2011 50th IEEE Conference on Decision and Control and
European Control Conference (CDC-ECC)
Orlando, FL, USA, December 12-15, 2011

Aggregate Observational Distinguishability is Necessary and Sufficient for Social Learning

Pooya Molavi and Ali Jadbabaie

Abstract— **We study a model of information aggregation and social learning recently proposed by Jadbabaie, Sandroni, and Tahbaz-Salehi, in which individual agents try to learn a correct state of the world by iteratively updating their beliefs using private observations and beliefs of their neighbors. No individual agent's private signal might be informative enough to reveal the unknown state. As a result, agents share their beliefs with others in their social neighborhood to learn from each other. At every time step each agent receives a private signal, and computes a Bayesian posterior as an intermediate belief. The intermediate belief is then averaged with the beliefs of neighbors to form the individual's belief at next time step. We find a set of necessary and sufficient conditions under which agents will learn the unknown state and reach consensus on their beliefs without any assumption on the private signal structure. The key enabler is a result that shows that using this update, agents will eventually forecast the indefinite future correctly.**

I. Introduction

Individuals often form opinions about economic, political, and social issues using both their personal experiences and information they obtain through communication with their friends, neighbors, and colleagues [1]. These beliefs determine the decisions they make when faced with different options. However, the "best" course of action available to each agent is, often, not obvious and depends on unknown variables (i.e., "states of the world"). Frequently, not all agents make the same observations, and not all observations are equally informative. Lack of access to all the relevant information is a motivation for individuals to share their opinions with others in order to learn from their personal experiences. Social networks expedite this process by enabling flow of information from informed agents to uninformed ones. In light of this, an important question one could ask is what are the least restrictive conditions (on the network and signal structures) that ensure learning by all agents? In this paper we use a non-Bayesian learning framework to provide an answer to this question.

We base our analysis on the model in [2] in which agents use a simple non-Bayesian rule to incorporate new information available to them. Each individual has two sources of information: her personal observations and those of her neighbors in a social network. However, agents might not have direct access to personal experiences of their neighbors. Instead, we assume that they can only observe their neighbors' beliefs, i.e., the subjective probabilities assigned to different feasible realizations of the unknown state of the world.[1] Agents repeatedly interact with their neighbors and use the following rule to update their beliefs. Each agent first forms the Bayesian posterior given her observed 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.

This model provides a tractable framework to study the opinion dynamics of agents who repeatedly receive private signals in addition to observing the opinions of their neighbors. This is our main motivation for considering a non-Bayesian protocol. Bayesian inference in social networks can be—except for certain simple scenarios—computationally complicated to carry out. Part of the complications is because there is no reason to believe that agents know the source of their neighbors' information. Rather, they have to infer it to be able to form an unbiased belief about the true state of the world. The complexities of Bayesian updating limit its applicability in practice.

In [2] the authors show that, under some assumptions, this update eventually leads to social learning even in finite networks. Namely, agents can eventually forecast the *immediate future* correctly. Furthermore, they will eventually learn the unknown state if for each agent there exists a signal that is the most probable under the true state of the world than any other state. This assumption "turns the deck" in favor of learning by assuming that agents are infinitely often notified, indirectly, of the true state of the world.

In Section III we show that agents will learn the state of the world under the much weaker assumption that they can distinguish the state collectively. Signals need not be independent among agents at the same time period. Instead, we require the signal structure to be such that the state is identifiable given the marginals of the likelihood function. We first prove that not only agents will forecast the immediate future correctly, but also they will eventually learn to forecast the indefinite future. We also show that there exists a signal sequence which is informative enough to let agents identify the true state of the world, even if no revealing signal exists. The results signify that even when none of agents have enough information to learn the true state of the world, and in spite of individual signals not being revealing enough, social interaction can aggregate pieces of information available to agents such that each and every one of them can distinguish the true state of the world.

In section IV we find a set of necessary and sufficient conditions for learning by relaxing some of the assumptions made in [2]. Namely, we prove that all agents learn the true state if and only if the state is observationally distinguishable by agents, and the social network can accommodate flow of information from informed agents to the uninformed ones. Furthermore, we show that if these assumptions are not satisfied, the social network can be partitioned into a number of "islands" together with the set of agents who do not belong to any island. With probability one, either all agents in an island learn the true state, or none do so. In absence of right paths for the flow of information, while some agents will almost surely learn the truth, some will almost surely not learn it. Furthermore,

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, jadbabai}@seas.upenn.edu

This research was supported in parts by ONR MURI No. N000140810747, and AFOSR Complex Networks Program.

[1]An alternative equally valid interpretation is that agents play a particular repeated game of imperfect information where each agent can only observe the actions of her neighbors, and the actions completely reveal the beliefs of acting agents.

978-1-61284-799-3/11/$26.00 ©2011 IEEE

in such instances agents will never reach consensus.

This paper is related to an extensive body of literature on social learning. For a survey of different models of social learning see [1]. To overcome the computational complications of Bayesian updating, various simplifications have been proposed. The first group of models assumes that agents interact sequentially. Examples include models in [3]–[7]. In such models each agent, having observed some of agents' previous actions, takes an action. However, each agent makes only one decision, and cannot reverse or change her choice. The other group of models suggests non-Bayesian rules of thumb for belief update. Examples include [2], [8]–[14]. In the same spirit is the seminal work of Tsitsiklis on decentralized decision making [15].

II. The Model

The social learning model we consider was first proposed in [2]. The model assumes that time is discrete and there are a finite number of agents, signals, and states of the world.

Let Θ be the finite set of possible states of the world, and let $\theta^* \in \Theta$ be the true state that is determined at time zero by nature and is unchanged thereafter. Without loss of generality we assume that $\Theta \neq \{\theta^*\}$.

