Learning Over Social Networks via Diffusion Adaptation

Xiaochuan Zhao and Ali H. Sayed
Department of Electrical Engineering
University of California, Los Angeles, CA 90095

***Abstract*—We propose a diffusion strategy to enable social learning over networks. Individual agents observe signals influenced by the state of the environment. The individual measurements are not sufficient to enable the agents to detect the true state of the environment on their own. Agents are then encouraged to cooperate through a diffusive process of self-learning and social-learning. We show that the diffusion algorithm converges almost surely to the true state. Simulation results also illustrate the superior convergence rate of the diffusion strategy over consensus-based strategies since diffusion schemes allow information to diffuse more thoroughly through the network.**

***Index Terms*—Diffusion adaptation, non-Bayesian learning, social networks, consensus, belief update.**

I. Introduction

One problem of particular interest is the study of mechanisms by which beliefs and opinions are formed and propagated through social networks [1]–[3]. There are two major approaches that are often considered to model social interactions and their related learning procedures. In the Bayesian learning approach, agents employ the Bayesian rule to incorporate information from their observations and from their neighbors into their private beliefs [1], [2], [4]–[6]. In the non-Bayesian learning approach, some heuristic updates as in (8) and (11) further ahead are used to guide the learning process by the agents [3], [7]–[16].

The work [3] is particularly relevant to our discussion since it proposes a *fully-distributed*, non-Bayesian learning procedure whereby beliefs of the agents are shown to converge almost surely to the true model. A distributed solution is attractive because connections over social networks tend to be dynamic and interactions need to be localized. The learning procedure employed by [3] relies on a consensus-type construction — see (8) further ahead. In their construction, agents combine their private beliefs with the *old* beliefs from their neighbors. In this article, we propose an alternative construction that relies on diffusion strategies [17]–[19] as opposed to consensus strategies. In the diffusion procedure, agents combine their beliefs with their neighbors' *updated* beliefs. In doing so, information ends up diffusing more thoroughly through the network. In the context of mean-square-error (MSE) estimation, it was shown in [20] that diffusion strategies outperform consensus strategies in terms of MSE performance. The simulations further ahead illustrate the faster convergence rate of the diffusion social learning process over the consensus alternative. To facilitate comparison of our results and algorithm with the presentation in [3], we use similar notation whenever possible.

This work was supported by NSF grant CCF-1011918. Email: {xzhao, sayed}@ee.ucla.edu.

II. Problem Formulation

Consider a group of N agents interconnected through a certain topology over a social network. Let $\mathcal{N} \triangleq \{1, 2, \dots, N\}$ denote the indexes of the agents in the network. Let Θ denote a finite set of all possible events that can be detected by the social network. Let $\theta^o \in \Theta$ denote an *unknown* event that has happened and the objective of the network is to select the event that is most likely to have occurred based on available observations.

Initially at time $i = 0$, each agent k in the network assumes a private prior belief, denoted by $\mu_{k,0}(\theta) \in [0, 1]$, which represents the probability distribution over the events $\theta \in \Theta$, i.e., $\mu_{k,0}(\theta) = \mathbb{P}(\boldsymbol{\theta} = \theta)$; note that in our notation, we use boldface letters to denote random variables and plain letters for their realizations. For subsequent time instants $i \geq 1$, the private belief of agent k is denoted by $\mu_{k,i}(\theta) \in [0, 1]$. All beliefs across all agents must be valid probability measures over the entire event set Θ. That is, they must obey the following constraint:

$$\sum_{\theta \in \Theta} \mu_{k,i}(\theta) = 1 \tag{1}$$

for any $k \in \mathcal{N}$ and $i \geq 0$. Agents continually update their private beliefs $\{\mu_{k,i}(\theta)\}$ over time based on the private signals they observe from the environment and the information shared by their social neighbors.

At each time $i \geq 1$, we assume that a signal profile $\boldsymbol{\xi}_i \triangleq (\boldsymbol{\xi}_{1,i}, \boldsymbol{\xi}_{2,i}, \dots, \boldsymbol{\xi}_{N,i})$ is randomly drawn from a probability distribution $L(\cdot)$ dependent on the true event θ^o:

$$\boldsymbol{\xi}_i \sim L(s_1, s_2, \dots, s_N | \theta^o) \tag{2}$$

where $\boldsymbol{\xi}_{k,i} \in \mathcal{S}_k$ and $\boldsymbol{\xi}_i \in \mathcal{S} \triangleq \mathcal{S}_1 \times \mathcal{S}_2 \times \cdots \times \mathcal{S}_N$. We assume that the signal space $\mathcal{S}$ is finite and therefore each partial signal space $\mathcal{S}_k$ is also finite. Each agent k can only observe its partial signal $\boldsymbol{\xi}_{k,i}$, which follows the marginal distribution $L_k(s_k|\theta^o)$:

$$\boldsymbol{\xi}_{k,i} \sim L_k(s_k|\theta^o) \triangleq \sum_{\{s_\ell \in \mathcal{S}_\ell, \ell \neq k\}} L(s_1, s_2, \dots, s_N|\theta^o) \tag{3}$$

