OPINION DYNAMICS IN MULTI-AGENT SYSTEMS WITH BINARY DECISION EXCHANGES

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

*Department of ECE and DTC, University of Minnesota, 55455, USA
†Department of ECE, Stony Brook University, Stony Brook, NY 11794

ABSTRACT

Opinion dynamics in social networks has been widely studied in recent years, mostly by considering exchanges of opinions among neighboring agents. This paper addresses a scenario where the agents make decisions repeatedly on two hypotheses and where agents only exchange decisions. Motivated by the Bayesian models in the literature of human cognition, we model this learning procedure by the Bayes' rule. The social belief of each agent is defined to be the posterior of one of the hypotheses conditioned on the information it obtained from the society. We show that under certain conditions, once the social belief evolves to some region, the agent will refuse to change its belief. We demonstrate the asymptotical properties of the proposed model by computer simulations.

Index Terms— Opinion dynamics, Bayesian learning, voter model,

1. INTRODUCTION

Opinion dynamics have been widely studied in the fields of economy [1], sociology [2], and engineering [3] with the aim of understanding how opinions among people in a social network evolve over time. Social networks are modeled as multi-agent systems described by connected graphs where each node represents one social agent. Unlike many cooperative systems where consensus is of interest and may be reached asymptotically, opinions in social networks often differ from each other and stay that way.

In one category of models, the interactions among agents are modeled by opinion exchanges; see [2, 4, 5, 6]. Two discrete-time models based on the idea of bounded confidence are the Krause model [2, 7] and the Deffuant-Weisbuch model [5]. There, the opinion exchanges take place only among agents whose differences in opinions are no larger than a threshold. According to these models, the opinions eventually evolve to one or several clusters depending on the model of the agents' observations.

Another category of models assumes that every agent makes decisions based on its opinion, and only decisions can be exchanged among neighboring agents [8, 9, 10, 11]. The opinions are hidden behind the decisions, and the agents cannot fuse them directly. Instead, they make inference about the opinions based on the decisions they observe. In other words, every agent learns from its neighbors. In addressing this learning procedure, several Bayesian [12, 13] and non-Bayesian learning [14, 15] models have been proposed.

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

Motivated by the similarity between human cognition and the Bayes's rule [16], we propose a model where every agent uses the Bayes' rule to learn from decisions. In particular, we consider a network of N agents that decide on one of two hypotheses. Similarly to the voter model in [17], at every time slot, one agent is randomly selected, and it chooses one of its neighbors's decision for learning. In [8, 9], the authors propose that once a decision is observed, the learner adjusts its log opinion ratio by adding or subtracting a fixed small value. In contrast to [8, 9], this paper proposes a scheme that quantifies this value by leveraging Bayesian inference. As a result, this value becomes dynamic. Moreover, under certain conditions, we prove the existence of two regions of "stubbornness." Once an agent's opinion evolves to such region, the opinion is not changed any more. By simulations, we show that the opinion of every agent evolves to these regions in finite time.

The paper is organized as follows. In the next section we explain the social learning process and describe the studied system. In Sections 3 and 4, we present and analyze the proposed Bayesian learning scheme. Simulation results are given in Section 5, and concluding remarks are made in Section 6.

2. PROBLEM FORMULATION

Consider a multi-agent system that consists of N agents A_i, $i \in \mathcal{N}_A = \{1, 2, ..., N\}$. The agents are capable of performing local computations and of exchanging binary decisions with other agents. Each agent A_i receives a random private observation $y_i \in \mathbb{R}$, which is generated according to either hypothesis $\mathcal{H}_0$ or $\mathcal{H}_1$. The probability distribution of the y_is under $\mathcal{H}_k$, $k \in \{0, 1\}$, are given by

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

where $\phi_k(\cdot)$ stands for the distribution of y_i under $\mathcal{H}_k$, which belongs to the exponential family of distributions, i.e.,

$$\phi_k(y_i) = h(y_i) \exp(\eta(\theta_k) M(y_i) - A(\theta_k)), \tag{2}$$

where the parameters θ_k and η are known to the agents. In this work, the log-likelihood ratio (LLR) of the hypotheses is defined by $\log\left(\frac{\phi_1(y_i)}{\phi_0(y_i)}\right)$. We distinguish two types of LLRs, bounded and unbounded. An LLR is bounded if there exist two finite real numbers z and Z, such that $\forall\, y_i$, $\log\left(\frac{\phi_1(y_i)}{\phi_0(y_i)}\right) \in [z, Z]$. Otherwise, the LLR is unbounded.

