BAYESIAN SOCIAL LEARNING IN LINEAR NETWORKS OF AGENTS WITH RANDOM BEHAVIOR

Yunlong Wang and Petar M.Djurić

Department of Electrical and Computer Engineering
Stony Brook University, Stony Brook, NY 11794
Email: {yunlong.wang,petar.djuric@stonybrook.edu}

ABSTRACT

In this paper, we consider the problem of social learning in a network of agents where the agents make decisions on K hypotheses sequentially and broadcast their decisions to others. Each agent in the system has a private observation that is generated by one of the hypotheses. All the observations are independently generated from the same hypothesis. We study a setting where the agents randomly choose to make decisions prudently or non-prudently. A prudent decision is based on the private observation of the agent and all the previous decisions, whereas a non-prudent decision relies only on the private observation of the agent. We present a Bayesian learning method for the agents that exploits the information from other decisions. We analyze the asymptotical property of this system. A proof is presented that with the proposed decision policy, the posterior probability of the true hypothesis converges to one in probability. Simulation results are also provided.

Index Terms— social learning, Bayesian learning, prudent agents, non-prudent agents, random behavior

1. INTRODUCTION

In this paper, we study social learning in presence of random behavior. The addressed system is represented by a network of agents with line topology. Thus, we refer to the network as a linear network. Every agent receives an independent private observation generated by a time invariant true hypothesis. They make decisions sequentially one at a time in a predefined order. At every time instant, one agent chooses one of the K hypotheses, i.e., $\mathcal{H}_k$, where $k \in \{1, 2, \cdots, K\}$, and broadcasts the decision to all its successors. Each agent uses the Bayes' rule to formulate its belief in the hypotheses as a posterior distribution conditioned on the previously made decisions and on its private observation.

In recent years, various issues of social learning have been widely studied including the one that addresses the problems of modeling the interaction of agents in social networks. In some of these networks, the agents do not reveal their private observations to others. Instead, they only share their decisions [1]. After receiving the decisions, one agent can exploit the information in these decisions by using either non-Bayesian [2, 3] or Bayesian approaches [4, 5, 6, 7]. In this work, our study is within the realm of Bayesian learning. Some works on this issue can be found in [1, 8], where the interest is the study of herd behavior and information cascades. We recall that herding represents a setting where the agents blindly follow the previous decisions (i.e., they ignore their private information). In [5, 9], the authors provide sufficient conditions for asymptotic learning in Bayesian social learning systems. In [4], the problem of social learning with Bayesian game is addressed. More recently, a gossip method for Bayesian social learning is proposed in [10], with its convergence being shown in [11]. In [12], a learning method that allows all the agents to reach a consensus within finite number of iterations is addressed. An overview of models and techniques for studying social learning can be found in [7, 13].

In networks where agents employ Bayesian learning, if they choose to ignore their private observations they may herd on the wrong decision [1]. In [14], the author shows that the herding can be delayed if some of the agents behave benevolently. Benevolent behavior is defined as the behavior of an agent which makes its decision by optimizing the welfare reward of the society of agents. In that system, if the public belief is within certain region, the agent will reveal its full observation to others; otherwise it makes a decision by simply maximizing its personal utility. In [15], we propose a random decision making policy with which the herd behavior can be avoided.

Here, we address a different type of randomness to avoid herding. Namely, we introduce a decision making policy where each agent randomly chooses to behave prudently or non-prudently. In the former case, an agent decides based on its private observation and all the previous decisions in the network to maximize the expected value of its personal utility. In the latter case, an agent does the same but by using its private observations only. We present a Bayesian learning method for the agents to formulate beliefs in the hypotheses from the decision history and the private observations. We show that due to the presence of the non-prudent agents, herding cannot happen. We also show that the posterior probability of the true hypothesis will converge to one as the number of agents goes to infinity.

The paper is organized as follows. In the next section we describe the models of the sequential system and explain the social learning process. In Section 3, we present the proposed Bayesian learning method. The analysis of the convergence of the social belief and the probability of decision error are provided in Section 4. Simulation results are given in Section 5, and concluding remarks in Section 6.

2. PROBLEM STATEMENT

The process of social learning is illustrated in Fig. 1. We consider the decision making problem in linear networks of agents A_n, $n \in \mathbb{N}^+$, where each agent A_n receives an independent private observation y_n that is generated according to one of the following K hypotheses:

$$\mathcal{H}_k: \qquad y_n \sim \phi_k(y_n), \tag{1}$$

This work was supported by NSF under Award CCF-1320626.