We assume that $L_k(s_k|\theta) > 0$ for any $s_k \in \mathcal{S}_k$, $\theta \in \Theta$, and $k \in \mathcal{N}$. We further assume that the signals $\{\boldsymbol{\xi}_i\}$ are temporally independent and define the history of the sequence $\{\boldsymbol{\xi}_i\}$ up to time $j \geq 1$ as the σ-algebra generated by $\{\boldsymbol{\xi}_i; 1 \leq i \leq j\}$:

$$\mathcal{F}_j \triangleq \sigma(\boldsymbol{\xi}_1, \boldsymbol{\xi}_2, \ldots, \boldsymbol{\xi}_j) \tag{4}$$

We assume a *directed* graph model in which each agent is connected to some local neighbors and hence can get information from them. We denote the neighborhood of agent k by $\mathcal{N}_k \subseteq \mathcal{N}$ such that any agent $\ell \in \mathcal{N}_k$ is connected to agent k. We assume $k \in \mathcal{N}_k$ for convenience. We assume that the network is strongly-connected, namely, there exists a path connecting any pair of agents in the network, and, moreover, there is at least one self-loop. We further assume that agents can only share their private beliefs, $\{\mu_{k,i}(\theta)\}$, rather than their private signals, $\{\xi_{k,i}\}$, with their local neighbors. This is because the support of beliefs, Θ, is common for all agents but the supports for the private signals, $\{\mathcal{S}_k\}$, may be quite different across the agents.

We further assume that the individual agents may not be able to identify the unknown event θ^o on their own [3]. Specifically, we assume that for every agent k, there exists a nonempty subset of events, $\Theta_k \subseteq \Theta$, with maybe more than one element such that

$$L_k(s_k|\theta) = L_k(s_k|\theta^o) \tag{5}$$

for any $\theta \in \Theta_k$ and $s_k \in \mathcal{S}_k$. Therefore, from agent k's perspective, all events in Θ_k are equivalent to θ^o and there is no way for agent k to distinguish between these events on its own. However, to guarantee that the identification problem is still feasible, we require,

$$\bigcap_{k \in \mathcal{N}} \Theta_k = \{\theta^o\} \tag{6}$$

We denote the complement of Θ_k by $\bar{\Theta}_k$ such that $\Theta_k \cap \bar{\Theta}_k = \varnothing$ and $\Theta_k \cup \bar{\Theta}_k = \Theta$.

III. Diffusion Social Learning

At every time $i \geq 1$, each agent k first updates its private belief based on its observed private signal $\xi_{k,i}$ (which is a realization of $\boldsymbol{\xi}_{k,i}$) by means of the Bayesian rule:

$$\psi_{k,i}(\theta) = \frac{\mu_{k,i-1}(\theta) L_k(\xi_{k,i}|\theta)}{\sum_{\theta' \in \Theta} \mu_{k,i-1}(\theta') L_k(\xi_{k,i}|\theta')} \tag{7}$$

This step leads to an intermediate belief $\psi_{k,i}(\theta)$. After learning from their observed signals, agents can then learn from their social neighbors through cooperation. For example, in [3] a convex combination rule was proposed for fusing the $\{\mu_{\ell,i-1}\}$:

$$\mu_{k,i}(\theta) = \alpha_{kk}\, \psi_{k,i}(\theta) + \sum_{\ell \in \mathcal{N}_k \setminus \{k\}} \alpha_{\ell k}\, \mu_{\ell,i-1}(\theta) \tag{8}$$

where

$$\alpha_{\ell k} > 0 \text{ if } \ell \in \mathcal{N}_k,\ \sum_{\ell \in \mathcal{N}_k} \alpha_{\ell k} = 1,\ \alpha_{\ell k} = 0 \text{ otherwise} \tag{9}$$

We collect the combination weights $\{\alpha_{\ell k}\}$ into an $N \times N$ matrix A such that its (ℓ, k)-th element is $\alpha_{\ell k}$. In general, matrix A is asymmetric due to the directed graph model. It follows from (9) that A is a left-stochastic matrix. Since the network is assumed to be strongly-connected, then it can be verified that A is primitive. By the Perron-Frobenius theorem [21], there exists a positive vector $y \in \mathbb{R}_+^{N \times 1}$ such that

$$A\, y = y \quad \text{and} \quad y^{\mathsf{T}} \mathbb{1} = 1 \tag{10}$$

where $\mathbb{1}$ denotes the vector with all entries equal to one.