Let $\mathcal{N} = \{1, 2, \dots, n\}$ be the set of agents. At time $t \geq 0$ each agent i has a belief about the true state denoted by $\mu_{i,t}(\theta)$, which is a probability distribution over Θ.

At each time period $t \geq 1$, each agent i observes a private random signal $\omega_{i,t} \in S_i$, where $S_i = \left\{s_i^1, s_i^2, \dots, s_i^{M_i}\right\}$ is the set of possible signals for agent i. Conditioned on $\theta \in \Theta$ being the state of the world, the observation profile $\omega_t = (\omega_{1,t}, \omega_{2,t}, \dots, \omega_{n,t})$ is generated according to likelihood function $\ell\left(\cdot|\,\theta\right)$ with $\ell_i\left(\cdot|\,\theta\right)$ as its ith marginal. Let $\mathbb{P}^\theta = \ell\left(\cdot|\,\theta\right)^{\mathbb{N}}$ be the product measure that determines the realization of signals conditioned on θ being the state of the world. This definition allows for signals to be correlated among agents at the same time period, but makes them independent over time. Without loss of generality we assume that $\ell_i\left(s_i\,|\theta^*\right) > 0$ for all $s_i \in S_i$. That is, S_i is only the set of signals that are realized with positive probabilities conditioned on the true state of the world being θ^*. Also let $\bar{\Theta}_i = \{\theta \in \Theta : \ell_i\left(s_i|\,\theta\right) = \ell_i\left(s_i|\,\theta^*\right) \text{ for all } s_i \in S_i\}$ be the set of states that are observationally equivalent to the true state θ^* from the point of view of agent i.

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

Agent i starts with the initial belief $\mu_{i,0}\left(\theta\right)$ that θ is the true state of the world. At the end of period t, each agent observes the beliefs of her neighbors. At the beginning of the next period, agent i receives the private signal $\omega_{i,t+1}$, and then uses the following rule to update her belief:

$$\mu_{i,t+1}\left(\theta\right) = a_{ii}\mu_{i,t}\left(\theta\right)\frac{\ell_i\left(\omega_{i,t+1}|\,\theta\right)}{m_{i,t}\left(\omega_{i,t+1}\right)} + \sum_{j \in \mathcal{N}_i} a_{ij}\mu_{j,t}\left(\theta\right), \quad (1)$$

where $m_{i,t}\left(s_i\right)$ is defined for any $s_i \in S_i$ as

$$m_{i,t}\left(s_i\right) = \sum_{\theta \in \Theta} \ell_i\left(s_i|\,\theta\right)\mu_{i,t}\left(\theta\right).$$

Using this rule, 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' previous period beliefs. a_{ij} is the weight 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 $a_{ij} > 0$ and $\sum_{j \in \mathcal{N}_i \cup \{i\}} a_{ij} = 1$ for the beliefs to remain a probability distribution over Θ after the update.

We use A to denote the $n \times n$ matrix whose ij element is a_{ij}, and use $\mu_t(\theta)$ to denote the n dimensional column vector whose ith element is $\mu_{i,t}(\theta)$.

$(\Omega, \mathcal{F}, \mathbb{P})$ is the probability triple, where $\Omega = \left(\prod_{i=1}^n S_i\right)^{\mathbb{N}}$, $\mathcal{F}$ is the smallest σ-field that makes all $\omega_{i,t}$ measurable, and $\mathbb{P} = \mathbb{P}^{\theta^*}$ is the probability distribution determining the realization of signals.[2] We use $\omega \in \Omega$ to denote the infinite signal sequence $(\omega_1, \omega_2, \dots)$, and $\mathbb{E}$ to denote the expected value operator with respect to the probability measure $\mathbb{P}$. Let $\mathcal{F}_t$ be the filtration generated by the observations of all agents up to time t.

We say that the adapted random variables X_t and Y_t are asymptotically $\mathbb{P}$-almost surely equal, denoted by $X_t \overset{aas}{=} Y_t$, if there exist $\tilde{\Omega} \in \mathcal{F}$ such that $\mathbb{P}(\tilde{\Omega}) = 1$, and for all $\omega \in \tilde{\Omega}$ and all $\epsilon > 0$, there exist $T\left(\omega, \epsilon\right)$ such that for all $t_1, t_2 > T\left(\omega, \epsilon\right)$,

$$|X_{t_1} - Y_{t_2}| < \epsilon.$$

It is an easy exercise to show that if $X_t \overset{aas}{=} Y_t$ and $Z_t \overset{aas}{=} W_t$, then $X_t \pm Z_t \overset{aas}{=} Y_t \pm W_t$ and $X_tZ_t \overset{aas}{=} Y_tW_t$.

For all $t > 0$, $\mu_{i,t}\left(\cdot\right)$ and $m_{i,t}\left(\cdot\right)$ are random functions adapted to $\mathcal{F}_t$, the former on Θ and the latter on S_i. $m_{i,t}\left(s_i\right)$ is the probability that agent i assigns, at time t, to signal s_i being observed in the next time step, hence, it is called agent i's *one step forecast*. We can extend this notion to define the *k-step forecast* $m_{i,t}^{(k)}\left(s_{i,1}, s_{i,2}, \dots, s_{i,k}\right)$ as the forecast at time t of agent i that the signal sequence $\left(s_{i,1}, s_{i,2}, \dots, s_{i,k}\right) \in \left(S_i\right)^k$ will be realized in the next k time steps. More formally,

$$m_{i,t}^{(k)}\left(s_{i,1}, \dots, s_{i,k}\right) = \sum_{\theta \in \Theta} \ell_i\left(s_{i,1}|\,\theta\right)\dots\left(s_{i,k}|\,\theta\right)\mu_{i,t}\left(\theta\right).$$

III. Asymptotic Learning

In this section we find a set of sufficient conditions for learning when agents use (1) to update their beliefs. We maintain the following assumptions throughout the section:

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

Assumption 2: There exists an agent with positive prior belief on the true parameter θ^*.

