BAYESIAN LEARNING IN A NETWORK WITH MULTI-HYPOTHESIS DECISION EXCHANGES

Lingqing Gan,[†] *Yunlong Wang,*[*] *and Petar M. Djurić* [†]

[†]Department of ECE, Stony Brook University, Stony Brook, NY 11794
[*]IMS Health Inc., Plymouth Meeting, PA 19462

ABSTRACT

Opinion dynamics and its understanding in social networks is an emerging field of research in recent years. Existing work mainly considers direct exchanges of opinions among agents under certain conditions. This paper addresses a problem where the agents of a network make and exchange decisions repeatedly in a multi-hypothesis scenario and learn from the neighbors' decisions. Two models are proposed where the agents of the network use quasi-Bayesian learning to extract information about the true hypothesis from the neighbors' decisions. Theoretical analysis is provided about the conditions of a setting when agents become stubborn, that is, when they do not change their opinions anymore. We have run computer simulations to demonstrate the asymptotical properties of the proposed models. With our simulations we also show that, under one of the models, the agents of the network reach a consensus, and under the other, they form clusters.

Index Terms— Bayesian learning, Opinion dynamics, Decision exchanges, DeGroot model

1. INTRODUCTION

With the popularization of online social networks such as Facebook and LinkedIn, there are many questions of interest about them. One of these questions is how the opinions in social networks are formed and how they spread. More specifically, in a network of social agents described by a connected graph where each node represents one social agent, the aim is to model the learning of agents from their neighbors by way of exchanging decisions in a multi-hypothesis setting. Furthermore, it is of interest to understand how the opinions of the agents in the network change with time [1, 2, 3, 4].

In this context, efforts have been made in different fields, including economy [5, 6], sociology [1, 7], engineering [8] and physics [9]. In [5, 6], the authors consider social learning problems by Bayesian and non-Bayesian models, respectively. In these papers, the authors also provide proofs of the asymptotic optimality of the proposed models. In [1], the Krause model for addressing bounded confidence was introduced. According to the model, each agent only exchanges opinions with those agents whose opinions do not differ too much from its own opinion. A comprehensive review of opinion dynamics can be found in [10].

In social networks, individuals cannot always exchange beliefs and opinions directly. In addressing this, one category of models assumes that only decisions can be exchanged among neighboring agents [11, 12, 13, 14]. Then the agents make inference on the true state of nature from the available information and make their decisions by maximizing their utility. In [11, 12], the authors propose that once a decision is observed, the learner adjusts its log opinion ratio by adding or subtracting a fixed small value. Recently, in [13], the authors consider social learning in a network with line topology and where the agents make random decisions and repeatedly update their opinions by relying on Bayes' rule. To alleviate the computational complexity of the Bayesian belief update in networks with general topology, in [15] an approximated Bayesian learning method is proposed.

In this paper, we consider a general network of N agents that decide on one of K hypotheses. At the beginning, each agent obtains a private signal from the true hypothesis, which is the same for all the agents in the network. Subsequently, in every time slot, each agent exchanges information with its neighbors and makes its own decision. Motivated by DeGroot's model in linear opinion pooling [16], we propose two discrete time models where the agents fuse the information from the decisions of their neighbors by using Bayes' rule. Ideally, rational agents in a social network would apply Bayes' rule successively based on their private signals and the decisions of their neighbors. However, as the agents are unaware of the global network structure, such repeated applications of Bayes rule in networks become very complex due to the increasing amount of hidden variables [17, 18]. In contrast to a fully Bayesian model, this paper proposes two alternative models for approximating the posterior beliefs of the agents and based on the history of the decisions of their neighbors. By simulations, we show that by Model 1, it is likely that a big portion of the agents in the system will choose the true hypothesis. For Model 2, we show that there is a very large probability that all the decisions of the agents converge to the true hypothesis.

This paper is organized as follows: in the next section we describe the studied system and explain the agents' social learning processes in the network. In Section 3, we present two models of Bayesian learning. An analytical result related to Model 1 is provided in Section 4. Simulation results are given in Section 5, and concluding remarks are made in Section 6.