If neighbors share instead their intermediate private beliefs $\{\psi_{\ell,i}\}$, agent k can fuse the information by using

$$\mu_{k,i}(\theta) = \alpha_{kk}\, \psi_{k,i}(\theta) + \sum_{\ell \in \mathcal{N}_k \setminus \{k\}} \alpha_{\ell k}\, \psi_{\ell,i}(\theta) \tag{11}$$

Rule (11), where the intermediate beliefs are aggregated, is an extension of the class of diffusion strategies [17]–[19] to the context of social learning. In comparison, rule (8) is reminiscent of consensus strategies where the old beliefs are aggregated. Combing (7) and (11) we arrive at the two-step diffusion social learning process that we study in this work:

$$\begin{cases} \psi_{k,i}(\theta) = \dfrac{\mu_{k,i-1}(\theta) L_k(\xi_{k,i}|\theta)}{\sum_{\theta' \in \Theta} \mu_{k,i-1}(\theta') L_k(\xi_{k,i}|\theta')} & \text{(self-learning)} \\ \mu_{k,i}(\theta) = \sum_{\ell \in \mathcal{N}_k} \alpha_{\ell k}\, \psi_{\ell,i}(\theta) & \text{(social-learning)} \end{cases} \tag{12}$$

Using the updated private belief $\mu_{k,i}(\theta)$, agent k can now make a forecast or prediction on the probability of a certain signal $s_k \in \mathcal{S}_k$ occurring in the next time instance $i+1$. This probability can be computed as follows

$$m_{k,i}(s_k) \triangleq \sum_{\theta \in \Theta} \mu_{k,i}(\theta) L_k(s_k|\theta) \tag{13}$$

In the sequel, we shall establish the fact that the agents in the social network eventually learn the truth, namely, the private beliefs of agents converge asymptotically and almost surely to an impulse of size one at the location $\theta = \theta^o$.

IV. Convergence Analysis

For the sake of the analysis, we view the signal sequence $\{\boldsymbol{\xi}_i; i \geq 1\}$ as a stochastic process. Therefore, the diffusion social learning process (12) becomes a stochastic system of equations, which we rewrite as

$$\begin{cases} \boldsymbol{\psi}_{k,i}(\theta) = \dfrac{\boldsymbol{\mu}_{k,i-1}(\theta) L_k(\boldsymbol{\xi}_{k,i}|\theta)}{\sum_{\theta' \in \Theta} \boldsymbol{\mu}_{k,i-1}(\theta') L_k(\boldsymbol{\xi}_{k,i}|\theta')} \\ \boldsymbol{\mu}_{k,i}(\theta) = \sum_{\ell \in \mathcal{N}_k} \alpha_{\ell k}\, \boldsymbol{\psi}_{\ell,i}(\theta) \end{cases} \tag{14}$$

Likewise, the forecast $\{\boldsymbol{m}_{k,i}(s_k); i \geq 1\}$ also becomes a stochastic process and we rewrite it as

$$\boldsymbol{m}_{k,i}(s_k) = \sum_{\theta \in \Theta} \boldsymbol{\mu}_{k,i}(\theta) L_k(s_k|\theta) \tag{15}$$

for any $k \in \mathcal{N}$. Let us denote the truth, namely, the true probability mass function (PMF) over the event set Θ, by $p(\theta)$ so that

$$p(\theta) = \delta_{\theta,\theta^\bullet} \tag{16}$$

where $\delta_{x,y}$ is the Kronecker delta function such that $\delta_{x,y}=1$ when $x=y$ and $\delta_{x,y}=0$ otherwise.

The structure of the diffusion strategy (12) relies on the use of updated beliefs. Therefore, some effort is needed to study the learning ability of the agents. In order to assess the performance of the diffusion learning strategy, we employ the expectation of the Kullback-Leibler divergence (KL-divergence) [22] from $p(\theta)$ to $\boldsymbol{\mu}_{k,i}(\theta)$ as a performance measure. Specifically, we define agent k's *regret* at time i as

$$Q(\boldsymbol{\mu}_{k,i}) \triangleq D(p||\boldsymbol{\mu}_{k,i}) \tag{17}$$

where $D(p||q)$ denotes the KL-divergence from p to q. Then, we define agent k's *risk* at time i as

$$J(\boldsymbol{\mu}_{k,i}) \triangleq \mathbb{E}_{\mathcal{F}_i} Q(\boldsymbol{\mu}_{k,i}) \tag{18}$$

where $\mathbb{E}_{\mathcal{F}_i}$ denotes the expectation with respect to $\mathbb{P}$ over the filtration $\mathcal{F}_i$ in (4). The overall performance at time i is defined as the weighted average risk across the network:

$$J(\boldsymbol{\mu}_i) \triangleq \sum_{k=1}^{N} y_k J(\boldsymbol{\mu}_{k,i}) \tag{19}$$

where y_k denotes the k-th element of y given in (10). We first establish the convergence of the overall average risk $J(\boldsymbol{\mu}_i)$.

Lemma 1: Assume that A is left-stochastic and primitive, and that there exists at least one agent with a positive prior belief about the true event θ^o. Then, the network risk $J(\boldsymbol{\mu}_i)$ converges as $i \to \infty$.