Assumption 1 allows for information to flow from any agent to any other one. Assumption 2 is similar to what is known as a "grain of truth" in agents' prior beliefs [16].

In [2] the authors show that if agents use the update in (1), the ones having positive self-reliance learn to forecast the *immediate future* correctly.[4]

Proposition 1 (Tahbaz-Salehi, Sandroni, and Jadbabaie): If Assumptions 1 and 2 are satisfied, for any agent i such that $a_{ii} > 0$,

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

with $\mathbb{P}$-probability one.

[2] $\mathbb{N}$ stands for the set of natural numbers.

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

[4] The proposition is proved assuming that all agents have positive self-reliance. However, even if not all agents have positive self-reliance, the proof still applies to those agents who do so, without any modifications.

The authors also show that agents will learn the true state asymptotically $\mathbb{P}-$almost surely. However, to prove asymptotic learning (in addition to maintaining Assumptions 1 and 2 and assuming that all agents have positive self-reliance), the authors assume that for any agent i, there exists a signal $\hat{s}_i \in S_i$ and a positive number δ_i such that

$$\frac{\ell_i\left(\hat{s}_i\right|\theta)}{\ell_i\left(\hat{s}_i\right|\theta^*)} \leq \delta_i < 1 \qquad \forall\theta \in \Theta \quad \text{such that} \quad \theta \neq \theta^*. \quad (2)$$

This assumption asks for existence of a signal that is more likely conditioned on θ^* being the true state of the world rather than conditioned on any other state in Θ being the true state of the world. Under this assumption, and provided that the conditions of Proposition 1 hold, the authors prove that agents asymptotically learn the true parameter θ^* with $\mathbb{P}$-probability one. The condition in (2) guarantees that there exists a "revealing signal" that is observed infinitely often.

To prove asymptotic learning no assumption should be made other than *observational distinguishability* of the true state of the world θ^* by agents who have positive self-reliance.

Assumption 3: There is no $\theta \in \Theta$ that is observationally equivalent to θ^* from the point of view of all agents who have positive self-reliance, that is,

$$\bar{\Theta} \stackrel{\text{def}}{=} \bigcap_{\{i\in\mathcal{N}:a_{ii}>0\}} \bar{\Theta}_i = \{\theta^*\}.$$

This is obviously a necessary condition for agents to learn the true state of the world. In Proposition 4, which is the main result of this section, we show that it is also sufficient. To this end, we first show in Proposition 2 that correct forecasts of agents can be extended into the future. To prove these results we first need to present a preliminary lemma.

Lemma 1: If Assumptions 1 and 2 are satisfied, for any agent i with positive self-reliance,

$$\mathbb{E}\left[\mu_{i,t+1}(\theta)\,|\,\mathcal{F}_t\right] \stackrel{aas}{=} a_{ii}\mu_{i,t}(\theta) + \sum_{j\in\mathcal{N}_i} a_{ij}\mu_{j,t}(\theta).$$

Proof: By Proposition 1,

$$\mathbb{E}\left[\frac{\ell_i\left(\omega_{i,t+1}\right|\theta)}{m_{i,t}\left(\omega_{i,t+1}\right)}\Big|\,\mathcal{F}_t\right] = \sum_{s_i\in S_i} \ell_i\left(s_i\right|\theta^*)\,\frac{\ell_i\left(s_i\right|\theta)}{m_{i,t}\left(s_i\right)} \stackrel{aas}{=} \sum_{s_i\in S_i} \ell_i\left(s_i\right|\theta) = 1. \quad (3)$$

On the other hand, taking expectations of both sides of (1) conditioned on $\mathcal{F}_t$,

$$\mathbb{E}\left[\mu_{i,t+1}(\theta)\,|\,\mathcal{F}_t\right] = a_{ii}\mu_{i,t}(\theta)\,\mathbb{E}\left[\frac{\ell_i\left(\omega_{i,t+1}\right|\theta)}{m_{i,t}\left(\omega_{i,t+1}\right)}\Big|\,\mathcal{F}_t\right] + \sum_{j\in\mathcal{N}_i} a_{ij}\mu_{j,t}(\theta).$$

Evaluating the above equation for large t and using (3) completes the proof. ■

The next result shows that not only agents eventually forecast the next step correctly, as shown in Proposition 1, but also they do so for the next k steps for any *finite* k.

Proposition 2: If Assumptions 1 and 2 are satisfied, for any agent i with positive self-reliance,

$$m_{i,t}^{(k)}\left(s_{i,1},\ldots,s_{i,k}\right) \to \prod_{r=1}^{k} \ell_i\left(s_{i,r}\right|\theta^*) \quad \text{as} \quad t\to\infty,$$

with $\mathbb{P}$-probability one for all $s_{i,1}, s_{i,2}, \ldots, s_{i,k} \in S_i$.