2. PROBLEM FORMULATION

Consider a network of N agents A_i, $i \in \mathcal{N}_A = \{1, 2, ..., N\}$. The agents do not know the topology of the network but they know their neighbors. The agents are also capable of performing local computations and of exchanging decisions with their neighbors. Suppose there are K possible hypotheses, listed as $\mathcal{H}_0, \mathcal{H}_1, \cdots, \mathcal{H}_{K-1}$. Under each hypothesis, a random categorical

This work has been supported by NSF under Awards CCF-1320626 and CCF-1618999.

signal y is generated, where $y \in \mathcal{K} = \{0, 1, \cdots, K-1\}$. The different hypotheses have their unique distributions, that is, there is a total of K different distributions each with the same support $\{0, 1, \cdots, K-1\}$. Each agent A_i receives a random signal y_i as its private observation from the same distribution as do the other agents in the network. The probability distribution of y_i under $\mathcal{H}_k$ is $\phi_k(y_i)$, i.e.,

$$\mathcal{H}_k \quad : \quad y_i \sim \phi_k(y_i). \tag{1}$$

In this model, we further assume that the distribution of the signals has the form

$$\phi_k(y) = \begin{cases} P, & \text{if } y = k, \\ \frac{1-P}{K-1}, & \text{otherwise,} \end{cases} \tag{2}$$

where P is a constant in the range $(0, 1)$. We also assume that the probability distributions of the signals under all the hypotheses are known to the agents.

At each time slot t, every agent A_i maintains its private belief according to the Bayes' rule. We denote the belief by the vector $\mathbf{B}_i^{(t)} = [\beta_{i,0}^{(t)}, \beta_{i,1}^{(t)}, \cdots, \beta_{i,K-1}^{(t)}]^\top$. The vector element $\beta_{i,k}^{(t)}$ denotes the posterior probability of the i^{th} agent about $\mathcal{H}_k$ at time instant t. Thus, the equation $\sum_{k=0}^{K-1} \beta_{i,k}^{(t)} = 1$ would always hold $\forall i, t$. Each posterior probability is given by

$$\beta_{i,k}^{(t)} = p(\mathcal{H}_k | y_i, \mathcal{I}_i^{(t)}), \tag{3}$$

$$= \frac{\pi_{i,k}^{(t)} p(y_i|\mathcal{H}_k)}{\sum_{l=0}^{K-1} \pi_{i,l}^{(t)} p(y_i|\mathcal{H}_l)}, \tag{4}$$

where $\pi_{i,k}^{(t)}$, refers to the *social belief* formed by A_i on $\mathcal{H}_k$ based on all of the information available to it, $\mathcal{I}_i^{(t)}$, other than y_i. We use the vector $\mathbf{\Pi}_i^{(t)} = [\pi_{i,0}^{(t)}, \pi_{i,1}^{(t)}, \cdots, \pi_{i,K-1}^{(t)}]^\top$ to represent the social belief that A_i holds at time t. It is worth noting that $\mathbf{\Pi}_i^{(t)} \in \mathbf{U} = \{\mathbf{\Pi}_i^{(t)} \mid \sum_{l=0}^{K-1} \pi_{i,l}^{(t)} = 1, \pi_{i,k}^{(t)} \in [0,1], \forall k \in \mathcal{K}\}$ holds true $\forall i \in \mathcal{N}_A$, $\forall t \in \mathbb{N}^+$. This $\mathbf{\Pi}_i^{(t)}$ serves as the prior distribution, and it is time varying as more and more decisions of the neighbors of A_i become known to A_i. At the initial time slot, $\mathbf{\Pi}_i^{(1)} = [1/K, 1/K, \cdots, 1/K]^\top$, $\forall i \in \mathcal{N}_A$. The likelihood, by definition, can be obtained by

$$p(y_i|\mathcal{H}_k) = \phi_k(y_i). \tag{5}$$

The A_i's decision on choosing one of the hypotheses is made according to

$$\alpha_i^{(t+1)} = \arg\max_{k \in \mathcal{K}} p(\mathcal{H}_k | y_i, \mathcal{I}_i^{(t)}). \tag{6}$$

In the next section, we present agents' Bayesian learning models for updating their respective social beliefs at each time slot.