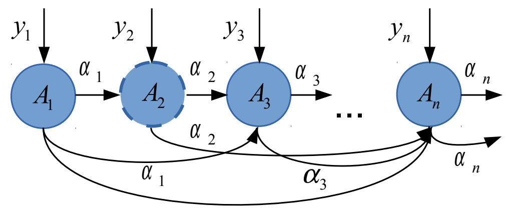

Fig. 1. Social learning in a sequential system. Agent A_2 ignores the decision of agent A_1 and makes its decision non-prudently.

where $k \in \mathcal{A} = \{1, 2, \cdots, K\}$. We assume that these K distributions are known by all the agents. Every agent has a uniform prior of the hypotheses, which means that $p(\mathcal{H}_k) = 1/K, \forall k \in \mathcal{A}$.

Each agent makes a once-in-life decision $\alpha_n \in \mathcal{A}$ about the hypothesis that generated the observation. The agents make these decisions in a predefined order. Once an agent makes a decision, it broadcasts it to all the subsequent agents. Let $\alpha_{1:n-1}$ denote the decision history from agent A_1 to agent A_{n-1}. Then the agent A_n can formulate its private belief on the hypothesis $\mathcal{H}_k$, $\beta_{k,n}$ as the posterior given by $p(\mathcal{H}_k|\alpha_{1:n-1}, y_n)$. By using the Bayes' rule, we have that $\forall n \in \mathbb{N}^+$ and $\forall k \in \mathcal{A}$,

$$\beta_{k,n} = p(\mathcal{H}_k|\alpha_{1:n-1}, y_n) \tag{2}$$

$$= \frac{\pi_{k,n-1}\phi_k(y_n)}{\sum_{i=1}^{K} \pi_{i,n-1}\phi_i(y_n)}, \tag{3}$$

where we define $\pi_{k,n}$ to be the social belief in $\mathcal{H}_k$ as a posterior on $\mathcal{H}_k$ conditioned on the decision history until A_n, i.e.,

$$\pi_{k,n} = p(\mathcal{H}_k|\alpha_{1:n}), \forall n \in \mathbb{N}^+, \tag{4}$$

with $\pi_{k,0}$ being initialized by the noninformative prior $\pi_{k,0} = 1/K$. Here we remark that the social belief serves as the prior knowledge for agent A_n before it has its private observation y_n.

For any agent A_n, after obtaining the private belief by the Bayes' rule in (2), we assume that it has two ways to make its decision. It can either choose to be non-prudent with probability ξ or be prudent with probability $1 - \xi$. Throughout this paper, it is assumed that ξ is known by all the agents. Let I_n be an indicator function that takes value one if A_n is non-prudent and zero otherwise; then it follows that $p(I_n = 1) = \xi$, $\forall n$. If A_n chooses to be non-prudent, it makes its decision by using its private information only. To that end, the agent A_n formulates the posterior on the hypotheses by y_n, and uses the maximum a posteriori (MAP) method to make its decision. Mathematically, we have

$$\alpha_n = \arg\max_k p\left(\mathcal{H}_k|y_n\right). \tag{5}$$

Here we remark that the decision is made non-prudently because it is independent to the decision history. Therefore, agent A_n provides information to the following agents from its observations only.

By contrast, a prudent agent makes a decision by maximizing its expected personal utility by using all the available information in the network. In this paper, if A_n is prudent, it makes a decision according to

$$\alpha_n = \arg\max_k p\left(\mathcal{H}_k|y_n, \alpha_{1:n-1}\right). \tag{6}$$

Thus, A_n chooses the hypothesis with the largest a posteriori probability conditioned on its observation and the decisions of the previous agents. If we set the reward of making a correct decision to be one and zero otherwise, by using this policy the expected reward is maximized.