Proof: To simplify notation, in the proof we drop the subscript i from $s_{i,1}, s_{i,2}, \ldots, s_{i,k}$ and $\omega_{i,t+1}$. We use induction on k. For $k = 1$ the result is proved in Proposition 1. Let $s_2, \ldots, s_k \in S_i$ be arbitrary. Multiplying both sides of (1) by $m_{i,t}\left(\omega_{t+1}\right)\prod_{r=2}^{k}\ell_i\left(s_r\right|\theta)$ and summing over $\theta \in \Theta$,

$$\begin{aligned} m_{i,t}\left(\omega_{t+1}\right) m_{i,t+1}^{(k-1)}\left(s_2,\ldots,s_k\right) &= a_{ii}m_{i,t}^{(k)}\left(\omega_{t+1},s_2,\ldots,s_k\right) \\ &\quad + m_{i,t}\left(\omega_{t+1}\right)\sum_{\theta\in\Theta}\prod_{r=2}^{k}\ell_i\left(s_r\right|\theta)\sum_{j\in\mathcal{N}_i}a_{ij}\mu_{j,t}(\theta). \end{aligned}$$

By Lemma 1,

$$\begin{aligned} m_{i,t}\left(\omega_{t+1}\right) m_{i,t+1}^{(k-1)}\left(s_2,\ldots,s_k\right) &\stackrel{aas}{=} a_{ii}m_{i,t}^{(k)}\left(\omega_{t+1},s_2,\ldots,s_k\right) \\ &\quad + m_{i,t}\left(\omega_{t+1}\right)\sum_{\theta\in\Theta}\prod_{r=2}^{k}\ell_i\left(s_r\right|\theta)\,\mathbb{E}\left[\mu_{i,t+1}(\theta)\,|\mathcal{F}_t\right] \\ &\quad - a_{ii}m_{i,t}\left(\omega_{t+1}\right)\sum_{\theta\in\Theta}\prod_{r=2}^{k}\ell_i\left(s_r\right|\theta)\,\mu_{i,t}(\theta). \end{aligned}$$

We can change the order of sum and expectation using Fubini's theorem to get

$$\begin{aligned} m_{i,t}\left(\omega_{t+1}\right) m_{i,t+1}^{(k-1)}\left(s_2,\ldots,s_k\right) &\stackrel{aas}{=} a_{ii}m_{i,t}^{(k)}\left(\omega_{t+1},s_2,\ldots,s_k\right) \\ &\quad + m_{i,t}\left(\omega_{t+1}\right)\mathbb{E}\left[m_{i,t+1}^{(k-1)}\left(s_2,\ldots,s_k\right)|\,\mathcal{F}_t\right] \\ &\quad - a_{ii}m_{i,t}\left(\omega_{t+1}\right)m_{i,t}^{(k-1)}\left(s_2,\ldots,s_k\right). \end{aligned}$$

By the induction hypothesis, $m_{i,t}^{(k-1)}\left(s_2,s_3,\ldots,s_k\right)$ is equal to $\prod_{r=2}^{k}\ell_i\left(s_r\right|\theta^*)$ and $m_{i,t}\left(\omega_{t+1}\right)$ is equal to $\ell_i\left(\omega_{t+1}\right|\theta^*)$, asymptotically $\mathbb{P}$-almost surely. Furthermore, $m_{i,t+1}^{(k-1)}(s_2,\ldots,s_k)$ is bounded above by one for all t. Hence, $\mathbb{E}\left[m_{i,t+1}^{(k-1)}\left(s_2,\ldots,s_k\right)|\,\mathcal{F}_t\right]$ converges $\mathbb{P}$-almost surely to $\prod_{r=2}^{k}\ell_i\left(s_r\right|\theta^*)$.[5] Therefore, reorganizing the terms of the above expression,

$$m_{i,t}^{(k)}\left(\omega_{t+1},s_2,\ldots,s_k\right) \stackrel{aas}{=} \ell_i\left(\omega_{t+1}|\theta^*\right)\prod_{r=2}^{k}\ell_i\left(s_r|\theta^*\right),$$

which implies that

$$\left|m_{i,t}^{(k)}\left(\omega_{t+1},s_2,\ldots,s_k\right) - \ell_i\left(\omega_{t+1}|\theta^*\right)\prod_{r=2}^{k}\ell_i\left(s_r|\theta^*\right)\right| \stackrel{aas}{=} 0.$$

Taking expectations of the above expression conditioned on $\mathcal{F}_t$ and using the bounded convergence theorem for conditional expectations,

$$\sum_{s_1\in S_i}\left|m_{i,t}^{(k)}\left(s_1,\ldots,s_k\right) - \prod_{r=1}^{k}\ell_i\left(s_r|\theta^*\right)\right|\ell_i\left(s_1|\theta^*\right) \stackrel{aas}{=} 0.$$

[5]This is due to the bounded convergence theorem for conditional expectations. For a proof see, for instance, page 263 of [17].

Since $\ell_i(s_1|\theta^*) > 0$ for all $s_1 \in S_i$,

$$m_{i,t}^{(k)}(s_1, s_2, \ldots, s_k) \overset{aas}{=} \prod_{r=1}^{k} \ell_i(s_r|\theta^*),$$

for all $s_1, \ldots, s_k \in S_i$ ■

The following proposition uses finiteness of S_i to show that for any agent i there exists a long enough signal sequence that is more probable under θ^* than any state $\theta \notin \bar{\Theta}_i$.

Proposition 3: For any agent i, there exists a finite number $\hat{k}_i$ and signals $\hat{s}_{i,1}, \hat{s}_{i,2}, \ldots, \hat{s}_{i,\hat{k}_i}$ such that

$$\frac{\prod_{r=1}^{\hat{k}_i} \ell_i(\hat{s}_{i,r}|\theta)}{\prod_{r=1}^{\hat{k}_i} \ell_i(\hat{s}_{i,r}|\theta^*)} \leq \delta_i < 1 \qquad \forall \theta \notin \bar{\Theta}_i. \tag{4}$$

Proof: First assume that $\ell_i(s_i|\theta^*)$ is a rational number for all $s_i \in S_i$. In this case we will prove that we can take $\hat{k}_i$ to be the least common denominator (LCD) of $\{\ell_i(s_i|\theta^*)\}_{s_i \in S_i}$ and $\left(\hat{s}_{i,1}, \hat{s}_{i,2}, \ldots, \hat{s}_{i,\hat{k}_i}\right)$ to be a sequence of signals in which the number of occurrences of each $s_i \in S_i$ is exactly equal to the numerator of the fractional representation of $\ell_i(s_i|\theta^*)$ when the denominator is equal to $\hat{k}_i$. In other words, we pick the signal sequence in which the frequency of each signal is equal to its probability under θ^*.