3. MODELS OF LEARNING

Before we proceed to the proposed models, we first sketch DeGroot's model from [16]. At time instant 1, agent A_i in the multi-agent system starts with a belief $\zeta_i^{(1)} \in [0,1]$, $\forall i \in \mathcal{N}_A$ about a certain hypothesis. At the first time slot, the beliefs of all the agents in the system are represented by a vector $\mathbf{Z}^{(1)} = [\zeta_1^{(1)}, \zeta_2^{(1)}, \cdots, \zeta_N^{(1)}]^\top$.

Meanwhile, the agents' trust in each other can be represented by an $N \times N$ stochastic matrix $\mathbf{P}$, where the element $p_{i,j} \geq 0$ represents the belief that A_i has in agent A_j. It is assumed that $\sum_{j=1}^{N} p_{i,j} = 1$. The matrix $\mathbf{P}$ is not necessarily doubly stochastic. Then, at time slot t $(t > 1)$, all the agents update their beliefs according to

$$\mathbf{Z}^{(t+1)} = \mathbf{P}\mathbf{Z}^{(t)}. \tag{7}$$

In this paper, we propose a modified DeGroot model in a network where the agents learn from the decisions of their neighbors by the Bayesian method in a multi-hypothesis setting. More precisely, instead of linearly combining the belief of its neighbors, agent A_i adjusts its belief in $\mathcal{H}_k$ based on $\alpha_j^{(t)}$ by Bayes' rule, as

$$\pi_{i,k}^{(t+1)} = \frac{\pi_{i,k}^{(t)} \prod_{j \in \mathcal{N}_i} p(\alpha_j^{(t)} \mid \mathcal{H}_k)}{\sum_{l=0}^{K-1} [\pi_{i,l}^{(t)} \prod_{j \in \mathcal{N}_i} p(\alpha_j^{(t)} \mid \mathcal{H}_l)]}, \tag{8}$$

where $\mathcal{N}_i$ denotes the neighbors of agent A_i. The "action likelihood", $p(\alpha_j^{(t)} \mid \mathcal{H}_k)$ denotes the probability of A_j making decision $\alpha_j^{(t)}$ given that $\mathcal{H}_k$ is true. We remark that, to avoid sudden drop of social belief, A_i would ignore $p(\alpha_j^{(t)} \mid \mathcal{H}_k)$ if it equals zero.

According to (3) and (6), $\alpha_j^{(t)} = k$ if and only if

$$\pi_{j,k}^{(t)} p(y_j \mid \mathcal{H}_k) \geq \pi_{j,l}^{(t)} p(y_j \mid \mathcal{H}_l) \tag{9}$$

holds $\forall l \in \mathcal{K}$. The action likelihood can be written as,

$$p(\alpha_j = k \mid \mathcal{H}_m) = \int_{y_j \in S_k} \phi_m(y_j) dy_j, \tag{10}$$

where the set $S_k = \{y_j \mid \pi_{j,k}^{(t)} p(y_j \mid \mathcal{H}_k) \geq \pi_{j,l}^{(t)} p(y_j \mid \mathcal{H}_l), \forall l \in \mathcal{K}\}$ is the decision region of making decision k by agent j.

In computing $p(\alpha_j^{(t)} \mid \mathcal{H}_k)$, agent A_i must know its neighbor's current social belief vector $\mathbf{\Pi}_j^{(t)}$. However, $\mathbf{\Pi}_j^{(t)}$ is unknown to A_i and consequently, A_i has to estimate it. We propose two models that we use for estimating $\mathbf{\Pi}_j^{(t)}$.

Model 1: The agent A_i considers $\mathbf{\Pi}_j^{(t)}$ to be a point estimate that is equal to its own social belief $\mathbf{\Pi}_i^{(t)}$. That is to say, in the inference procedure of A_i, $\pi_{j,k}^{(t)} = \pi_{i,k}^{(t)}$, i.e., $p(\pi_{j,k}^{(t)}) = \delta(\pi_{j,k}^{(t)} - \pi_{i,k}^{(t)})$ is true $\forall k \in \mathcal{K}$. If A_j is a neighbor of A_i, it is convenient and intuitive for A_i to assume that A_j has the same social belief. Such methodology is useful for agents with limited computation ability.

