents' beliefs to remain a probability distribution over Θ after they perform the update. The update in (1) can be written more explicitly as

$$\mu_{i,t+1}(\theta) = a_{ii}\frac{\ell_i(\omega_{i,t+1}|\theta)\mu_{i,t}(\theta)}{\sum_{\tilde{\theta} \in \Theta} \ell_i(\omega_{i,t+1}|\tilde{\theta})\mu_{i,t}(\tilde{\theta})} + \sum_{j \in \mathcal{N}_i} a_{ij}\mu_{j,t}(\theta), \qquad (2)$$

for all $\theta \in \Theta$.

C. Information

Let $(\Omega, \mathcal{F})$ be the measurable space where $\Omega = S^{\mathbb{N}}$ and $\mathcal{F}$ is the Borel σ-field over Ω. Let $\mathbb{P}^\theta = \ell(\cdot|\theta)^{\mathbb{N}}$ be the probability measure over Ω given that the realized state of the world is θ, and let $\mathbb{E}^\theta$ be the corresponding expectation operator.[2]

Agents' private observations are not necessarily informative about the realized state of the world. We use the *expected discrimination information* of agents' observations as a measure of how informative they are. The expected discrimination information (or simply information) of agent i's observation for θ over θ' is defined as

$$\mathcal{I}_i^\theta(\theta') = D_{\mathrm{KL}}(\ell_i(\cdot|\theta) \| \ell_i(\cdot|\theta')) = \mathbb{E}^\theta\left[\log\frac{\ell_i(\omega_{i,t}|\theta)}{\ell_i(\omega_{i,t}|\theta')}\right], \qquad (3)$$

where D_{KL} is the Kullback-Leibler divergence and ℓ_i is the marginal of ℓ over S_i. This is the expected information per observation for discriminating in favor of state θ against state θ', when state θ is realized. Note that $\mathcal{I}_i^\theta(\theta) = 0$. Gibbs' inequality implies that $\mathcal{I}_i^\theta(\theta') \geqslant 0$, and $\mathcal{I}_i^\theta(\theta') = 0$ if and only if θ is observationally equivalent to θ' from the point of view of agent i; that is,

$$\ell_i(s_i|\theta) = \ell_i(s_i|\theta'),$$

for all $s_i \in S_i$.

Let the *total expected discrimination information* (or simply total information) of agents' observations for θ over θ' be defined as

$$\mathcal{I}_{\mathrm{soc}}^\theta(\theta') = \sum_{i=1}^{n} \mathcal{I}_i^\theta(\theta').$$

This is the expected information per *observation profile* in favor of θ against θ', when the realized state is θ. $\mathcal{I}_{\mathrm{soc}}^\theta(\theta') = 0$ if and only if θ is observationally equivalent to θ' from the point of view of all agents.

III. Assumptions

We maintain the following assumptions throughout the paper. These are sufficient to guarantee that all agents eventually learn the realized state of the world. For a full characterization of necessary and sufficient conditions for learning see [7].

[2]More generally, we use a superscript θ to denote conditioning given that the realized state is θ.

Assumption 1: The social network is *strongly connected*.[3]

This assumption allows for information to flow from any agent to any other one. One can always assume connectivity without loss of generality, since otherwise each connected component could be analyzed separately. However, strong connectivity requires that any agent that influences other agents be influenced back by them, either directly or indirectly. This assumption excludes the scenarios where some stubborn agents place zero total weight on the beliefs of all other agents in the network.

Assumption 2: For any state $\theta \in \Theta$, there exists at least one agent with positive prior belief in θ.

This assumption requires that, no matter which state is realized, at least one of the agents' beliefs includes a "grain of truth". If this assumption is violated, we could face the rather uninteresting scenario where all agents continue to have zero belief in the realized state at all time periods.

Assumption 3: All agents have strictly positive self-reliance.

This assumption is a convenient way to guarantee that there is sufficient flow of new information into the social network. It requires agents to incorporate their private observations into their posterior beliefs. If all agents fail to do so, information cannot be accumulated over time, excluding the possibility of learning.

Assumption 4: For any two states $\theta \neq \theta'$, total information of agents' observations for θ over θ' is positive.

This assumption is necessary to ensure that agents' observations are sufficiently informative to let them distinguish θ from θ'. If this assumption is violated, even sophisticated Bayesian learners would not be able to learn the realized state.

IV. LEARNING

We are interested in characterizing the rate of learning for agents who use (2) to update their beliefs. In this section we formalize the notions of learning and rate of learning, and summarize the earlier results that establish learning.

The following proposition states that Assumptions 1–4 are sufficient to guarantee that agents' beliefs are asymptotically almost surely correct. The proof and a thorough discussion of the result's implications can be found in [1].