Let $\hat{k}_i = \text{LCD}\left(\{\ell_i(s_i|\theta^*)\}_{s_i \in S_i}\right)$ and $k_i^j = \ell(s_i^j|\theta^*)\hat{k}_i$ for $1 \leq j \leq M_i$. We prove that $\ell_i(\cdot|\theta^*)$ is the unique probability measure for which the probability of the signal sequence

$$\Big(\underbrace{s_i^1, \ldots, s_i^1}_{k_i^1 \text{times}}, \ldots, \underbrace{s_i^{M_i}, \ldots, s_i^{M_i}}_{k_i^{M_i} \text{times}}\Big)$$

is maximized. As a result, for this sequence $\ell_i(\cdot|\theta)/\ell_i(\cdot|\theta^*)$ is strictly less than one for any $\theta \notin \bar{\Theta}_i$.

Let $p_i^j = \mathbb{P}_i(s_i^j)$ for $1 \leq j \leq M_i$, where $\mathbb{P}_i$ is some probability measure on S_i. To simplify notation we drop the subscript i whenever there is no risk of confusion. We solve the following concave maximization problem:

$$\begin{aligned} \max_{p^1,\ldots,p^M} \quad & \left(p^1\right)^{k^1}\left(p^2\right)^{k^2}\cdots\left(p^M\right)^{k^M} \\ \text{subject to} \quad & p^1 + p^2 + \cdots + p^M = 1, \end{aligned} \tag{5}$$

where by $\left(p^j\right)^{k^j}$ we mean p^j to the power of k^j. This problem has the unique solution

$$p^j = \frac{k^j}{\sum_{j=1}^{M} k^j} \qquad 1 \leq j \leq M.$$

By construction, this solution corresponds to the probability measure $\ell_i(\cdot|\theta^*)$.

For the case that $\ell_i(s_i|\theta^*)$ is irrational for some $s_i \in S_i$, the proof follows from continuity of the objective of the optimization problem (5) with respect to $\left(p^1, p^2, \ldots, p^j\right)$ and the fact that rational numbers are dense in reals. ■

The condition in (4) is the k-step counterpart of the condition in (2). We have shown that even though a single signal revealing the true state might not exist, but there are signal sequences that will serve to do so.

We are now ready to prove the main result of this section.

Proposition 4: If Assumptions 1–3 are satisfied,

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

with $\mathbb{P}$-probability one for all $i \in \mathcal{N}$.

Proof: Fix an arbitrary $\theta \neq \theta^*$. By Assumption 3, there exists an agent, call her i, such that both $a_{ii} > 0$ and $\theta \notin \bar{\Theta}_i$. Let $\hat{k}_i$ and $\left(\hat{s}_{i,1}, \hat{s}_{i,2}, \ldots, \hat{s}_{i,\hat{k}_i}\right)$ be a positive integer and a sequence of signals, respectively, that satisfy (4) for agent i. By Proposition 2, $m_{i,t}(s_{i,1}, \cdots, s_{i,k}) \rightarrow \ell_i(s_{i,1}, \cdots, s_{i,k}|\theta^*)$ with $\mathbb{P}$-probability one for any sequence of finite length. We can use this result for $\hat{s}_{i,1}, \hat{s}_{i,2}, \ldots, \hat{s}_{i,\hat{k}_i}$ to conclude that

$$\sum_{\tilde{\theta} \in \Theta} \mu_{i,t}(\tilde{\theta}) \frac{\ell_i\left(\hat{s}_{i,1}, \ldots, \hat{s}_{i,\hat{k}_i}|\tilde{\theta}\right)}{\ell_i\left(\hat{s}_{i,1}, \ldots, \hat{s}_{i,\hat{k}_i}|\theta^*\right)} - 1 \longrightarrow 0,$$

and therefore,

$$\sum_{\tilde{\theta} \in \Theta \setminus \bar{\Theta}_i} \mu_{i,t}(\tilde{\theta}) \frac{\ell_i\left(\hat{s}_{i,1}, \ldots, \hat{s}_{i,\hat{k}_i}|\tilde{\theta}\right)}{\ell_i\left(\hat{s}_{i,1}, \ldots, \hat{s}_{i,\hat{k}_i}|\theta^*\right)} + \sum_{\tilde{\theta} \in \bar{\Theta}_i} \mu_{i,t}(\tilde{\theta}) - 1 \rightarrow 0,$$

with $\mathbb{P}$-probability one. Using the result of Proposition 3 we can conclude there exist $\delta_i > 0$ such that for $\mathbb{P}$-almost all ω,

$$0 \leq (1 - \delta_i) \sum_{\tilde{\theta} \in \Theta \setminus \bar{\Theta}_i} \mu_{i,t}(\tilde{\theta}) \rightarrow 0.$$

This proves that $\mu_{i,t}(\theta) \rightarrow 0$ as $t \rightarrow \infty$. Taking limits of (1) as $t \rightarrow \infty$ and using the result proved above shows that $\sum_{j \in \mathcal{N}_i} a_{ij}\mu_{j,t}(\theta)$ converges to zero as $t \rightarrow \infty$ and so does $\mu_{j,t}(\theta)$ for all $j \in \mathcal{N}_i$, with $\mathbb{P}$-probability one. Proceeding in the same way and using the strong connectivity assumption, for all $j \in \mathcal{N}$, $\mu_{j,t}(\theta) \rightarrow 0$. Thus, with $\mathbb{P}$-probability one,

$$\mu_{i,t}(\theta) \rightarrow 0 \quad \text{as} \quad t \rightarrow \infty \qquad \forall i \in \mathcal{N}.$$