Model 2: The agent A_i draws $\mathbf{\Pi}_j^{(t)} = [\pi_{j,0}^{(t)}, \pi_{j,1}^{(t)}, \cdots, \pi_{j,K-1}^{(t)}]^\top$ from a Dirichlet distribution, i.e., $\mathbf{\Pi}_j^{(t)} \sim \mathrm{Dir}(\mathbf{a}_j^{(t)})$ given by

$$p(\pi_{j,0}^{(t)}, \pi_{j,1}^{(t)}, \cdots, \pi_{j,K-1}^{(t)} | \mathbf{a}_j^{(t)}) = \frac{\Gamma\left(\sum_{k=0}^{K-1} a_{j,k}^{(t)}\right)}{\prod_{k=0}^{K-1} \Gamma\left(a_{j,k}^{(t)}\right)} \times \prod_{k=0}^{K-1} \pi_{j,k}^{a_{j,k}^{(t)} - 1} \tag{11}$$

with a parameter $\mathbf{a}_j^{(t)} = [a_{j,0}^{(t)}, a_{j,1}^{(t)}, \cdots, a_{j,K-1}^{(t)}]^\top$ satisfying

$$a_{j,k}^{(t)} = n_{j,k}^{(t)} + 1 \tag{12}$$

for all $k \in \mathcal{K}$, where $n_{j,k}^{(t)}$ is the number of times A_i observes that A_j has made a decision k up until t. To explain the motivation of such estimation, suppose we have a Multinomial random vector $\mathbf{S} \sim \text{Multi}(n, \mathbf{p})$ that is an outcome of n trials with K possible mutually exclusive outcomes with corresponding probabilities $\mathbf{p} = [p_1, p_2, \cdots, p_K]^\top$, which are unknown. If we estimate $\mathbf{p}$ with a Dirichlet prior with parameters $\mathbf{a}_0 = [1, 1, \cdots, 1]_{1\times K}$, being a uniform prior of $\mathbf{p}$ on its domain, the posterior $\hat{\mathbf{p}}$ also follows a Dirichlet distribution but with parameters $\mathbf{a}_n$, which is a vector whose k^{th} element a_k is equal to one plus the number of occurrences of the k^{th} outcome.

With both models, A_i computes $p(\alpha_j^{(t)} \mid \mathcal{H}_k)$ for all $k \in \mathcal{K}$, according to

$$p(\alpha_j^{(t)} \mid \mathcal{H}_k) = \int_{\mathbf{\Pi}_j \in \mathbf{U}} p(\alpha_j^{(t)} | \mathcal{H}_k, \mathbf{\Pi}_j^{(t)})\, p(\mathbf{\Pi}_j^{(t)})\, \mathrm{d}\mathbf{\Pi}_j^{(t)}. \tag{13}$$

With model one, (13) simplifies to

$$p(\alpha_j^{(t)} \mid \mathcal{H}_k) \quad = \quad Pr(\alpha_j^{(t)} | \mathcal{H}_k, \mathbf{\Pi}_j^{(t)} = \mathbf{\Pi}_i^{(t)}). \tag{14}$$

With model two, (13) turns into

$$\begin{aligned} &p(\alpha_j^{(t)} \mid \mathcal{H}_k) \qquad (15)\\ &= \int_{\mathbf{U}} \left(\int_{-\infty}^{+\infty} p(\alpha_j^{(t)} \mid y_j, \mathbf{\Pi}_j^{(t)}) p(y_j \mid \mathcal{H}_k) \mathrm{d}y_j \right) p(\mathbf{\Pi}_j^{(t)}) \mathrm{d}\mathbf{\Pi}_j^{(t)} \\ &= \int_{\mathbf{U}} \left(\int_{S_{\alpha_j^{(t)}}} \phi_k(y_j) \mathrm{d}y_j \right) p(\mathbf{\Pi}_j^{(t)}) \mathrm{d}\mathbf{\Pi}_j^{(t)}. \end{aligned}$$

4. ANALYSIS OF MODEL 1

In this section, we show that for Model 1, if the ratio of the largest element and the second largest element in an agent's social belief becomes greater than a given threshold, this agent will become a stubborn agent. In other words, irrespectively of the neighbors' decisions, this agent will keep its social belief unchanged forever.