Proof: From (16) and (17), we get

$$Q(\boldsymbol{\mu}_{k,i}) = \sum_{\theta\in\Theta} p(\theta)\log\frac{p(\theta)}{\boldsymbol{\mu}_{k,i}(\theta)} = -\log \boldsymbol{\mu}_{k,i}(\theta^o) \tag{20}$$

where we adopt the convention $0\log 0=0$. Without loss of generality, we assume that agent ℓ^o has a positive prior belief $\mu_{\ell^\bullet,0}(\theta^o)>0$. By the nonnegativity of A and (14), the private belief about θ^o of any agent k with ℓ^o in its neighborhood will become positive in the next time instance $i=1$ because

$$\boldsymbol{\mu}_{k,1}(\theta^o) \geq \alpha_{\ell^o k}\frac{\mu_{\ell^o,0}(\theta^o)L_{\ell^o}(\boldsymbol{\xi}_{\ell^o,1}|\theta^o)}{m_{k,0}(\boldsymbol{\xi}_{\ell^o,1})} > 0 \tag{21}$$

Repeating this argument for a finite time, every agent in the network ends up with a positive belief about θ^o due to the primitivity of A. For large enough i, from (17)–(19) we get

$$J(\boldsymbol{\mu}_i) = \mathbb{E}_{\mathcal{F}_i}\sum_{k=1}^{N} y_k D(p||\boldsymbol{\mu}_{k,i}) \geq 0 \tag{22}$$

due to Gibbs' inequality [22], namely, $D(p||q)\geq 0$ for any two PMFs $\{p,q\}$ over Θ satisfying the absolute-continuity condition, $p(\theta)=0$ if $q(\theta)=0$ for any $\theta\in\Theta$. It is obvious that our p and $\boldsymbol{\mu}_{k,i}$ in (22) satisfy the absolute-continuity condition because of (16) and (21). Then, the risk of agent k at time i can be upper bounded by

$$\begin{aligned}
J(\boldsymbol{\mu}_{k,i}) &\overset{(a)}{=} -\mathbb{E}_{\mathcal{F}_i}\log\left(\boldsymbol{\mu}_{k,i}(\theta^o)\right)\\
&\overset{(b)}{=} -\mathbb{E}_{\mathcal{F}_i}\log\left(\sum_{\ell\in\mathcal{N}_k}\alpha_{\ell k}\frac{\boldsymbol{\mu}_{\ell,i-1}(\theta^o)L_\ell(\boldsymbol{\xi}_{\ell,i}|\theta^o)}{\boldsymbol{m}_{\ell,i-1}(\boldsymbol{\xi}_{\ell,i})}\right)\\
&\overset{(c)}{\leq} -\mathbb{E}_{\mathcal{F}_i}\sum_{\ell\in\mathcal{N}_k}\alpha_{\ell k}\log\left(\frac{\boldsymbol{\mu}_{\ell,i-1}(\theta^o)L_\ell(\boldsymbol{\xi}_{\ell,i}|\theta^o)}{\boldsymbol{m}_{\ell,i-1}(\boldsymbol{\xi}_{\ell,i})}\right)\\
&\overset{(d)}{=} \sum_{\ell\in\mathcal{N}_k}\alpha_{\ell k}\left[-\mathbb{E}_{\mathcal{F}_{i-1}}\log\boldsymbol{\mu}_{\ell,i-1}(\theta^o)\right]\\
&\quad - \mathbb{E}_{\mathcal{F}_{i-1}}\sum_{\ell\in\mathcal{N}_k}\alpha_{\ell k}\mathbb{E}_{\boldsymbol{\xi}_{\ell,i}}\left[\log\left(\frac{L_\ell(\boldsymbol{\xi}_{\ell,i}|\theta^o)}{\boldsymbol{m}_{\ell,i-1}(\boldsymbol{\xi}_{\ell,i})}\right)\Big|\mathcal{F}_{i-1}\right]\\
&\overset{(e)}{\leq} \sum_{\ell\in\mathcal{N}_k}\alpha_{\ell k}\,J(\boldsymbol{\mu}_{\ell,i-1})
\end{aligned} \tag{23}$$

where step (a) is by (18) and (20), step (b) by (14), step (c) by the convexity of $-\log(\cdot)$, step (d) by the conditional expectation, and step (e) by the nonnegativity of the KL-divergence from $L_\ell(s_\ell|\theta^o)$ to $\boldsymbol{m}_{\ell,i-1}(s_\ell)$. By (9), (19), and (23), we get

$$\begin{aligned}
J(\boldsymbol{\mu}_i) &\leq \sum_{k=1}^{N} y_k \sum_{\ell=1}^{N}\alpha_{\ell k}\,J(\boldsymbol{\mu}_{\ell,i-1})\\
&\overset{(a)}{=} \sum_{\ell=1}^{N} y_\ell\,J(\boldsymbol{\mu}_{\ell,i-1}) = J(\boldsymbol{\mu}_{i-1})
\end{aligned} \tag{24}$$

where step (a) is due to (10). From (22) and (24), we conclude that $\{J(\boldsymbol{\mu}_i)\}$ is a nonnegative, monotonically decreasing real sequence. By the monotone convergence theorem of real sequences [23], $\{J(\boldsymbol{\mu}_i)\}$ converges to a real number. ■

From Lemma 1, we arrive at the following result

$$\lim_{i\to\infty} J(\boldsymbol{\mu}_i) = \inf_{i\geq 0} -\mathbb{E}_{\mathcal{F}_i}\sum_{k=1}^{N} y_k \log\boldsymbol{\mu}_{k,i}(\theta^o) \tag{25}$$

We shall show in the sequel that the convergence of the overall average risk implies the asymptotic correctness of the forecast of the incoming signal.