Since $\theta \neq \theta^*$ was arbitrary and $\mu_{i,t}(\cdot)$ is a probability distribution over Θ for all i and t,

$$\mu_{i,t}(\theta^*) \rightarrow 1 \quad \text{as} \quad t \rightarrow \infty \qquad \forall i \in \mathcal{N},$$

with $\mathbb{P}$-probability one. ■

IV. Necessary and Sufficient Conditions

Which one of the Assumptions 1–3 are necessary for learning? Clearly Assumptions 2 and 3 are necessary; violation of Assumption 2 constrains agents to hold zero beliefs over θ^* for all t and ω, whereas, if Assumption 3 is violated, then there exists a state $\hat{\theta}$ that is not distinguishable from θ^* by any of agents, i.e., $\theta^* \neq \hat{\theta} \in \bar{\Theta}$. In the latter case, agents' signals contain no information that could help them distinguish θ^* from $\hat{\theta}$.

In the following proposition we formalize this intuition by showing that if Assumption 3 is violated, agents will *not* learn the true state, $\mathbb{P}$-almost surely. However, if $\mu_{i,0}(\theta^*) = 1$ for all agents, then they will learn the true state for all ω. To exclude this (non-generic) case, we make an assumption similar to Assumption 2.

Assumption 4: There exists at least one agent i such that $\mu_{i,0}(\hat{\theta}) > 0$ for some $\theta^* \neq \hat{\theta} \in \bar{\Theta}$.

Proposition 5: If Assumptions 1 and 4 are satisfied and Assumption 3 is *not* satisfied,

$$\mu_{i,t}(\theta^*) \nrightarrow 1 \quad \text{as} \quad t \rightarrow \infty \qquad \forall i \in \mathcal{N},$$

with $\mathbb{P}$-probability one.

The proof is along the lines of that of Lemma 2 in [2].

Proof: Let $\theta^* \neq \hat{\theta} \in \bar{\Theta}$ be the state that satisfies Assumption 4. Since the network is strongly connected, after at most n

time steps all agents will have positive beliefs over $\hat{\theta}$. Taking logarithms of (1) evaluated at $\hat{\theta}$ and using Jensen's inequality,

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

Taking expectations conditioned on $\mathcal{F}_t$,

$$\mathbb{E}\left[\log \mu_{i,t+1}(\hat{\theta})|\mathcal{F}_t\right] \geq a_{ii} \log \mu_{i,t}(\hat{\theta}) + \sum_{j\in\mathcal{N}_i} a_{ij} \log \mu_{j,t}(\hat{\theta}) + a_{ii}\mathbb{E}\left[\log \frac{\ell_i(\omega_{i,t+1}|\,\hat{\theta})}{m_{i,t}(\omega_{i,t+1})}|\mathcal{F}_t\right]. \quad (6)$$

On the other hand, by Jensen's inequality and since $\ell_i(\cdot|\,\hat{\theta}) = \ell_i(\cdot|\,\theta^*)$ for all agents i such that $a_{ii} > 0$,

$$\mathbb{E}\left[\log \frac{\ell_i(\omega_{i,t+1}|\,\hat{\theta})}{m_{i,t}(\omega_{i,t+1})}|\mathcal{F}_t\right] = -\mathbb{E}\left[\log \frac{m_{i,t}(\omega_{i,t+1})}{\ell_i(\omega_{i,t+1}|\,\hat{\theta})}|\mathcal{F}_t\right] \geq -\log \mathbb{E}\left[\frac{m_{i,t}(\omega_{i,t+1})}{\ell_i(\omega_{i,t+1}|\,\hat{\theta})}|\mathcal{F}_t\right] = 0,$$

for all agents with positive self-reliance. Therefore, (6) can be written in vector form as

$$\mathbb{E}\left[\log \mu_{t+1}(\hat{\theta})|\mathcal{F}_t\right] \geq A \log \mu_t(\hat{\theta}).$$

Since the network is strongly connected, by Perron-Frobenius theorem, A has a positive left eigenvector v corresponding to the unit eigenvalue. Left multiplying the above equation by v,

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

Therefore, $\sum_{i\in\mathcal{N}} v_i \log \mu_{i,t}(\hat{\theta})$ is a submartingale, with respect to filtration $\mathcal{F}_t$, which is bounded above by zero. Hence, it converges $\mathbb{P}$-almost surely. Since $v_i > 0$ for all $i \in \mathcal{N}$, this implies that $\mu_{i,t}(\hat{\theta})$ is $\mathbb{P}$-almost surely uniformly bounded away from zero for all i. Therefore, since $\mu_{i,t}(\cdot)$ is a probability distribution, $\mu_{i,t}(\theta^*)$ is $\mathbb{P}$-almost surely uniformly bounded away from one for all i. ■

Assumption 1, on the other hand, is not necessary for learning. For instance, if the network is not strongly connected, but each agent can distinguish θ^* from all the other states by herself, then agents will still learn the true state. What we really require is existence of information paths to any agent from informed agents. In Theorem 6, which is the main result of this paper, we find a set of necessary and sufficient conditions for learning by dropping the strong connectivity assumption and changing Assumptions 2 and 3 accordingly. To do so, first, we have to introduce some definitions and simple results. The following can be found in [13].