Proposition 1 *In the proposed Model 1, let $\pi_{i,m_1}^{(t)}$ be the largest element in vector $\Pi_i^{(t)}$, and let $\pi_{i,m_2}^{(t)}$ be the second largest element. If $\frac{\pi_{i,m_1}^{(t)}}{\pi_{i,m_2}^{(t)}} > \frac{P(1-N)}{1-P}$, then $\Pi_i^{(\tau)} = \Pi_i^{(t)}, \forall \tau > t$.*

Proof: From (6), we can have that if $\frac{\pi_{i,m_1}^{(t)}}{\pi_{i,m_2}^{(t)}} > \frac{P(1-N)}{1-P}$ becomes true, A_i will choose $\mathcal{H}_{m_1}$ regardless of y_i. Thus, the A_i's estimation of its neighbor A_j will use

$$p(\alpha_j^{(t)} = m \mid \Pi_j^{(t)} = \Pi_i^{(t)}) = \begin{cases} 1 & \text{if } m = m_1, \\ 0 & \text{otherwise.} \end{cases} \tag{16}$$

If $\alpha_j^{(t)} \neq m_1$, A_i will ignore $p(\alpha_j^{(t)} \mid \Pi_i^{(t)})$, since $p(\alpha_j^{(t)} \mid \Pi_i^{(t)}) = 0$; if $\alpha_j^{(t)} = m_1$, because $p(\alpha_j^{(t)} \mid \Pi_i^{(t)}) = 1$. Then A_i's updated social belief is $\Pi_i^{(t+1)} = \Pi_i^{(t)}$. □

With this proposition, we can argue that the system evolves into a stable state (no changes in decisions) once all the agents in it become stubborn. Though all the agents may not become stubborn, the number of stubborn agents in the network will not decrease with time. It is likely, as shown in the following simulation, that the majority of agents become stubborn in finite time.

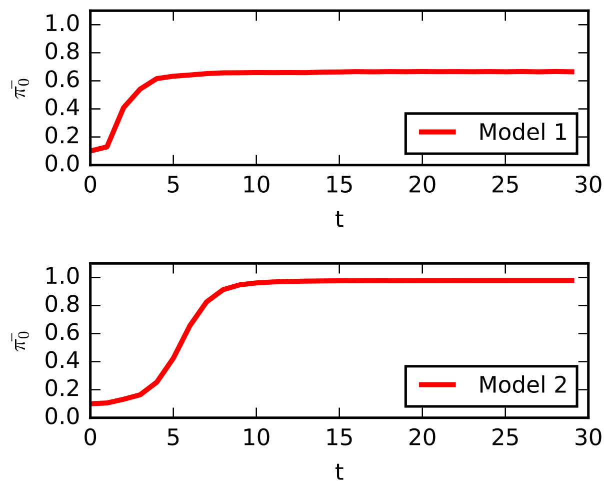

Fig. 1. Evolution of the average of social belief of all agents over time. The top subplot was obtained by running Model 1, with a total number of agents $N = 500$, number of hypotheses $K = 10$, probability of correct hypothesis $P = 0.25$, and number of iterations $T = 30$. The bottom subplot was run with Model 2, with the same parameters as Model 1. In both models the average social belief is nondecreasing with time.

5. SIMULATION RESULTS

In this section, we provide several simulation experiments to demonstrate the properties of the proposed models. In all experiments, a total number of N agents were modeled on a random Erdős-Rényi network of which the connectivity is checked. At the beginning of each experiment, every agent receives a random signal that is generated from hypothesis k by the distribution given in (2). While k is not known to the agents, P and K are known. Without loss of generality, we set $k = 0$ in all experiments.

In the first experiment, we simulated Models 1 and 2, and demonstrated the evolution of the agents' social belief as well as the agents' decisions. In Figure 1, we show the evolution of the network's average social belief over time, where the network's average social belief was obtained by an average of the social belief of all the agents in the network. We can see that in both methods the network's average social belief is nondecreasing over time. Though only one experiment is shown, from multiple experiments we know that the magnitude of P affects the rate of convergence. To be more specific, the greater P is, the more agents will eventually choose the correct hypothesis.