Proposition 1: If Assumptions 1–4 are satisfied and given that the realized state of the world is θ,

$$\mu_{i,t}(\cdot) \to \mu^{\theta}_{i,\infty}(\cdot) \quad \text{as} \quad t \to \infty \qquad \mathbb{P}^{\theta}\text{-a.s.},$$

for all $i \in \mathcal{N}$, where $\mu^{\theta}_{i,\infty}(\cdot)$ is the probability distribution over Θ defined as

$$\mu^{\theta}_{i,\infty}(\theta') = \begin{cases} 1 & \text{if} \quad \theta' = \theta, \\ 0 & \text{if} \quad \theta' \neq \theta. \end{cases}$$

We refer to the rate by which $\mu_{i,t}(\cdot)$ approaches $\mu^{\theta}_{i,\infty}(\cdot)$ in total variation (TV) distance as the rate of learning. Define

$$D^{\theta}_{i,t} = \|\mu_{i,\infty}(\cdot) - \mu^{\theta}_{i,t}(\cdot)\|_{\mathrm{TV}} = \sum_{\theta' \in \Theta \setminus \{\theta\}} \mu_{i,t}(\theta').$$

[3] A network is called strongly connected if there exists a directed path from any vertex to any other one.

Note that $D^{\theta}_{i,t}$ is equal to the total probability that agent i assigns to all states other than θ, given that θ is the realized state of the world. One can think of $D^{\theta}_{i,t}$, therefore, as agent i's *disbelief* in the realized state at time t. Let $\bar{D}^{\theta}_t$ be the average of all $D^{\theta}_{i,t}$; that is,

$$\bar{D}^{\theta}_t = \frac{1}{n} \sum_{i=1}^{n} D^{\theta}_{i,t}.$$

Note that the consequent of Proposition 1 can be written more compactly as $\bar{D}^{\theta}_t \to 0$ with $\mathbb{P}^{\theta}$-probability one, given that the realized state is θ. The next result shows that $\bar{D}^{\theta}_t$ converges to zero exponentially fast (equivalently, $\mu_{i,t}(\cdot)$ converges to $\mu^{\theta}_{i,\infty}(\cdot)$ exponentially fast in the TV distance), and provides a lower bound for the rate of learning. A proof can be find in [8].

Proposition 2: If Assumptions 1–4 are satisfied and given that the realized state is θ, for all $\epsilon > 0$,

$$\limsup_{t\to\infty} \frac{1}{t} \log \bar{D}^{\theta}_t \leqslant \lambda^{\theta}_1 + \epsilon, \qquad \mathbb{P}^{\theta}\text{-a.s.},$$

for some $\lambda^{\theta}_1 < 0$.

The lower bound on the rate of learning λ^{θ}_1 can be written in terms of the *top Lyapunov exponent* of a set of i.i.d. matrices. Let $M^{\theta}_t(\theta')$ be the $n \times n$ matrix defined as

$$M^{\theta}_t(\theta') = A + \mathrm{diag}\left(\left(a_{ii}\left(\frac{\ell_i(\omega_{i,t}|\,\theta')}{\ell_i(\omega_{i,t}|\,\theta)} - 1\right)\right)_{i=1,\dots,n}\right),$$

where $\mathrm{diag}(v_1, \dots, v_n)$ is the diagonal matrix with the ith diagonal element given by v_i. Note that since $\omega_{i,t}$ are i.i.d., the matrices $M^{\theta}_t(\theta')$ are i.i.d. as well. The top Lyapunov exponent (TLE) of $(M^{\theta}_t(\theta'))_{t=1,\dots,\infty}$ is defined as

$$\lambda^{\theta}_1(\theta') = \lim_{t\to\infty} \frac{1}{t} \log\|M^{\theta}_t(\theta') \cdots M^{\theta}_1(\theta')\|.$$

The Furstenberg-Kesten theorem [9] implies that for $\mathbb{P}^{\theta}$-almost all ω, the above limit exists and is independent of the realization of ω. The authors in [8] show that the lower bound on the rate of learning is given by

$$\lambda^{\theta}_1 = \max_{\theta' \in \Theta \setminus \{\theta\}} \lambda^{\theta}_1(\theta'). \tag{4}$$

V. BOUNDS ON THE RATE OF LEARNING

In this section we find analytic bounds on the rate of learning that depend on agents' centralities in the social network and information of their observations. Proposition 2 suggests that when the realized state is θ, $-\lambda^{\theta}_1$ could be used as a lower bound on the rate of learning. However, $\lambda^{\theta}_1(\theta')$ is the top Lyapunov exponent of a set of matrices, and determination of the TLE is known to be a difficult problem [10], [11]. Techniques for finding lower and upper bounds for the TLE of i.i.d. matrices have been proposed under some additional hypotheses [12], [13].

We use the upper bound on the TLE suggested by Gharavi and Anantharam in [13] to find a lower bound on the rate of learning. The bound is expressed as the solution to a concave maximization problem over a finite-dimensional convex set of probability distributions. In the appendix, we show that when $\mathcal{I}^{\theta}_{\text{soc}}(\theta')$ is small, the upper bound on $\lambda^{\theta}_1(\theta')$ can be approximated by

$$-\sum_{i\in\mathcal{N}} \pi_i a_{ii} \mathcal{I}^{\theta}_i(\theta'),$$

where $\pi_i > 0$ is the *eigenvector centrality* of agent i; that is, $(\pi_1, \ldots, \pi_n)$ is the left eigenvector of A corresponding to the unit eigenvalue normalized such that $\sum_{i=1}^n \pi_i = 1$.[4]

We can also find an upper bound on the rate of learning that depends on agents' centralities and the information of their observations. The following proposition establishes one such upper bound.