A group of agents $\mathcal{N}' \subseteq \mathcal{N}$ is *closed* if there exist no two agents $i \in \mathcal{N}'$ and $j \in \mathcal{N} \setminus \mathcal{N}'$ such that $(j, i) \in E$. A closed group of agents $\mathcal{N}'$ is *minimal* if no non-empty strict subset of $\mathcal{N}'$ is closed. The induced subgraph on any minimal closed group of agents is strongly connected. Moreover, if a graph $(\mathcal{N}, E)$ is strongly connected, then its only minimal closed group is $\mathcal{N}$.

The following assumption is sufficient for learning (together with Assumption $2'$ below) when agents use the update in (1).

Assumption $3'$: For any minimal closed group $\mathcal{N}'$ and any $\theta \neq \theta^*$, there exists at least one agent with positive self-reliance who can distinguish θ from θ^*, that is,

$$\bar{\Theta}_{\mathcal{N}'} \stackrel{\text{def}}{=} \bigcap_{\{i\in\mathcal{N}':a_{ii}>0\}} \bar{\Theta}_i = \{\theta^*\},$$

for all $\mathcal{N}'$ minimal closed.

We also have to modify Assumptions 2 and 4 by requiring them to be satisfied for each closed and minimal group separately.

Assumption $2'$: For any minimal closed group $\mathcal{N}'$, there exists at least one agent $i \in \mathcal{N}'$ such that $\mu_{i,0}(\theta^*) > 0$.

Assumption $4'$: For any minimal closed group $\mathcal{N}'$, there exists at least one agent $i \in \mathcal{N}'$ and one state $\theta^* \neq \hat{\theta} \in \bar{\Theta}_{\mathcal{N}'}$ such that $\mu_{i,0}(\hat{\theta}) > 0$.

The following theorem generalizes Propositions 4 and 5 by summarizing the necessary and sufficient conditions for learning (or absence thereof).

Theorem 6: The following statements are true for $\mathbb{P}$-almost all ω:

1) All agents will asymptotically learn the true state if Assumptions $2'$ and $3'$ are satisfied.
2) At least one agent will *not* asymptotically learn the true state if Assumption $3'$ is *not* satisfied and Assumption $4'$ is satisfied.

Proof: The key observation enabling the proof is that the evolution of beliefs of agents belonging to a minimal closed group is independent of that of other agents—or their existence in the network. Therefore, every minimal closed group can be analyzed ignoring the rest of agents. Every minimal closed group is strongly connected. Therefore, Propositions 4 and 5 can be directly applied to them.

If Assumption $3'$ is not satisfied and Assumption $4'$ is satisfied, by Proposition 5, there exists at least one closed minimal group whose members do not learn the true state. This proves the second part of the theorem.

On the other hand, if Assumptions $2'$ and $3'$ are satisfied, Proposition 4 implies that agents in minimal closed groups will learn the true state with $\mathbb{P}$-probability one. Let $\bar{\mathcal{N}}$ be the set of agents who do not belong to any minimal closed group. If $\bar{\mathcal{N}}$ is the empty set, the proof is complete. In what follows we assume that $\bar{\mathcal{N}} \neq \emptyset$.

Evaluating (1) at θ^* and taking conditional expectations,

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

where the inequality is Jensen's. Note that this does not follow from Lemma 1 as it is obtained without assuming strong connectivity. Taking expectation of the above equation, in vector form,

$$\mathbb{E}\left[\mu_{t+1}(\theta^*)\right] \geq A\mathbb{E}\left[\mu_t(\theta^*)\right],$$

which by induction implies,

$$\mathbb{E}\left[\mu_{2t}(\theta^*)\right] \geq A^t\mathbb{E}\left[\mu_t(\theta^*)\right]. \quad (7)$$

Agents can be reordered such that A takes the following canonical form:

$$A = \begin{pmatrix} A_{11} & A_{12} \\ 0 & A_{22} \end{pmatrix},$$

where A_{11}, A_{12}, and A_{22} are $|\bar{\mathcal{N}}| \times |\bar{\mathcal{N}}|$, $|\bar{\mathcal{N}}| \times |\mathcal{N} \setminus \bar{\mathcal{N}}|$, and $|\mathcal{N} \setminus \bar{\mathcal{N}}| \times |\mathcal{N} \setminus \bar{\mathcal{N}}|$ matrices respectively. Furthermore, A_{22} is a block diagonal matrix with diagonal blocks corresponding to different minimal closed groups of agents. By Assumption $3'$ each of the diagonal blocks of A_{22} has at least one positive diagonal element. Therefore, each such diagonal block is primitive. Consequently, the limit of A^t exists, as $t \to \infty$, and is equal to

$$\lim_{t\to\infty} A^t = \begin{pmatrix} 0 & B \\ 0 & C \end{pmatrix}, \tag{8}$$

where B and C are both stochastic matrices. All the above definitions and results can be found in Chapter 8 of [18].

Taking lim inf's of (7) and using (8) implies that for all $i \in \bar{\mathcal{N}}$,

$$\liminf_{t\to\infty} \mathbb{E}\left[\mu_{i,t}(\theta^*)\right] \geq \sum_{j\in\mathcal{N}\setminus\bar{\mathcal{N}}} b_{ij} \liminf_{t\to\infty} \mathbb{E}\left[\mu_{j,t}(\theta^*)\right],$$

where b_{ij} is the ij element of B. For all $j \in \mathcal{N} \setminus \bar{\mathcal{N}}$, $\mu_{j,t}(\theta^*)$ converges to one with $\mathbb{P}$-probability one. Therefore, by Fatou's lemma, for all $i \in \bar{\mathcal{N}}$,

$$\liminf_{t\to\infty} \mathbb{E}\left[\mu_{i,t}(\theta^*)\right] \geq \sum_{j\in\mathcal{N}\setminus\bar{\mathcal{N}}} b_{ij} = 1.$$