Lemma 2: Under the same conditions of Lemma 1, agents develop correct forecasts of the incoming signals, namely,

$$\lim_{i\to\infty} \boldsymbol{m}_{k,i}(s_k) \overset{a.s.}{=} L_k(s_k|\theta^o) \tag{26}$$

for any $s_k\in\mathcal{S}_k$ and $k\in\mathcal{N}$, where $\overset{a.s.}{=}$ denotes almost surely convergence.

Proof: From steps (d) and (e) in (23) we get

$$\begin{aligned}
&\sum_{\ell\in\mathcal{N}_k}\alpha_{\ell k}J(\boldsymbol{\mu}_{\ell,i-1}) - J(\boldsymbol{\mu}_{k,i})\\
&\quad\geq \mathbb{E}_{\mathcal{F}_{i-1}}\sum_{\ell\in\mathcal{N}_k}\alpha_{\ell k}\,\mathbb{E}_{\boldsymbol{\xi}_{\ell,i}}\left[\log\frac{L_\ell(\boldsymbol{\xi}_{\ell,i}|\theta^o)}{\boldsymbol{m}_{\ell,i-1}(\boldsymbol{\xi}_{\ell,i})}\Big|\mathcal{F}_{i-1}\right]
\end{aligned} \tag{27}$$

Scaling by y_k, then summing over k on both sides of (27) and using (24) gives

$$J(\boldsymbol{\mu}_{i-1}) - J(\boldsymbol{\mu}_i) \geq \mathbb{E}_{\mathcal{F}_{i-1}}\sum_{\ell=1}^{N} y_\ell\mathbb{E}_{\boldsymbol{\xi}_{\ell,i}}\left[\log\frac{L_\ell(\boldsymbol{\xi}_{\ell,i}|\theta^o)}{\boldsymbol{m}_{\ell,i-1}(\boldsymbol{\xi}_{\ell,i})}\Big|\mathcal{F}_{i-1}\right] \tag{28}$$

By the Cauchy criterion for convergence [23], we get

$$0 = \lim_{i\to\infty} [J(\boldsymbol{\mu}_{i-1}) - J(\boldsymbol{\mu}_i)] \geq$$

$$\lim_{i\to\infty} \mathbb{E}_{\mathcal{F}_{i-1}} \sum_{\ell=1}^{N} y_\ell \mathbb{E}_{\boldsymbol{\xi}_{\ell,i}} \left[\log \frac{L_\ell(\boldsymbol{\xi}_{\ell,i}|\theta^o)}{\boldsymbol{m}_{\ell,i-1}(\boldsymbol{\xi}_{\ell,i})} \Big| \mathcal{F}_{i-1}\right] \geq 0 \quad (29)$$

where the last inequality is because $J(\boldsymbol{\mu}_i)$ is monotonically decreasing as shown in (24). Therefore, it holds that

$$\lim_{i\to\infty} \mathbb{E}_{\mathcal{F}_{i-1}} \sum_{\ell=1}^{N} y_\ell \mathbb{E}_{\boldsymbol{\xi}_{\ell,i}} \left[\log \frac{L_\ell(\boldsymbol{\xi}_{\ell,i}|\theta^o)}{\boldsymbol{m}_{\ell,i-1}(\boldsymbol{\xi}_{\ell,i})} \Big| \mathcal{F}_{i-1}\right] = 0 \quad (30)$$

By Gibbs' inequality, the conditional expectations in (30) are nonnegative random variables, i.e.,

$$\mathbb{E}_{\boldsymbol{\xi}_{\ell,i}} \left[\log \frac{L_\ell(\boldsymbol{\xi}_{\ell,i}|\theta^o)}{\boldsymbol{m}_{\ell,i-1}(\boldsymbol{\xi}_{\ell,i})} \Big| \mathcal{F}_{i-1}\right] \geq 0 \quad (31)$$

for any ℓ. Since $y_\ell > 0$, from (30) and (31), we obtain

$$\lim_{i\to\infty} \sum_{s_\ell \in \mathcal{S}_\ell} L_\ell(s_\ell|\theta^o) \log \frac{L_\ell(s_\ell|\theta^o)}{\boldsymbol{m}_{\ell,i-1}(s_\ell)} \overset{a.s.}{=} 0 \quad (32)$$

for any ℓ. According to Gibbs's inequality, equation (32) (almost surely) equals to zero if, and only if,

$$\lim_{i\to\infty} \boldsymbol{m}_{\ell,i-1}(s_\ell) \overset{a.s.}{=} L_\ell(s_\ell|\theta^o) \quad (33)$$