Proposition 3: If Assumptions 1–4 are satisfied, and given that the realized state is θ, for all $\epsilon > 0$,

$$\liminf_{t\to\infty} \frac{1}{t} \log \bar{D}^{\theta}_t \geqslant -r^{\theta} - \epsilon, \qquad \mathbb{P}^{\theta}\text{-a.s.},$$

where

$$r^{\theta} = \min_{\theta'\in\Theta\setminus\{\theta\}} \sum_{i\in\mathcal{N}} \pi_i a_{ii} \mathcal{I}^{\theta}_i(\theta'). \tag{5}$$

Proof: Let the realized state be denoted by θ. By Assumption 2, for any $\theta' \in \Theta$ there exists an agent j such that $\mu_{j,0}(\theta') > 0$. This implies by (2) that $\mu_{i,1}(\theta') > 0$ for all i such that $j \in \mathcal{N}_i$. Since the network is assumed to be strongly connected, one can repeat the same argument to conclude that $\mu_{i,t}(\theta') > 0$ for all $i \in \mathcal{N}$, $\theta' \in \Theta$ and $t > n$. We take logarithms of (2) evaluated at $t > n$ and use Jensen's inequality to get

$$\begin{aligned}\log \mu_{i,t+1}(\theta') &\geqslant a_{ii} \log\left(\frac{\ell_i(\omega_{i,t+1}|\theta')\mu_{i,t}(\theta')}{\sum_{\tilde\theta\in\Theta} \ell_i(\omega_{i,t+1}|\tilde\theta)\mu_{i,t}(\tilde\theta)}\right) \\ &\quad + \sum_{j\in\mathcal{N}_i} a_{ij} \log \mu_{j,t}(\theta') \\ &= a_{ii} \log \mu_{i,t}(\theta') + \sum_{j\in\mathcal{N}_i} a_{ij} \log \mu_{j,t}(\theta') \\ &\quad + a_{ii} \log \frac{\ell_i(\omega_{i,t+1}|\,\theta')}{\sum_{\tilde\theta\in\Theta} \ell_i(\omega_{i,t+1}|\tilde\theta)\mu_{i,t}(\tilde\theta)}.\end{aligned} \tag{6}$$

By Proposition 1, $\mu_{i,t}(\cdot) \to \mu^{\theta}_{i,t}(\cdot)$ with $\mathbb{P}^{\theta}$-probability one; hence,

$$\sum_{\tilde\theta\in\Theta} \ell_i(\cdot|\tilde\theta)\mu_{i,t}(\tilde\theta) \longrightarrow \ell_i(\cdot|\theta) \qquad \mathbb{P}^{\theta}\text{-a.s.} \tag{7}$$

Since $\ell_i(\cdot|\cdot) > 0$, equation (7) implies that for all $\epsilon > 0$ and $\mathbb{P}^{\theta}$-almost all ω, there exists T_i such that for all $t \geqslant T_i$,

$$\left|\log \frac{\ell_i(\omega_{i,t+1}|\,\theta')}{\sum_{\tilde\theta\in\Theta} \ell_i(\omega_{i,t+1}|\tilde\theta)\mu_{i,t}(\tilde\theta)} - \log \frac{\ell_i(\omega_{i,t+1}|\,\theta')}{\ell_i(\omega_{i,t+1}|\theta)}\right| < \epsilon. \tag{8}$$

Let $(\pi_1, \ldots, \pi_n)$ be the left eigenvector of A corresponding to the unit eigenvalue normalized such that $\sum_{i=1}^n \pi_i = 1$. Multiplying both sides of (6) by π_i, summing over $i \in \mathcal{N}$, and using (8) implies that for all $\epsilon > 0$ and $\mathbb{P}^{\theta}$-almost all ω, there exists T such that for all $t \geqslant T$,

$$\begin{aligned}\sum_{i=1}^n \pi_i \log \mu_{i,t+1}(\theta') &\geqslant \sum_{i=1}^n \pi_i \log \mu_{i,t}(\theta') \\ &\quad + \sum_{i\in\mathcal{N}} \pi_i a_{ii} \log \frac{\ell_i(\omega_{i,t+1}|\,\theta')}{\ell_i(\omega_{i,t+1}|\theta)} - \epsilon,\end{aligned}$$

which recursively implies that for all $t > t_0 \geqslant T$,

$$\begin{aligned}\sum_{i=1}^n \pi_i \log \mu_{i,t}(\theta') &\geqslant \sum_{i=1}^n \pi_i \log \mu_{i,t_0}(\theta') \\ &\quad + \sum_{\tau=t_0}^{t-1} \sum_{i\in\mathcal{N}} \pi_i a_{ii} \log \frac{\ell_i(\omega_{i,\tau+1}|\,\theta')}{\ell_i(\omega_{i,\tau+1}|\theta)} \\ &\quad - \epsilon(t - t_0).\end{aligned}$$