In the following sections, we present the Bayesian learning method for agents to update their beliefs. We show that if the probability for each agent to be non-prudent is $\xi > 0$, then the social belief converges to the true hypothesis in probability, i.e., if $\mathcal{H}_k$ is the true hypothesis,

$$\lim_{n\to\infty} p(\pi_{k,n} = 1) = 1, \ \forall k \in \mathcal{A}. \tag{7}$$

3. THE BAYESIAN LEARNING

In this system, once agent A_n makes its decision α_n, all the following agents should update the social belief by Bayes' rule from $\pi_{k,n-1}$ to $\pi_{k,n}$. However, when the decision α_n is made, the following agents do not know whether this decision is made by a prudent or non-prudent agent. Therefore the hidden random variable I_n should be marginalized out when updating the social belief. By the definition in (4), given that $\alpha_n = i$, then $\forall k$, $i \in \mathcal{A}$, the social belief is updated according to

$$\pi_{k,n} = \frac{\pi_{k,n-1} l_{k,n}^{(i)}}{\sum_{j=1}^{K} \pi_{j,n-1} l_{j,n}^{(i)}}, \tag{8}$$

where the $l_{k,n}^{(i)}$ is defined by

$$l_{k,n}^{(i)} = p(\alpha_n = i|\alpha_{1:n-1}, \mathcal{H}_k) \tag{9}$$

$$= \sum_{I_n=0}^{1} p(\alpha_n = i|\alpha_{1:n-1}, \mathcal{H}_k, I_n)p(I_n). \tag{10}$$

It denotes the probability of agent A_n making decision $\alpha_n = i$ given the decision sequence up to α_{n-1} and the true state of nature being $\mathcal{H}_k$. As in [13], we refer to $l_{k,n}^{(i)}$ as *action likelihood* in the following sections. When $n = 1$, we set $l_{k,1}^{(i)} = p(\alpha_1 = i|\mathcal{H}_k)$.

As is shown in (10), the action likelihood is obtained by using the total probability theorem conditioned on two cases that $I_n = 0$ or $I_n = 1$. For both cases, the agents can obtain $p(\alpha_n = i|\alpha_{1:n-1}, \mathcal{H}_k, I_n)$ by marginalizing out y_n. Given the decision making policy in (5) and $I_n = 1$, we have

$$p(\alpha_n = i|\alpha_{1:n-1}, \mathcal{H}_k, I_n = 1) = \int_{y_n \in \mathrm{D}_i} \phi_k(y_n) dy_n, \tag{11}$$

where the set $\mathrm{D}_i = \{y_n|\phi_i(y_n) \geq \phi_j(y_n), \forall j \neq i\}$ is the decision region for a non-prudent agent to decide $\alpha_n = i$. Similarly, if $I_n = 0$, it follows that

$$p(\alpha_n = i|\alpha_{1:n-1}, \mathcal{H}_k, I_n = 0) = \int_{y_n \in \mathrm{S}_i} \phi_k(y_n) dy_n, \tag{12}$$

where the set $\mathrm{S}_i = \{y_n|\pi_{i,n-1}\phi_i(y_n) \geq \pi_{j,n-1}\phi_j(y_n), \forall j \neq i\}$ is the decision region for a non-prudent agent to decide $\alpha_n = i$. Here we remark that the sets S_i are defined by both the data model and the social belief $\pi_{k,n-1}$.

We summarize the behavior of each agent A_n as follows: When $t < n$, agent A_n calculates the action likelihood by (10), and updates its social belief by (8) after α_t is made. When $t = n$, A_n first randomly chooses one of the decision policies in (5) and (6), and then it makes a decision according to the chosen policy. When $t > n$, the agents become inactive.

4. THE ANALYSIS

In this section, we analyze the asymptotic property of the proposed Bayesian social learning system. By (4), $\pi_{k,n}(\alpha_{1:n}) = p(\mathcal{H}_k|\alpha_{1:n})$, shows that π_n is a function defined by the action sequence $\alpha_{1:n}$ (there is a deterministic mapping from $\alpha_{1:n}$ to π_n). Considering that $\alpha_{1:n}$ is random due to the randomness in the decision making and the observations , $\pi_{k,n}$ is a discrete random process. Our main result is about the expected value of the social belief denoted by $\mathbb{E}\,\pi_{k,n}$, which is presented by the following theorem:

Theorem 1 *In the proposed system, let $\mathcal{H}_k$ be the true state of nature. Then $\forall k \in \mathcal{A}$, the expected value of the social belief in $\mathcal{H}_k$ converges to one, i.e.,*

$$\lim_{n\to\infty} \mathbb{E}\,\pi_{k,n} \quad = \quad 1. \tag{13}$$

Proof: The proof is in the appendix. There, we first prove that the $\mathbb{E}\,\pi_{k,n} \geq \mathbb{E}\,\pi_{k,n-1}$. We then show that $\mathbb{E}\,\pi_{k,n} - \mathbb{E}\,\pi_{k,n-1} = 0$ if and only if $\mathbb{E}\,\pi_{k,n} = 1$. □