In Figure 2, we also show histograms of the distribution of the

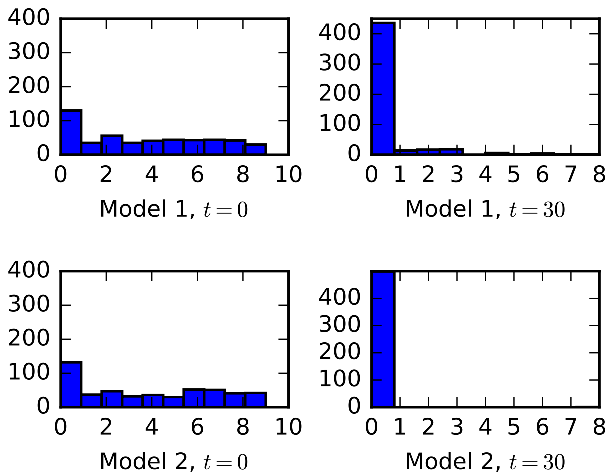

Fig. 2. The upper two subplots show the histogram of agents' decision at the starting and ending time slots for Model 1, while the lower two subplots show that for Model 2. The bars show the number of agents choosing the corresponding hypothesis. The experiments were run with the same parameters as Figure 1. At the beginning, the agents' decisions were distributed slightly in favor of the true hypothesis ($\mathcal{H}_0$) for both models. In the last time slot, for Model 1, the true hypothesis held the biggest cluster of agents, while for Model 2, all agents eventually chose $\mathcal{H}_0$.

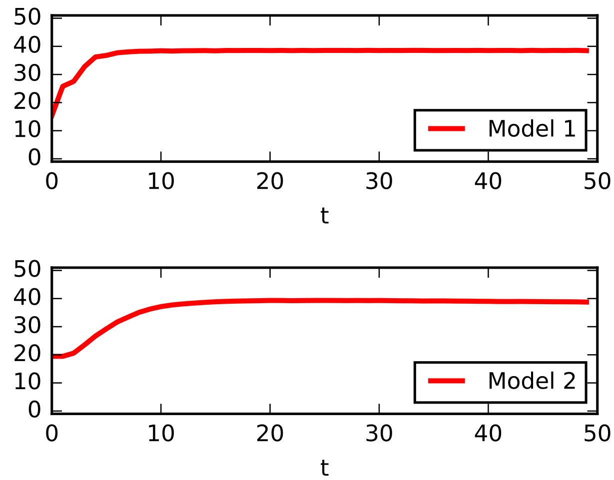

Fig. 3. Evolution of average number of agents choosing the true hypothesis over time. The upper subgraph shows the result of Model 1, while the lower subgraph shows the result of Model 2. Both experiments had parameters with total number of agents $N = 50$, number of hypotheses $K = 5$, probability of the correct hypothesis $P = 0.4$, and number of iterations $T = 50$. The Monte Carlo tests were run for 100 times. For both models, $\bar{\alpha}$ is generally nondecreasing with time.

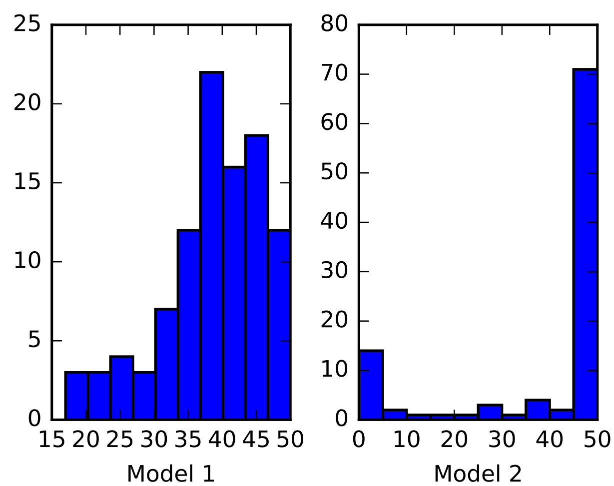

Fig. 4. Distribution of the final number of agents choosing the correct hypothesis ($\mathcal{H}_0$). The experiments were run with the same parameters as those in the caption of Figure 3. The left subplot is for Model 1 and the right for Model 2. The histograms show that in Model 1 the agents had the inclination to choose $\mathcal{H}_0$, while in Model 2, it is more likely that all the agents eventually clustered to $\mathcal{H}_0$.

agents' decisions, both, at the starting and ending time slots. We can see that according to Model 1, the true hypothesis ended up with the greatest number of believers, while the other hypotheses were still adopted by some agents. By contrast, all the agents of Model 2 eventually chose the correct hypothesis.