for any $s_\ell \in \mathcal{S}_\ell$ and ℓ. ∎

It is worth noting that the correct forecasting result (26) is guaranteed by requiring A to be a primitive matrix. This condition is satisfied if at least one diagonal entry of A is strictly positive. That is, we only require at least one agent to have a positive self-reliance α_{kk}. In contrast, Proposition 1 in [3] imposes the stronger condition that *all* agents must have positive self-reliances. The reason behind this difference is that even though some agent k may not integrate its own intermediate belief $\boldsymbol{\psi}_{k,i}$ into the new belief $\boldsymbol{\mu}_{k,i}$ by having $\alpha_{kk} = 0$, the new-information-bearing belief $\boldsymbol{\psi}_{k,i}$ will still be utilized by some other agent ℓ who has agent k in its neighborhood (i.e., $\alpha_{\ell k} > 0$). This fact explains why the diffusion strategy (11) can distribute information more thoroughly over the network than the consensus strategy (8).

In the following we shall show that as $i \to \infty$, each agent k will assign zero belief to any event $\theta \in \bar{\Theta}_k$. In order to proceed with the analysis, we assume that for each agent k, there exists at least one prevailing signal s_k^o such that

$$L_k(s_k^o|\theta^o) - L_k(s_k^o|\theta) \geq \delta_k^o > 0, \quad \forall\, \theta \in \bar{\Theta}_k \quad (34)$$

Lemma 3: Given the existence of prevailing signals, asymptotically correct forecasting implies that each agent k eventually identifies its "unlikely" event set $\bar{\Theta}_k$, namely,

$$\lim_{i\to\infty} \boldsymbol{\mu}_{k,i}(\theta) \overset{a.s.}{=} 0 \quad (35)$$

for any $\theta \in \bar{\Theta}_k$ and $k \in \mathcal{N}$.

Proof: Assuming $\bar{\Theta}_k \neq \varnothing$. From (26) we get

$$\lim_{i\to\infty} [L_k(s_k|\theta^o) - \boldsymbol{m}_{k,i-1}(s_k)] \overset{a.s.}{=} 0$$

$$\overset{(a)}{\Longrightarrow} \lim_{i\to\infty} \sum_{\theta\in\Theta} \boldsymbol{\mu}_{k,i-1}(\theta)\, [L_k(s_k|\theta^o) - L_k(s_k|\theta)] \overset{a.s.}{=} 0$$

$$\overset{(b)}{\Longrightarrow} \lim_{i\to\infty} \sum_{\theta\in\bar{\Theta}_k} \boldsymbol{\mu}_{k,i-1}(\theta) [L_k(s_k|\theta^o) - L_k(s_k|\theta)] \overset{a.s.}{=} 0 \quad (36)$$

for any $s_k \in \mathcal{S}_k$, where step (a) is by (15) and step (b) by (5). By (34), applying signal s_k^o to (36) yields

$$\lim_{i\to\infty} \sum_{\theta\in\bar{\Theta}_k} \boldsymbol{\mu}_{k,i-1}(\theta) [L_k(s_k^o|\theta^o) - L_k(s_k^o|\theta)] \overset{a.s.}{=} 0 \implies$$

$$0 \geq \lim_{i\to\infty} \delta_k^o \sum_{\theta\in\bar{\Theta}_k} \boldsymbol{\mu}_{k,i-1}(\theta) \geq 0, \quad a.s. \quad \implies$$

$$\lim_{i\to\infty} \sum_{\theta\in\bar{\Theta}_k} \boldsymbol{\mu}_{k,i-1}(\theta) \overset{a.s.}{=} 0 \implies \lim_{i\to\infty} \boldsymbol{\mu}_{k,i-1}(\theta) \overset{a.s.}{=} 0 \quad (37)$$

for any $\theta \in \bar{\Theta}_k$ and $k \in \mathcal{N}$. ∎

With Lemma 3, we can now proceed to show that all agents learn the truth asymptotically.

Lemma 4: Assume that (a) the combination matrix A is left-stochastic and primitive; (b) there exists at least one agent with positive prior belief about the true event; and (c) there exists at least one prevailing signal for each agent. Then, all agents in the social network learn the truth asymptotically, namely,

$$\lim_{i\to\infty} \boldsymbol{\mu}_{k,i}(\theta^o) \overset{a.s.}{=} 1 \quad (38)$$

for any $k \in \mathcal{N}$.