Hence, for all $\epsilon > 0$ and $\mathbb{P}^{\theta}$-almost all ω,

$$\begin{aligned}&\liminf_{t\to\infty} \frac{1}{t} \sum_{i=1}^n \pi_i \log \mu_{i,t}(\theta') \\ &\geqslant \liminf_{t\to\infty} \frac{1}{t} \sum_{\tau=t_0}^{t-1} \sum_{i\in\mathcal{N}} \pi_i a_{ii} \log \frac{\ell_i(\omega_{i,\tau+1}|\,\theta')}{\ell_i(\omega_{i,\tau+1}|\theta)} - \epsilon \\ &= \mathbb{E}^{\theta}\left[\sum_{i\in\mathcal{N}} \pi_i a_{ii} \log \frac{\ell_i(\omega_{i,t+1}|\,\theta')}{\ell_i(\omega_{i,t+1}|\theta)}\right] - \epsilon \\ &= -\sum_{i\in\mathcal{N}} \pi_i a_{ii} \mathcal{I}^{\theta}_i(\theta') - \epsilon,\end{aligned}$$

where the first equality is by the strong law of large numbers and the second one is by (3). Therefore, for all $\epsilon > 0$ and $\mathbb{P}^{\theta}$-almost all ω,

$$\begin{aligned}&\liminf_{t\to\infty} \frac{1}{t} \max_{\theta'\in\Theta\setminus\{\theta\}} \sum_{i=1}^n \pi_i \log \mu_{i,t}(\theta') \\ &\geqslant \max_{\theta'\in\Theta\setminus\{\theta\}} \left(-\sum_{i\in\mathcal{N}} \pi_i a_{ii} \mathcal{I}^{\theta}_i(\theta')\right) - \epsilon \\ &= -r^{\theta} - \epsilon.\end{aligned}$$

On the other hand,

$$\begin{aligned}\frac{1}{t} \log \bar{D}^{\theta}_t &= \frac{1}{t} \log\left(\frac{1}{n} \sum_{i=1}^n \sum_{\theta'\in\Theta\setminus\{\theta\}} \mu_{i,t}(\theta')\right) \\ &\geqslant \frac{1}{t} \log\left(\max_{i\in\mathcal{N}} \max_{\theta'\in\Theta\setminus\{\theta\}} \mu_{i,t}(\theta')\right) - \frac{1}{t} \log n \\ &= \frac{1}{t} \max_{\theta'\in\Theta\setminus\{\theta\}} \max_{i\in\mathcal{N}} \log \mu_{i,t}(\theta') - \frac{1}{t} \log n \\ &\geqslant \frac{1}{t} \max_{\theta'\in\Theta\setminus\{\theta\}} \sum_{i=1}^n \pi_i \log \mu_{i,t}(\theta') - \frac{1}{t} \log n,\end{aligned}$$

where the last inequality is since $\sum_{i=1}^n \pi_i = 1$. Hence, for all $\epsilon > 0$ and $\mathbb{P}^{\theta}$-almost all ω,

$$\liminf_{t\to\infty} \frac{1}{t} \log \bar{D}^{\theta}_t \geqslant -r^{\theta} - \epsilon - \limsup_{t\to\infty} \frac{1}{t} \log n = -r^{\theta} - \epsilon.$$

■

[4] Since A is the weighted adjacency matrix of a strongly connected graph, the Perron-Frobenius theorem implies A has a unique left eigenvector corresponding to the unit eigenvalue; furthermore, this eigenvector is elementwise positive.

Note that r^θ is equal to the approximate lower bound found earlier assuming that $\mathcal{I}^\theta_{\text{soc}}(\theta')$ is small for all θ'. Therefore, the upper and lower bounds on the rate of learning are both tight when the total information of agents' observations for θ over θ' is small, an assumption we maintain in the rest of the paper.

Equation (5) implies that for an agent to be influential in accelerating learning in the social network, three conditions need to be met. First, she needs to make observations which are highly informative in favor of the realized state. Second, she needs to have a high self-reliance. This is required for the agent to incorporate her observations in her belief with a large weight. Finally, she needs to be centrally located in the social network, as measured by eigenvector centrality. This allows her to influence the beliefs of other agents to a large extent.

VI. Network Topology and Rate of Learning

We are interested in comparing network structures in terms of how quickly they lead agents to learn the realized state. We can use r^θ defined in (5) as an approximation of the rate of learning. The minimum in (5) might be obtained, however, by different θ in different networks. This could complicate the comparison between rates of learning in different networks and obscure the resulting insights. To avoid this issue, in the rest of the paper we assume that $|\Theta| = 2$; that is, there are only two possibilities for the state. We let $\Theta = \{\theta, \theta'\}$, and let θ denote the realized state of the world. We also use $\mathcal{I}_i$ as shorthand for $\mathcal{I}^\theta_i(\theta')$ and r as shorthand for r^θ. Moreover, we make the dependence of the rate of learning on the network topology and information of agents' observations explicit by letting $r(A, \mathcal{I})$ denote the rate of learning over the network with adjacency matrix A when the information of agents' observations is given by $\mathcal{I} = (\mathcal{I}_1, \dots, \mathcal{I}_n)$.

A. Examples

In what follows we compute the rate of learning for three classes of networks in terms of their structural properties.