In the following analysis, let $\mathcal{H}_k$ be the true hypothesis. Then by Theorem 1, we immediately get that the social belief in the true hypothesis will converge to one in probability, which is given by the (7). Furthermore, considering that the prudent agents use the social belief as the prior for formulating the private belief, then we have that its private belief in the true hypothesis $\mathcal{H}_k$ also converges to one, given by

$$\lim_{n\to\infty} p(\beta_{k,n} = 1) = 1. \tag{14}$$

By the decision rule of the prudent agents in (6), the probability for a prudent agent making a correct decision converges to one, i.e.,

$$\lim_{n\to\infty} p(\alpha_n = k|I_n = 0) = 1. \tag{15}$$

Therefore, as n goes to infinity, the probability for an agent to make the correct decision is given by

$$\lim_{n\to\infty} p(\alpha_n = k) = (1-\xi) + b_k\xi, \tag{16}$$

where $b_k = p(\alpha_n = k|\mathcal{H}_k, I_n = 1)$ can be obtained by (11), and ξ is the probability that an agent in the network is non-prudent. Here we remark that this limit becomes smaller when ξ becomes larger. We also note that if $\xi = 0$, the system becomes identical to the systems in [1]. There, it is shown that the above asymptotical performances may not be achieved because of the herd behavior of the agents.

5. SIMULATIONS

In this section, we present simulation results on the evolution of social beliefs and the probability for right decisions. We also provide some numerical comparisons between different values of ξ.

In the experiment, the data of the kth hypothesis were generated by

$$\mathcal{H}_k : \qquad y_n \sim \mathrm{B}(m, p_k), \tag{17}$$

where $\mathrm{B}(m,p)$ represents a binomial distribution parameterized by m and p_k. We validated the analytical result of the expected social belief by Monte Carlo simulations, which were conducted with 2,000 trials. In each trial, we set the number of agents to be $N = 1,000$, and they observed data generated from the data model in (17) and parameterized by $m = 6$, $[p_1, p_2, p_3, p_4] = [0.5, 0.6, 0.7, 0.4]$.

We tested the system with values $\xi = 0$, $\xi = 0.02$, $\xi = 0.05$ and $\xi = 0.1$. The private signals of the agents were generated according to $\mathcal{H}_1$. When $\xi = 0$, all the agents were prudent. Then [1] shows that there is a positive probability that the agents herd on a wrong hypothesis. Once this herding starts, the social belief becomes unchanged and information cascade occurs.

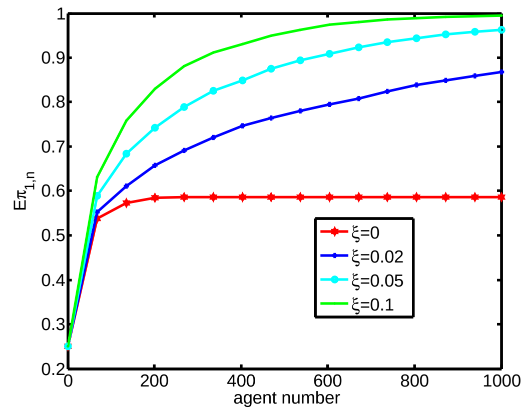

Fig. 2. The convergence of social beliefs with different ξs, where the agents get observations from the binomial data model.

We show the results of the proposed method in Fig. 2. On the abscissa, we plotted the agent index and on the ordinate the estimate of the expected social belief, which was given by the average social belief from all the 2,000 trials. From the figure, we see that the expected social beliefs kept increasing when $\xi > 0$, whereas it leveled off when $\xi = 0$ because of the herd behavior. With all the ξs which are greater than zero, it can be expected that $\mathbb{E}\,\pi_{1,n}$ converges to one if more and more agents are present in the network. We also observe that a larger ξ resulted in a larger convergence speed of the expected social belief.