In the second experiment, we simulated Models 1 and 2 multiple times under the same settings. Figure 3 shows the evolution of the average number of agents choosing the correct hypothesis. We can see that the average number of agents choosing the correct hypothesis increased with time and then reached a stable state.

We also analyzed the convergence properties of Models 1 and 2. From Figure 4, we can see that in Model 1 not all the agents finally chose the true hypothesis, though usually a large numbers of them did. However, in Model 2, it is very likely that all the agents ended up with the correct decision. Yet, exceptional outcomes exist, where only a small number of agents chose the true hypothesis. A preliminary explanation of this phenomenon is that the final decisions of the agents depend heavily on the random signals they received at the beginning. Therefore, if the signals show less correlation with the correct hypothesis, it is possible that an agent's decision does not converge to the true one.

6. CONCLUSION

In this paper, we introduced a Bayesian learning model to study decision exchanges in a random network. We proposed two quasi-Bayesian models for the agents in the network to estimate the social beliefs of their neighbors. With Model 1, the agents are inclined to end up with the true hypothesis, with many agents still sticking to other hypotheses. With Model 2, the agents eventually converge to the true hypothesis. In the future, this work will be extended to (i) finding analytical results for the convergence properties of the network's average social belief and the expected probability of the agents choosing the true hypothesis and (ii) applying the model to real-life data and trying to predict outcomes.

7. REFERENCES

[1] R. Hegselmann and U. Krause, "Opinion dynamics and bounded confidence models, analysis, and simulation," *Journal of Artificial Societies and Social Simulation*, vol. 5, no. 3, 2002.

[2] W. Jager and F. Amblard, "Uniformity, bipolarization and pluriformity captured as generic stylized behavior with an agent-based simulation model of attitude change," *Computational & Mathematical Organization Theory*, vol. 10, no. 4, pp. 295–303, 2005.

[3] G. Deffuant, D. Neau, F. Amblard, and G. Weisbuch, "Mixing beliefs among interacting agents," *Advances in Complex Systems*, vol. 3, no. 01n04, pp. 87–98, 2000.

[4] 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.

[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] 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.

[7] V. D. Blondel, J. M. Hendrickx, and J. N. Tsitsiklis, "On Krause's multi-agent consensus model with state-dependent connectivity," *IEEE Transactions on Automatic Control*, vol. 54, no. 11, pp. 2586–2597, 2009.

[8] K. R. Varshney, "Bounded confidence opinion dynamics in a social network of Bayesian decision makers," *IEEE Journal of Selected Topics in Signal Processing*, vol. 8, no. 4, pp. 576–585, 2014.

[9] K. Sznajd-Weron and J. Sznajd, "Opinion evolution in closed community," *International Journal of Modern Physics C*, vol. 11, no. 06, pp. 1157–1165, 2000.

[10] D. Acemoglu and A. Ozdaglar, "Opinion dynamics and learning in social networks," *Dynamic Games and Applications*, vol. 1, no. 1, pp. 3–49, 2011.

[11] A. C. R. Martins, "Continuous opinions and discrete actions in opinion dynamics problems," *International Journal of Modern Physics C*, vol. 19, no. 04, pp. 617–624, 2008.

[12] K. Fan and W. Pedrycz, "Emergence and spread of extremist opinions," *Physica A: Statistical Mechanics and its Applications*, vol. 436, pp. 87–97, 2015.

[13] Y. Wang and P. M. Djurić, "Social learning with Bayesian agents and random decision making," *IEEE Transactions on Signal Processing*, vol. 63, no. 12, pp. 3241–3250, 2015.

[14] 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.

[15] M. A. Rahimian and A. Jadbabaie, "Bayesian learning without recall," *arXiv preprint arXiv:1601.06103*, 2016.

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

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

[18] 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.