Example 1 (symmetric networks): A network is called symmetric if $a_{ij} = a_{ji}$ for all $i, j \in \mathcal{N}$. In a symmetric network $\pi_i = 1/n$ for all i. Therefore, the rate of learning for the symmetric network with adjacency matrix A_{sym} is given by

$$r(A_{\text{sym}}, \mathcal{I}) = \frac{1}{n}\sum_{i=1}^{n} a_{ii}\mathcal{I}_i.$$

Example 2 (k-regular network with equal weights): A network is called k-regular if each node has exactly k neighbors. Also assume that $a_{ii} = 1 - \epsilon$ for all i and $a_{ij} = \epsilon/k$ for all $j \in \mathcal{N}_i$; that is, agents all have the same self-reliance and trust all their neighbors equally. In this case $\pi_i = 1/n$ for all i, which implies that the rate of learning over such a k-regular network with adjacency matrix $A_{k\text{-reg}}$ is given by

$$r(A_{k\text{-reg}}, \mathcal{I}) = \frac{1}{n}\sum_{i=1}^{n}(1-\epsilon)\mathcal{I}_i.$$

Note that the rate of learning is independent of k.

Example 3 (star network): A star network is one where there is a central node that is a neighbor of all other nodes, whereas any other node is only neighbors with the central node. Let agent $i = 1$ denote the central agent in the star network, and let the adjacency matrix be given by

$$A_{\text{star}} = \begin{pmatrix} 1-\delta & \frac{\delta}{n-1} & \frac{\delta}{n-1} & \dots & \frac{\delta}{n-1} \\ 1-\epsilon & \epsilon & 0 & \dots & 0 \\ 1-\epsilon & 0 & \epsilon & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1-\epsilon & 0 & 0 & \dots & \epsilon \end{pmatrix}.$$

It is straightforward to show that agents' eigenvector centralities are given by[5]

$$\pi_i = \begin{cases} \dfrac{1-\epsilon}{1-\epsilon+\delta} & \text{if} \quad i = 1, \\ \dfrac{\delta}{(n-1)(1-\epsilon+\delta)} & \text{if} \quad i \neq 1. \end{cases}$$

Therefore, the rate of learning is given by

$$r(A_{\text{star}}, \mathcal{I}) = \frac{1}{1-\epsilon+\delta}\left[(1-\epsilon)(1-\delta)\,\mathcal{I}_1 + \frac{\delta\epsilon}{n-1}\sum_{i=2}^{n}\mathcal{I}_i\right].$$

The networks of Examples 1 and 2 are examples of "balanced" networks where all agents are (roughly) equally central, whereas in Example 3 agent $i = 1$ is much more central than other agents. Over balanced networks, information of agents' observations receives similar weights in the expression for the rate of learning. Consequently, the rate of learning is not sensitive to variations in the information of individual agents' observations—as long as the total information is constant. On the contrary, in a star network with a large number of agents, the rate of learning is very sensitive to variations in the information of the central agent's observations (even if the total information available to agents is fixed). This is because in a large star network the information of central agent's observations receives a much larger weight than that of peripheral agents.

B. Efficiency

Which social networks result in more efficient aggregation of the information contained in agents' observations? To answer this question, we use the rate of learning for a central observer with access to all agents' observations and likelihood functions $\ell_i(\cdot|\cdot)$ (but not the joint likelihood function $\ell(\cdot|\cdot)$) as a benchmark. The rate of learning for a such central observer is given by the Chernoff-Stein lemma as

$$r_c(\mathcal{I}) = \sum_{i=1}^{n}\mathcal{I}_i.$$

This is the fastest one can hope to learn the realized state, given all the information available to agents in the social

[5] See, e.g., [5].

network. For $\mathcal{I} \neq 0$ let $\alpha(A, \mathcal{I})$ be defined as

$$\alpha(A, \mathcal{I}) = \frac{r(A, \mathcal{I})}{r_c(\mathcal{I})}.$$

This is a measure of deceleration in the rate of learning as a result of decentralization of observations. A larger α corresponds to less deceleration, and hence, more efficient social learning. Note that if A corresponds to a strongly connected network, then $\alpha(A, \mathcal{I}) \in (0, 1)$.

The efficiency of social learning depends not only on the network topology, but also on which agents make the most informative observations; however, one can compare different networks directly in terms of the highest level of efficiency that is possible over them. Let $\alpha^*(A)$ be defined as

$$\alpha^*(A) = \sup_{\mathcal{I} > 0} r(A, \mathcal{I}).$$

That is, $\alpha^*(A)$ is the maximum achievable efficiency of social learning over the network corresponding to A. The networks in Examples 1–3 result in significantly different values for α^*:

$$\alpha^*(A_{\text{sym}}) = \frac{1}{n} \max_{i \in \mathcal{N}} a_{ii} < \frac{1}{n},$$
$$\alpha^*(A_{k\text{-reg}}) = \frac{1}{n}(1 - \epsilon),$$

whereas if n is sufficiently large

$$\alpha^*(A_{\text{star}}) = \frac{(1 - \epsilon)(1 - \delta)}{1 - \epsilon + \delta}.$$

While the rate of learning for the star network is independent of the number of agents, for the networks of Examples 1 and 2 it is a decreasing function of n. This implies that the balanced networks of Examples 1 and 2 cannot result in efficient learning when the number of agents is large.