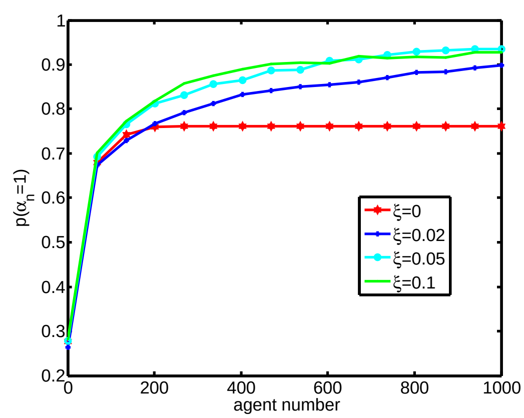

Fig. 3. The proportion of agents making the correct decision with a binomial data model.

Another performance metric of interest is the probability of correct decision. As we showed in (16), the limit of this probability is smaller when ξ is greater. The results of this investigation are plotted in Fig. 3, where the proportion of agents that made the correct decision is displayed. As shown in the figure, with $\xi > 0$, the probability for agents to make the correct decision keeps increasing. By (16), $p(\alpha_n = 1)$ converges to 0.9855, 0.9637, 0.9273 with ξ

equal to 0.02, 0.05, and 0.1 respectively. The figure indicates that larger ξs produce faster convergence rates.

Summarizing the results in the two figures, we conclude the following: Although the asymptotical probability for correct decision in a system with more non-prudent agents is smaller, future decision making agents can benefit more from them because the social belief has faster convergence to one.

6. CONCLUSIONS

In this paper we presented a Bayesian social learning system constituted by two types of agents. In this system, each agent randomly chooses to make its decision by using either a prudent or non-prudent policy. A prudent agent makes its decision by maximizing the expected personal utility based on all its available information. By contrast, a non-prudent makes a decision based on private observations only. We investigated a Bayeisan learning method for the agents that exploits the information from the previous decisions. In analyzing the system, we proved that the social belief in the true hypothesis converges to one in probability as long as not all the agents are prudent. We demonstrated the performance of the system by Monte Carlo simulations. We proved the convergence of the expected social belief and the probability of a correct decision. We also discussed the trade-off between the proportion of non-prudent agents in the network and the probability for correct decision by simulations.

7. APPENDIX

Before proving the Theorem, we first prove a lemma, which clarifies some properties of the proposed method.

Lemma 1 *In the proposed system, let $\mathcal{H}_k$ denote the true state of nature. If $0 < \pi_{1,n-1} < 1$, and if $\xi > 0$, then we have that $\forall i \in \mathcal{A}$, and $\forall \alpha_{1:n-1} \in \mathcal{A}^{n-1}$*

$$p(\alpha_n = i|\alpha_{1:n-1}, \mathcal{H}_k = \mathcal{H}_i) > p(\alpha_n = i|\alpha_{1:n-1}, \mathcal{H}_k \neq \mathcal{H}_i). \tag{18}$$

Proof: This lemma simply states that when a decision $\alpha_n = i$ is made, it is more likely that $\mathcal{H}_k = \mathcal{H}_i$ than $\mathcal{H}_k \neq \mathcal{H}_i$. To that end, by (10), we have that this statement can be shown by considering two cases. Without loss of generality, we assume that $\alpha_n = 1$.

First by [1], if the decision $\alpha_n = 1$ is made by a prudent agent and if $0 < \pi_{1,n-1} < 1$, we have that when $\alpha_n = 1$, the social belief in $\mathcal{H}_1$ is nondecreasing, i.e., $\pi_{1,n}(\alpha_{1:n-1}, 1) \geq \pi_{1,n-1}(\alpha_{1:n-1})$. Therefore from (8), this result implies that when $\alpha_n = 1$, the likelihood of $\mathcal{H}_k = \mathcal{H}_1$ is greater or equal to that of $\mathcal{H}_k \neq \mathcal{H}_1$. Moreover, the equal sign holds when information cascade emerges in the system.

Second, if the decision $\alpha_n = 1$ is made by a non-prudent agent, we can show that the likelihood of $\mathcal{H}_k = \mathcal{H}_1$ is strictly greater than $\mathcal{H}_k \neq \mathcal{H}_1$. Considering that the non-prudent decision is made by using the information in y_n only, then by (11), in the decision region D_1, $\phi_1(y_n) \geq \phi_j(y_n), \forall j \neq 1$. Thus, it follows that $\forall j \neq 1$,

$$\begin{aligned} p(\alpha_n &= 1|\alpha_{1:n-1}, \mathcal{H}_k = \mathcal{H}_1, I_n = 1) \geq \\ & p(\alpha_n = 1|\alpha_{1:n-1}, \mathcal{H}_k = \mathcal{H}_j, I_n = 1). \end{aligned} \tag{19}$$