Since $\mu_{i,t}(\theta^*)$ takes value only in $[0, 1]$, this implies that

$$\mu_{i,t}(\theta^*) \to 1 \quad \text{as} \quad t \to \infty \qquad \forall i \in \bar{\mathcal{N}}.$$

■

Proof of Theorem 6 also provides some insight about the outcome when Assumption 3 is satisfied only for some of the minimal closed groups. If so, the set of agents can be partitioned into "islands" of minimal closed groups plus a set of agents that belong to no minimal closed group. With $\mathbb{P}$-probability one and in each island, either all agents learn the true state or no agent learns the true state, depending on whether Assumption 3 is satisfied in that island or not. Therefore, if the social network contains more than one group and Assumption $3'$ is not satisfied, then agents in different islands will (generically) never reach consensus.

Assumptions $2'$ and $4'$ are both satisfied for generic prior beliefs. That is, if $\mathbb{Q}$ is an absolutely continuous probability measure on agents' prior beliefs, these assumptions are satisfied with $\mathbb{Q}$-probability one. Theorem 6 is the most general characterization of necessary and sufficient conditions—even for non-generic prior beliefs. However, if agents' prior beliefs are probabilistic, one can simplify Theorem 6 as follows.

Corollary 7: Let agents in a social network use (1) to update their beliefs. Then the following statements hold for $\mathbb{Q}$-almost all prior beliefs and $\mathbb{P}$-almost all observation sequences:

1) All agents will asymptotically learn the true state if and only if Assumption $3'$ is satisfied.
2) *Not* all agents will asymptotically learn the true state if and only if Assumption $3'$ is *not* satisfied.

This second part of the corollary does not follow from the first one directly. The first statement implies that the second one holds, only with a positive probability. One could then prove the second part using a zero-one law. However, here we have taken the direct approach.

V. Conclusion

We found a set of necessary and sufficient conditions for learning in a non-Bayesian model of social learning. Learning happens if agents have non-degenerate prior beliefs, the state is observationally distinguishable, and there exists a path of information flow to any agent from informed agents. This is in contrast to the results regarding consensus which require strong connectivity. As a result of constant inflow of information about the truth in our model, agents can learn the truth even in networks that fail to be strongly connected. In spite of our result being "what one would hope for", due to the highly non-linear nature of our update, the proof required establishing an intermediate result first: Once the immediate future forecasts of agents become approximately correct, they can be extended to indefinite future with a negligible error.

Under the same assumptions, not only agents eventually learn the true state of the world, but also they do so exponentially fast with an exponent that depends both on network topology and on signal structure. Moreover, as a corollary of exponential learning, eventually the consensus dynamic becomes dominant. Hence, even if the true state is not distinguishable, agents eventually reach consensus in a strongly connected network. On the other hand, if the network is not strongly connected, we will eventually have several islands such that agents in each island will reach consensus, whereas agents belonging to different islands might not. We defer the presentation of these results to a future paper.

Acknowledgments

The authors would like to thank Alireza Tahbaz-Salehi for careful review of our proofs and many helpful comments and discussions.

References

[1] S. Goyal. (2010) Learning in networks. [Online]. Available: http://www.econ.cam.ac.uk/faculty/goyal/wp10/surveypapergoyal.pdf

[2] A. Tahbaz-Salehi, A. Sandroni, and A. Jadbabaie, "Learning under social influence," in *Proceedings of the 48th IEEE Conference on Decision and Control, held jointly with the 28th Chinese Control Conference*, 2009, pp. 1513–1519.

[3] A. V. Banerjee, "A simple model of herd behavior," *The Quarterly Journal of Economics*, vol. 107, no. 3, pp. 797–817, 1992.

[4] S. Bikhchandani, D. Hirshleifer, and I. Welch, "A theory of fads, fashion, custom, and cultural change as informational cascades," *The Journal of Political Economy*, vol. 100, no. 5, pp. 992–1026, 1992.

[5] L. Smith and P. Sørensen, "Pathological outcomes of observational learning," *Econometrica*, vol. 68, no. 2, pp. 371–398, 2000.

[6] A. Banerjee and D. Fudenberg, "Word-of-mouth learning," *Games and Economic Behavior*, vol. 46, no. 1, pp. 1–22, Jan. 2004.

[7] D. Acemoglu, M. Dahleh, I. Lobel, and A. Ozdaglar, "Bayesian learning in social networks," 2008, NBER Working Paper No. W14040. [Online]. Available: http://www.nber.org/papers/w14040

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

[9] G. Ellison and D. Fudenberg, "Rules of thumb for social learning," *The Journal of Political Economy*, vol. 101, no. 4, pp. 612–643, 1993.

[10] ——, "Word-of-mouth communication and social learning," *The Quarterly Journal of Economics*, vol. 110, no. 1, pp. 93–125, 1995.

[11] V. Bala and S. Goyal, "Learning from neighbours," *Review of Economic Studies*, vol. 65, no. 3, pp. 595–621, 1998.

[12] 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.

[13] 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.

[14] 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.

[15] J. Tsitsiklis, "Problems in decentralized decision making and computation," Ph.D. dissertation, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, 1984.

[16] E. Kalai and E. Lehrer, "Rational learning leads to nash equilibrium," *Econometrica*, vol. 61, no. 5, pp. pp. 1019–1045, 1993.

[17] R. Durrett, *Probability: Theory and Examples*, 4th ed. Cambridge University Press, 2010.

[18] C. Meyer, *Matrix analysis and applied linear algebra: solutions manual*. Society for Industrial and Applied Mathematics, 2000, vol. 2.