More generally, the maximum achievable efficiency of social learning over the network with adjacency matrix A is given by

$$\alpha^*(A) = \max_i \pi_i a_{ii}.$$

This implies that for a network to be efficient in aggregating information it is necessary that the maximum eigenvector centrality among all agents is large; that is, there needs to exist an "opinion leader" who is disproportionally central in the network. This is not sufficient, however. For learning to be efficient, opinion leaders need to also be the agents who make the most informative observations. Even though strong connectivity is sufficient to ensure that the belief of even a non-central agent is eventually heard by others, a substantial part of the information contained in her observations is lost in the process. The following example illustrates this point. Suppose that the peripheral agent i observes a very informative signal that tilts her belief strongly in favor of state θ; however, as a result of her small centrality, her belief is observed only by few other agents. Let k be one such agent who observes the belief of agent i in addition to that of an uninformed agent j. Since k cannot know the quality of the observations made by i and j, she treats their beliefs as two equally valid "sides of the story". Therefore, agent k's belief will be less strongly in favor of θ than that of i. By the same argument, neighbors of k will be even less strongly in favor of θ compared to k, and so on.

A comparison with the result obtained by Golub and Jackson in [5] is due. The paper studies social learning for the case where agents receive only one noisy signal each about the unknown real-valued state of the world. Each agent repeatedly updates her estimate of the state by taking a convex combination of her estimate with those of her neighbors. The authors show that for agents in a large network to be able to form an unbiased estimate of the state, it is necessary that no agent has an eigenvector centrality which is non-vanishing in the number of agents. The authors call networks that satisfy this property wise. However, these are exactly the networks that in our model result in inefficient learning. This illustrates the tradeoff between correctness of agents' opinions when they have limited information, and the rate of learning when there is constant flow of new information. While if a central agent repeatedly makes informative observations she can lead all agents to quickly learn the state, if no agent makes repeated informative observations, all agents' could be misled by the central agent as a result of her excessive influence on others' beliefs.

VII. Conclusion

In this paper we studied a model of social learning where agents repeatedly update their beliefs to incorporate the information they obtain through both private observations and communication with their neighbors. We focused on scenarios where agents are unaware of (and cannot learn) the origin or quality of their neighbors' observations. We found lower and upper bounds on the rate of learning, and showed that the bounds are tight when the total information of agents' observations is small. We also compared networks in terms of their maximum achievable social learning efficiency. The analysis showed that even though in unbalanced networks agents could learn the realized state efficiently regardless of the network size, in large balanced networks social learning is always inefficient.

Throughout the paper our point of view was positive; however, the insights obtained are transferable to the problem of sensor network design. Consider the problem of decentralized detection where each sensor could sense a possibly different relevant variable. As we argued, in order for sensors to detect the event, they do not need to communicate to other sensors their entire observations; rather, it is sufficient for sensors to individually estimate the probability of the event and communicate their private estimates. This could save sensors valuable communication resources especially when their observations are large objects such as videos. Proposition 2 implies that, if the event is identifiable, this protocol leads sensors to learn whether the event has occurred or not exponentially fast. Moreover, the expression obtained for the rate of learning in (5) can be used to maximize the speed of detection subject to design constraints.

APPENDIX

We use the upper bound on the TLE found in [13] to approximate $\lambda_1^\theta(\theta')$. To make the derivation simpler we introduce some new notation. Let s denote the signal profile $(s_1, s_2, \dots, s_n)$ where $s_i \in S_i$ for all $i \in \mathcal{N}$. Let $\mathcal{S} = \{1, 2, \dots, |S|\}$, and let $h : \mathcal{S} \mapsto S$ be an enumeration of S.[6] For all $k \in \mathcal{S}$, we use $h_i(k)$ to denote the observation of agent i when signal profile $h(k)$ is realized; i.e., $h(k) = (h_1(k), \dots, h_n(k))$. For all $k \in \mathcal{S}$, let $p^k = \ell_1(h_1(k)|\theta) \cdots \ell_1(h_n(k)|\theta)$, and let N^k be the realization of $M_t^\theta(\theta')$ given $\omega_t = h(k)$; that is, N^k is the $n \times n$ matrix defined as

$$N^k = A + \text{diag}\left(\left(a_{ii}\left(\frac{\ell_i(h_i(k)|\theta')}{\ell_i(h_i(k)|\theta)} - 1\right)\right)_{i=1,\dots,n}\right).$$

Let $H(p)$ be the entropy of p defined as

$$H(p) = -\sum_{k \in \mathcal{S}} p^k \log p^k.$$

Let $\mathcal{M}$ be the set of all probability measures on $(\mathcal{N} \times \mathcal{S}) \times (\mathcal{N} \times \mathcal{S})$. With slight abuse of notation, for any $\eta \in \mathcal{M}$, let $H(\eta)$ be the entropy of η defined as

$$H(\eta) = -\sum_{\substack{i,j \in \mathcal{N} \\ k,l \in \mathcal{S}}} \eta_{i,j}^{k,l} \log \frac{\eta_{i,j}^{k,l}}{\eta_{i,*}^{k,*}},$$

where

$$\eta_{i,*}^{k,*} = \sum_{\substack{j \in \mathcal{N} \\ l \in \mathcal{S}}} \eta_{i,j}^{k,l}.$$