Proof: From (14) and (35), we obtain that for any $\theta \in \bar{\Theta}_k$,

$$\lim_{i\to\infty} \sum_{\ell\in\mathcal{N}_k} \alpha_{\ell k}\, \boldsymbol{\psi}_{\ell,i}(\theta) \overset{a.s.}{=} 0 \overset{(a)}{\Longrightarrow} \lim_{i\to\infty} \boldsymbol{\psi}_{\ell,i}(\theta) \overset{a.s.}{=} 0$$

$$\overset{(b)}{\Longrightarrow} \lim_{i\to\infty} \boldsymbol{\mu}_{\ell,i-1}(\theta) \overset{a.s.}{=} 0 \quad (39)$$

for any $\ell \in \mathcal{N}_k$ and $k \in \mathcal{N}$, where step (a) is because both $\alpha_{\ell k}$ and $\boldsymbol{\psi}_{\ell,i}$ are nonnegative and step (b) is by the absolute-continuity of $\boldsymbol{\psi}_{\ell,i}$ with respect to $\boldsymbol{\mu}_{\ell,i-1}$ because of the Bayesian update step in (12). Since the combination matrix A is primitive, propagating the argument in (39) by at most $N-1$ iterations will cross over every agent and enforce the following relation

$$\lim_{i\to\infty} \boldsymbol{\mu}_{\ell,i}(\theta) \overset{a.s.}{=} 0 \quad (40)$$

for any $\theta \in \bar{\Theta}_k$, $\ell \in \mathcal{N}$, and $k \in \mathcal{N}$. Repeating this argument for all agents in the network yields

$$\lim_{i\to\infty} \boldsymbol{\mu}_{k,i}(\theta) \overset{a.s.}{=} 0, \quad \forall\, \theta \in \bigcup_{\ell\in\mathcal{N}} \bar{\Theta}_\ell, \quad k \in \mathcal{N} \quad (41)$$

By (6), we obtain

$$\bigcup_{\ell\in\mathcal{N}} \bar{\Theta}_\ell = \overline{\bigcap_{\ell\in\mathcal{N}} \Theta_\ell} = \Theta \backslash \{\theta^o\} \quad (42)$$

From (41) and (42), we therefore arrive at

$$\lim_{i\to\infty} \boldsymbol{\mu}_{k,i}(\theta) \overset{a.s.}{=} 0, \quad \forall\, \theta \in \Theta\backslash\{\theta^o\}, \quad k \in \mathcal{N} \quad (43)$$

or, equivalently,

$$\lim_{i\to\infty} \boldsymbol{\mu}_{k,i}(\theta^o) \overset{a.s.}{=} 1, \quad \forall\, k \in \mathcal{N} \quad (44)$$

which completes the proof. ∎

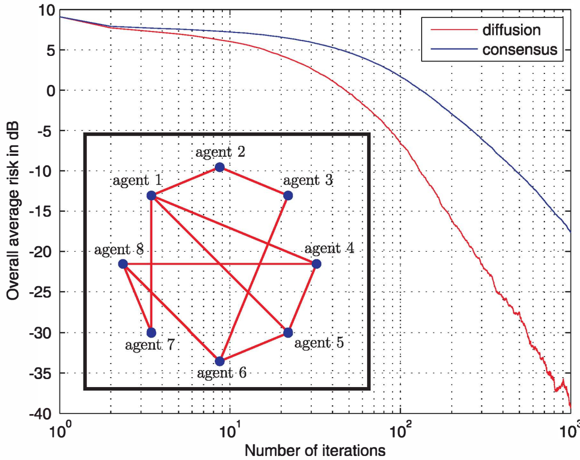

Fig. 1. Network topology and comparison of the learning curves: diffusion (12) vs. consensus (7) and (8).

V. Simulation Results

We illustrate the performance of the diffusion learning strategy by a simulation. For a fair comparison with the consensus-based learning strategy (7) and (8) from [3], we adopt a simulation setup similar to Example 2 in [3] and apply it to a more arbitrary topology rather than a ring. We assume that the social network consists of $N = 8$ agents interconnected via the topology shown in Fig. 1. The combination matrix A is formed by the Metropolis rule [24], i.e.,

$$\alpha_{\ell k} = \begin{cases} \dfrac{1}{\max\{n_k, n_\ell\}}, & \ell \in \mathcal{N}_k \backslash \{k\} \\ 1 - \sum_{\ell \in \mathcal{N}_k \backslash \{k\}} \alpha_{\ell k}, & \ell = k \end{cases} \tag{45}$$

We assume that there are $M = 9$ possible events, $\Theta = \{\theta^o, \theta_1, \theta_2, \ldots, \theta_8\}$, where θ^o is the true event of interest. Agents continuously observe binary signals randomly drawn from $\mathcal{S}_k = \{H, T\}$ over time $i \geq 1$. Signals for different agents are generated independently. The likelihood function for agent k is given by

$$L_k(s_k = H|\theta) = \begin{cases} \dfrac{k}{k+1}, & \text{if } \theta = \theta_k \\ \dfrac{1}{(k+1)^2}, & \text{otherwise} \end{cases} \tag{46}$$

We simulate both the diffusion social learning strategy (12) and the consensus social learning strategy (8) and plot the results in Fig. 1. The learning curves for the overall average risk $J(\boldsymbol{\mu}_i)$ are obtained by averaging over 100 independent experiments, or sample paths. Each sample path is obtained by feeding the social network a sequence of signals of length 1000. The prior beliefs are randomly generated by the uniform distribution $U(0, 1)$. From Fig. 1, we see that the diffusion strategy (12) converges faster than the consensus strategy (8).