Noting that the equal sign could not hold for each of the remaining $K-1$ hypotheses, then we must have that

$$\begin{aligned} p(\alpha_n &= 1|\alpha_{1:n-1}, \mathcal{H}_k = \mathcal{H}_1, I_n = 1) > \\ & p(\alpha_n = 1|\alpha_{1:n-1}, \mathcal{H}_k \neq \mathcal{H}_1, I_n = 1). \end{aligned} \tag{20}$$

With the analysis of the two cases, by marginalizing I_n, we have that (18) holds when $\xi > 0$. □

Now, by using this lemma, we prove Theorem 1.

Proof of Theorem 1: Without loss of generality, we assume that the true hypothesis $\mathcal{H}_k = \mathcal{H}_1$. Then we have that

$$\begin{aligned} \mathbb{E}\,\pi_{1,n} &= \sum_{\alpha_{1:n} \in \mathcal{A}^n} \pi_{1,n}\, p(\alpha_{1:n}) \\ &= \sum_{\alpha_{1:n-1} \in \mathcal{A}^{n-1}} p(\alpha_{1:n-1}) \left(\sum_{\alpha_n \in \mathcal{A}} p(\alpha_n|\alpha_{1:n-1}, \mathcal{H}_1)\pi_{1,n} \right). \end{aligned} \tag{21}$$

Next we show that $\mathbb{E}\,\pi_{s,n}$ is nondecreasing. By (21), we can write

$$\mathbb{E}\,\pi_{1,n} - \mathbb{E}\,\pi_{1,n-1} = \sum_{\alpha_{1:n-1} \in \mathcal{A}^{n-1}} p(\alpha_{1:n-1})\Delta(\alpha_{1:n}), \tag{22}$$

where $\Delta(\alpha_{1:n})$ is a function of the action sequence $\alpha_{1:n}$ given by,

$$\Delta(\alpha_{1:n}) = \sum_{\alpha_n \in \mathcal{A}} p(\alpha_n|\alpha_{1:n-1}, \mathcal{H}_1)\pi_{1,n} - \pi_{1,n-1}. \tag{23}$$

By (8), we can further expand $\Delta(\alpha_{1:n})$ by

$$\Delta(\alpha_{1:n}) = \pi_{1,n-1} \left(\sum_{i=1}^{K} \frac{l_{1,n}^{(i)}}{\sum_{j=1}^{K} \pi_{j,n-1} l_{j,n}^{(i)}} l_{1,n}^{(i)} - 1 \right). \tag{24}$$

Noting that $\sum_{i=1}^{K} l_{1,n}^{(i)} = 1$ and $\sum_{i=1}^{K} (\sum_{j=1}^{K} \pi_{j,n-1} l_{j,n}^{(i)}) = 1$, $\Delta(\alpha_{1:n})$ can be reformulated as follows:

$$\Delta(\alpha_{1:n}) = \pi_{1,n-1} \sum_{i=1}^{K} \frac{(l_{1,n}^{(i)} - \sum_{j=1}^{K} \pi_{j,n-1} l_{j,n}^{(i)})^2}{\sum_{j=1}^{K} \pi_{j,n-1} l_{j,n}^{(i)}}, \tag{25}$$

which shows that $\Delta(\alpha_{1:n}) \geq 0$. Because $p(\alpha_{1:n-1}) \geq 0$, $\forall \alpha_{1:n-1} \in \mathcal{A}^{n-1}$, we have proved that the expected social belief in $\mathcal{H}_1$ is nondecreasing. Noting that $0 \leq \mathbb{E}\,\pi_{1,n} \leq 1$ is bounded, then it must converge.