The solution to the following optimization problem is an upper bound on the TLE of the set of i.i.d. matrices $(M_t^\theta(\theta'))_{t=1,\dots,\infty}$ when $M_t^\theta(\theta') \in \{N^k\}_{k \in \mathcal{S}}$ and N^k is realized with probability p^k:

$$\hat{\lambda}_1^\theta(\theta') = \max_{\eta \in \mathcal{M}} H(\eta) + F(\eta) - H(p) \tag{9}$$

subject to

$$\eta_{*,*}^{k,l} = p^k p^l \quad \forall\, k,l \in \mathcal{S}, \tag{10}$$

$$\eta_{i,*}^{k,*} = \eta_{*,i}^{*,k} \quad \forall\, i \in \mathcal{N} \quad \forall k \in \mathcal{S}, \tag{11}$$

$$\eta_{i,j}^{k,l} = 0 \quad \forall\, i,j \in \mathcal{N} \quad \forall k,l \in \mathcal{S} \text{ s.t. } N_{j,i}^k = 0, \tag{12}$$

where $N_{j,i}^k$ is the element of N^k in the ith row and jth column, and $F : \mathcal{M} \mapsto \mathbb{R}$ is defined as

$$F(\eta) = \sum_{\substack{i,j \in \mathcal{N} \\ k,l \in \mathcal{S}}} \eta_{i,j}^{k,l} \log N_{j,i}^k.$$

We assume that the last set of constraints do not bind, and use Lagrange multipliers $\rho^{k,l}$ and ν_i^k to incorporate the first and second set of constraints into the objective function, respectively. The first order optimality conditions are given by

$$\log N_{j,i}^k - \log \eta_{i,j}^{k,l} + \log \eta_{i,*}^{k,*} + \rho^{k,l} + \nu_i^k - \nu_j^l = 0, \tag{13}$$

[6]An enumeration of a finite set S is a bijective mapping from $\{1, 2, ..., |S|\}$ to S.

for all $i,j \in \mathcal{N}$ and $k,l \in \mathcal{S}$. The above equation cannot be solved analytically, except for special cases. One such case is when the information of all agents' observations for θ over θ' is equal to zero; that is, $\mathcal{I}_i^\theta(\theta') = 0$ for all $i \in \mathcal{N}$. First, we solve for η assuming this. Then, we analyse the effect of a small perturbation in $\mathcal{I}_i^\theta(\theta')$ on $\hat{\lambda}_1^\theta(\theta')$.

First, assume that $\mathcal{I}_i^\theta(\theta') = 0$ for all i. In this case $\ell_i(h_i(k)|\theta) = \ell_i(h_i(k)|\theta')$ for all $i \in \mathcal{N}$ and $k \in \mathcal{S}$; consequently, $N^k = A$ for all $k \in \mathcal{S}$. It is easy to verify that the solution $(\hat{\eta}, \hat{\rho}, \hat{\nu})$ given below satisfies the KKT conditions (10)–(13).

$$\hat{\eta}_{i,j}^{k,l} = \pi_j a_{ji} p^k p^l$$
$$\hat{\rho}^{k,l} = \log p^k$$
$$\hat{\nu}_i^k = -\log(\pi_i p^k)$$

Note that since the set of constraints (12) are automatically satisfied, they are not binding in this case. Since the optimization problem (9) has a strictly concave cost function and a set of linear constraints, $\hat{\eta}$ is its unique solution. Substituting $\hat{\eta}$ in (9), one can easily see that $\hat{\lambda}_1^\theta(\theta') = 0$ when $\mathcal{I}_i^\theta(\theta') = 0$ for all i.

Next, for some agent i we perturb $\ell_i(\cdot|\theta')$ by the infinitesimal function $\delta\ell_i(\cdot|\theta')$ and find the resulting change in $\hat{\lambda}_1^\theta(\theta')$. For the perturbed likelihood function to remain a probability distribution over S_i, it is necessary that $\sum_{s_i \in S_i} \delta\ell_i(s_i|\theta') = 0$. Let $\nabla_{\delta\ell_i(\cdot|\theta')}\hat{\lambda}_1^\theta(\theta')$ denote the directional derivative of $\hat{\lambda}_1^\theta(\theta')$ along $\delta\ell_i(\cdot|\theta')$. Since the only constraint that depends on $\ell_i(\cdot|\theta')$ is not binding, by the envelope theorem, the derivative of $\hat{\lambda}_1^\theta(\theta')$ with respect to $\ell_i(\cdot|\theta')$ is given by the partial derivative of the $H(\eta) + F(\eta) - H(p)$ with respect to $\ell_i(\cdot|\theta')$ holding η fixed, and then evaluating the result at the optimal solution $\hat{\eta}$. While $H(p)$ and $H(\eta)$ do not explicitly depend on $\ell_i(\cdot|\theta')$, $F(\eta)$ is given by

$$F(\eta) = \sum_{\substack{i,j \in \mathcal{N} \\ k,l \in \mathcal{S}}} \eta_{i,j}^{k,l} \log a_{j,i} + \sum_{\substack{j \in \mathcal{N} \\ k,l \in \mathcal{S}}} \eta_{j,j}^{k,l} \log \frac{\ell_j(h_j(k)|\theta')}{\ell_j(h_j(k)|\theta)}.$$