VI. Conclusion

In this work we proposed a non-Bayesian social learning strategy of the diffusion-type given by (12). We showed that the diffusion strategy endows social networks with the ability to learn the truth. Simulation results indicate that the diffusion learning strategy has a faster convergence rate than the consensus learning strategy proposed in [3]. This is because the diffusion strategy allows information to diffuse through the network more thoroughly.

References

[1] C. P. Chamley, *Rational Herds*, Cambridge Univ. Press, Cambridge, UK, 2004.
[2] D. Acemoglu and A. Ozdaglar, "Opinion dynamics and learning in social networks," *Dyn. Games Appl.*, vol. 1, no. 1, pp. 3–49, Mar. 2011.
[3] A. Jadbabaie, P. Molavi, A. Sandroni, and A. Tahbaz-Salehi, "Non-Bayesian social learning," *Game. Econ. Behav.*, vol. 76, no. 1, pp. 210–225, Sept. 2012.
[4] A. V. Banerjee, "A simple model of herd behavior," *Quart. J. Econ.*, vol. 107, no. 3, pp. 797–817, Aug. 1992.
[5] L. Smith and Sorensen, "Pathological outcomes of observational learning," *Econometrica*, vol. 68, no. 2, pp. 371–398, Mar. 2000.
[6] D. Acemoglu, M. Dahleh, A. Ozdaglar, and A. Tahbaz-Salehi, "Observational learning in an uncertain world," in *Proc. IEEE Conf. Decision Control (CDC)*, Atlanta, GA, Dec. 2010, pp. 6645–6650.
[7] M. H. DeGroot, "Reaching a consensus," *J. Amer. Statist. Assoc.*, vol. 69, no. 345, pp. 118–121, Jan. 1974.
[8] G. Ellison and D. Fudenberg, "Rules of thumb for social learning," *J. Polit. Econ.*, vol. 101, no. 4, pp. 612–643, Aug. 1993.
[9] V. Bala and S. Goyal, "Learning from neighbours," *Rev. Econ. Stud.*, vol. 65, no. 3, pp. 595–621, July 1998.
[10] V. Bala and S. Goyal, "Conformism and diversity under social learning," *Econom. Theory*, vol. 17, no. 1, pp. 101–120, Jan. 2001.
[11] P. M. DeMarzo, D. Vayanos, and J. Zwiebel, "Persuasion bias, social influence, and unidimensional opinions," *Quart. J. Econ.*, vol. 118, no. 3, pp. 909–968, Aug. 2003.
[12] L. G. Epstein, J. Noor, and A. Sandroni, "Non-Bayesian updating: a theoretical framework," *Theoret. Econ.*, vol. 3, no. 2, pp. 193–229, June 2008.
[13] L. G. Epstein, J. Noor, and A. Sandroni, "Non-Bayesian learning," *B.E. J. Theor. Econ.*, vol. 10, no. 1, pp. 1–16, Jan. 2010.
[14] D. Acemoglu, A. Ozdaglar, and A. ParandehGheibi, "Spread of (mis)information in social networks," *Game. Econ. Behav.*, vol. 70, no. 2, pp. 194–227, Nov. 2010.
[15] B. Golub and M. O. Jackson, "Naive learning in social networks and the wisdom of crowds," *Am. Econ. J. Microecon.*, vol. 2, no. 1, pp. 112–149, Feb. 2010.
[16] R. M. Frongillo, G. Schoenebeck, and O. Tamuz, "Social learning in a changing world," in *Internet and Network Economics*, N. Chen, E. Elkind, and E. Koutsoupias, Eds., vol. 7090 of *Lecture Notes in Computer Science*, pp. 146–157. Springer, 2011.
[17] C. G. Lopes and A. H. Sayed, "Diffusion least-mean squares over adaptive networks: Formulation and performance analysis," *IEEE Trans. Signal Process.*, vol. 56, no. 7, pp. 3122–3136, July 2008.
[18] F. S. Cattivelli and A. H. Sayed, "Diffusion LMS strategies for distributed estimation," *IEEE Trans. Signal Process.*, vol. 58, no. 3, pp. 1035–1048, Mar. 2010.
[19] A. H. Sayed, "Diffusion adaptation over netowrks," to appear in *E-Reference Signal Processing*, R. Chellapa and S. Theodoridis, Eds. Elsevier, 2013. Also available online at http://arxiv.org/abs/1205.4220 [cs.MA], May 2012.
[20] S-Y. Tu and A. H. Sayed, "Diffusion strategies outperform consensus strategies for distributed estimation over adaptive netowrks," *IEEE Trans. Signal Process.*, vol. 60, no. 12, pp. 6217–6234, Dec. 2012.
[21] A. Berman and R. J. Plemmons, *Nonnegative Matrices in the Mathematical Sciences*, SIAM, PA, 1994.
[22] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, Wiley, NY, 1991.
[23] W. Rudin, *Principles of Mathematical Analysis*, McGraw-Hill, NY, 3rd edition, 1976.
[24] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, "Equations of state calculations by fast computing machines," *J. Chem. Phys.*, vol. 21, no. 6, pp. 1087–1092, 1953.