Next, we show that this limit is one. By the boundedness of $\mathbb{E}\,\pi_{1,n}$, we have that the limit of (22) must be zero, i.e.,

$$\lim_{n\to\infty} (\mathbb{E}\,\pi_{1,n} - \mathbb{E}\,\pi_{1,n-1}) = 0, \tag{26}$$

which implies that $\forall \alpha_{1:n-1} \in \mathcal{A}^{n-1}$,

$$\lim_{n\to\infty} p(\alpha_{1:n-1})\Delta(\alpha_{1:n-1}) = 0. \tag{27}$$

By Lemma 1, we have that if $0 < \pi_{1,n-1} < 1$,

$$p(\alpha_n = 1|\alpha_{1:n-1}, \mathcal{H}_k = \mathcal{H}_1) > \frac{\sum_{j\neq 1} \pi_{j,n-1} l_{j,n}^{(i)}}{1 - \pi_{1,n-1}}. \tag{28}$$

Then it follows that $(1-\pi_{1,n-1}) l_{1,n}^{(1)} - \sum_{j\neq 1} \pi_{j,n-1} l_{j,n}^{(i)} > 0$, which shows that $(l_{1,n}^{(1)} - \sum_{j=1}^{K} \pi_{j,n-1} l_{j,n}^{(1)})^2 > 0$ unless $\pi_{1,n-1} = 1$ or $\pi_{1,n-1} = 0$. Therefore, $\Delta(\alpha_{1:n}) = 0$ if and only if $\pi_{1,n-1} = 1$ or $\pi_{1,n-1} = 0$. By [1], the Bayesian learning cannot be totally wrong. Namely, if $\mathcal{H}_1$ is true, the probability that $\pi_{1,n-1} = 0$ converges to zero. Then from (27), we get that

$$\lim_{n\to\infty} \sum_{\alpha_{1:n-1} \in \mathcal{A}^{n-1}} p(\alpha_{1:n-1})(1 - \pi_{1,n-1}) = 0. \tag{29}$$

Then by (21), we show that the limit $\mathbb{E}\,\pi_n$ is one, given by

$$\lim_{n\to\infty} \mathbb{E}\,\pi_{1,n} = 1. \tag{30}$$

□

8. REFERENCES

[1] C. Chamley, *Rational Herds: Economic Models of Social Learning*, Cambridge University Press, 2004.

[2] A. Jadbabaie, P. Molavi, A. Sandroni, and A. Tahbaz-Salehi, "Non-Bayesian social learning," *Games and Economic Behavior*, vol. 76, no. 1, pp. 210–225, 2012.

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

[4] V. Krishnamurthy and H. V. Poor, "Social learning and Bayesian games in multiagent signal processing: How do local and global decision makers interact?," *IEEE Signal Processing Magazine*, vol. 30, no. 3, pp. 43–57, 2013.

[5] D. Acemoglu, M. A. Dahleh, I. Lobel, and A. Ozdaglar, "Bayesian learning in social networks," *The Review of Economic Studies*, vol. 78, no. 4, pp. 1201–1236, 2011.

[6] D. Gale and S. Kariv, "Bayesian learning in social networks," *Games and Economic Behavior*, vol. 45, no. 2, pp. 329–346, 2003.

[7] P. M. Djurić and Y. Wang, "Distributed Bayesian learning in multiagent systems: Improving our understanding of its capabilities and limitations," *IEEE Signal Processing Magazine*, vol. 29, no. 2, pp. 65–76, 2012.

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

[9] E. Mossel, A. Sly, and O. Tamuz, "Asymptotic learning in Bayesian social networks," *Probability Theory and Related Fields*, vol. 1, no. 1, pp. 1–31, 2012.

[10] Y. Wang and P. M. Djurić, "Reaching consensus on a binary state by exchanging binary actions," in *Acoustics, Speech and Signal Processing, 2012 IEEE International Conference on*, 2012, pp. 3297–3300.

[11] Y. Wang and P. M. Djurić, "A gossip method for optimal consensus on a binary state from binary actions," *IEEE Journal of Selected Topics in Signal Processing*, vol. 7, no. 2, pp. 274–283, 2013.

[12] Z. Weng and P. M. Djurić, "Efficient learning by consensus over regular networks," in *Acoustics, Speech and Signal Processing, IEEE International Conference on*, 2014, pp. 7253–7257.

[13] V. Krishnamurthy and H. V. Poor, "A tutorial on interactive sensing in social networks," *IEEE Transactions on Computational Social Systems*, pp. 3–21, 2014.

[14] V. Krishnamurthy, "Bayesian sequential detection with phase-distributed change time and nonlinear penalty pomdp lattice programming approach," *IEEE Transactions on Information Theory*, vol. 57, no. 10, pp. 7096–7124, 2011.

[15] Y. Wang and P. M. Djurić, "Sequential Bayesian learning in linear networks with random decision making," in *Acoustics, Speech and Signal Processing, 2014 IEEE International Conference on*, 2014, pp. 6404–6408.