Therefore,

$$\begin{aligned}
\nabla_{\delta\ell_i(\cdot|\theta')}\hat{\lambda}_1^\theta(\theta') &= \nabla_{\delta\ell_i(\cdot|\theta')} \sum_{\substack{j \in \mathcal{N} \\ k,l \in \mathcal{S}}} \hat{\eta}_{j,j}^{k,l} \log \frac{\ell_j(h_j(k)|\theta')}{\ell_j(h_j(k)|\theta)} \\
&= \sum_{\substack{j \in \mathcal{N} \\ k,l \in \mathcal{S}}} \hat{\eta}_{j,j}^{k,l} \nabla_{\delta\ell_i(\cdot|\theta')} \log \frac{\ell_j(h_j(k)|\theta')}{\ell_j(h_j(k)|\theta)} \\
&= \pi_i a_{ii} \sum_{k \in \mathcal{S}} p^k \nabla_{\delta\ell_i(\cdot|\theta')} \log \frac{\ell_i(h_i(k)|\theta')}{\ell_i(h_i(k)|\theta)} \\
&= \pi_i a_{ii} \sum_{s_i \in S_i} \ell_i(s_i|\theta) \nabla_{\delta\ell_i(\cdot|\theta')} \log \frac{\ell_i(s_i|\theta')}{\ell_i(s_i|\theta)},
\end{aligned}$$

where both of the derivatives are evaluated at $\ell_i(\cdot|\theta') = \ell_i(\cdot|\theta)$. Since $\ell_i(s_i|\theta)$ is not a function of $\ell_i(\cdot|\theta')$ and by linearity of the derivative operator, the above equation can be written as

$$\nabla_{\delta\ell_i(\cdot|\theta')}\hat{\lambda}_1^\theta(\theta') = \pi_i a_{ii} \nabla_{\delta\ell_i(\cdot|\theta')} \sum_{s_i\in S_i} \ell_i(s_i|\theta) \log \frac{\ell_i(s_i|\theta')}{\ell_i(s_i|\theta)}$$
$$= -\pi_i a_{ii} \nabla_{\delta\ell_i(\cdot|\theta')} \mathcal{I}_i^\theta(\theta').$$

Note that the above equation is valid for all $i \in \mathcal{N}$. On the other hand, $\hat{\lambda}_1^\theta(\theta') = 0$ when $\mathcal{I}_i^\theta(\theta') = 0$ for all $i \in \mathcal{N}$. Therefore, when $\mathcal{I}_i^\theta(\theta')$ is small for all $i \in \mathcal{N}$, $\hat{\lambda}_1^\theta(\theta')$ can be approximated by the linear term in its Taylor expansion with respect to $(\mathcal{I}_i^\theta(\theta'))_{i=1,\ldots,n}$ as

$$\hat{\lambda}_1^\theta(\theta') \approx -\sum_{i=1}^n \pi_i a_{ii} \mathcal{I}_i^\theta(\theta').$$

References

[1] A. Jadbabaie, P. Molavi, A. Sandroni, and A. Tahbaz-Salehi, "Non-bayesian social learning," *Games and Economic Behavior*, 2012. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0899825612000851

[2] M. H. DeGroot, "Reaching a consensus," *Journal of the American Statistical Association*, vol. 69, no. 345, pp. 118–121, 1974.

[3] P. M. DeMarzo, D. Vayanos, and J. Zwiebel, "Persuasion bias, social influence, and unidimensional opinions," *The Quarterly Journal of Economics*, vol. 118, no. 3, pp. 909–968, 2003.

[4] D. Acemoglu, A. Ozdaglar, and A. Parandeh-Gheibi, "Spread of (mis)information in social networks," *Games and Economic Behavior*, vol. 70, no. 2, pp. 194–227, 2010.

[5] B. Golub and M. Jackson, "Naïve learning in social networks and the wisdom of crowds," *American Economic Journal: Microeconomics*, vol. 2, no. 1, pp. 112–149, 2010.

[6] ——, "How homophily affects diffusion and learning in networks," *Arxiv preprint arXiv:0811.4013*, 2008.

[7] P. Molavi and A. Jadbabaie, "Aggregate observational distinguishability is necessary and sufficient for social learning," in *Proceedings of the 50th IEEE Conference on Decision and Control and European Control Conference*, 2011, pp. 2335–2340.

[8] P. Molavi, K. Rahnama Rad, A. Tahbaz-Salehi, and A. Jadbabaie, "On consensus and exponentially fast social learning," in *Proceedings of the 2012 American Control Conference*, 2012, to appear.

[9] H. Furstenberg and H. Kesten, "Products of random matrices," *The Annals of Mathematical Statistics*, vol. 31, no. 2, pp. 457–469, 1960.

[10] L. Arnold and V. Wihstutz, "Lyapunov exponents: a survey," *Lyapunov Exponents*, pp. 1–26, 1986.

[11] J. Tsitsiklis and V. Blondel, "The Lyapunov exponent and joint spectral radius of pairs of matrices are hard when not impossible to compute and to approximate," *Mathematics of Control, Signals, and Systems (MCSS)*, vol. 10, no. 1, pp. 31–40, 1997.

[12] E. Key, "Lower bounds for the maximal Lyapunov exponent," *Journal of Theoretical Probability*, vol. 3, no. 3, pp. 477–488, 1990.

[13] R. Gharavi and V. Anantharam, "An upper bound for the largest Lyapunov exponent of a Markovian product of nonnegative matrices," *Theoretical computer science*, vol. 332, no. 1-3, pp. 543–557, 2005.