At each time slot t, every agent A_i maintains an opinion on $\mathcal{H}_1$ quantified by the posterior and according to the Bayes' rule. This posterior is given by

$$\beta_i^{(t)} = p(\mathcal{H}_1|y_i, \mathcal{I}_i^{(t)}), \quad (3)$$
$$\propto \pi_i^{(t)} p(y_i|\mathcal{H}_1), \quad (4)$$

where $\pi_i^{(t)}$, referred to as "social belief", stands for the belief of A_i on $\mathcal{H}_1$ based on all of the information available to it other than y_i and denoted by $\mathcal{I}_i^{(t)}$. This $\pi_i^{(t)}$ 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, $\pi_i^{(1)} = 0.5, \forall i \in \mathcal{N}_A$. Equivalently, the Bayes' rule in the log-form can be expressed as

$$\log \frac{1-\beta_i^{(t)}}{\beta_i^{(t)}} = \log \frac{1-\pi_i^{(t)}}{\pi_i^{(t)}} + \log \frac{\phi_0(y_i)}{\phi_1(y_i)}. \quad (5)$$

We remark that the opinion of A_i is formed by (a) its social belief that summarizes the information from the society, and (b) the log-likelihood ratio determined by its private observation.

With the opinion $\beta_i^{(t)}$, the A_i's decision on choosing one of the hypotheses is made according to

$$\alpha_n^{(t+1)} = I(\beta_i^{(t)} > 0.5), \quad (6)$$

where $I(\cdot) \in \{0, 1\}$ is the indicator function on the truthfulness of the statement inside the parentheses.

In summary, the private observation y_i is constant and therefore, the dynamics of the opinion of A_i is fully defined by the dynamics of its social belief. Thus, modeling the evolution of the social beliefs is equivalent to modeling the dynamics of the opinion in the multi-agent system. In the next section, we will build a Bayesian learning model for updating the social belief at each time slot.

3. VOTER MODEL WITH BAYESIAN LEARNING

Before we proceed to the proposed model, we first sketch the voter model from [17]. According to the voter model, at each time slot, one agent A_i is randomly selected and it randomly picks one of its neighbors, say $A_j \in \mathcal{N}_i$, to follow. Namely, A_i makes its decision as $\alpha_i^{(t+1)} = \alpha_j^{(t)}$, and the the remaining $N-1$ agents keep their opinions unchanged.

Here we propose a modified voter model where the agent A_i learns from the decisions of its neighbors by the Bayesian method. More precisely, instead of blindly following the decision of A_j, A_i adjusts its private belief based on $\alpha_j^{(t)}$. By Bayes' rule, we have

$$\log \frac{1-\pi_i^{(t+1)}}{\pi_i^{(t+1)}} = \log \frac{1-\pi_i^{(t)}}{\pi_i^{(t)}} + \log \frac{p(\alpha_j^{(t)}|\mathcal{H}_0)}{p(\alpha_j^{(t)}|\mathcal{H}_1)}. \quad (7)$$

Due to the property of the $\log(\cdot)$ function, it can be shown that

$$\log \frac{1-\beta_i^{(t+1)}}{\beta_i^{(t+1)}} = \log \frac{1-\beta_i^{(t)}}{\beta_i^{(t)}} + \log \frac{p(\alpha_j^{(t)}|\mathcal{H}_0)}{p(\alpha_j^{(t)}|\mathcal{H}_1)}, \quad (8)$$

where the "action likelihood" $p(\alpha_j^{(t)}|\mathcal{H}_k)$ denotes the probability that A_j makes the decision $\alpha_j^{(t)}$ under $\mathcal{H}_k$.

According to (3) and (6), $\alpha_j^{(t)} = 1$ or $\alpha_j^{(t)} = 0$ if and only if

$$\frac{p(y_j|\mathcal{H}_1)}{p(y_j|\mathcal{H}_0)} \gtrless \frac{1-\pi_j^{(t)}}{\pi_j^{(t)}}, \quad (9)$$

where $A \gtrless B$ means that A is greater or less than B, respectively. When the distributions of the observations are given by (2), we have

$$\log\left(\frac{p(y_j|\mathcal{H}_1)}{p(y_j|\mathcal{H}_0)}\right) = \Big(\eta(\theta_1)-\eta(\theta_0)\Big)T_j - \Big(A(\theta_1)-A(\theta_0)\Big), \quad (10)$$

where $T_j = M(y_j)$. Without loss of generality, we assume that $\eta(\theta_1) - \eta(\theta_0) > 0$. Then, if and only if

$$T_j \gtrless \gamma_j^{(t)} \quad (11)$$
$$= \frac{(A(\theta_1) - A(\theta_0)) + \log(\frac{1-\pi_j^{(t)}}{\pi_j^{(t)}})}{\eta(\theta_1) - \eta(\theta_0)}, \quad (12)$$

the decision $\alpha_j^{(t)}$ is one or zero, respectively. Provided $\mathcal{H}_k$ is true, the probability of $\alpha_j^{(t)} = 1$ is defined by

$$l_{j,k}^{(t)} = Pr(\alpha_j^{(t)} = 1|\mathcal{H}_k) = Pr(T_j > \gamma_j^{(t)}|\mathcal{H}_k), \quad (13)$$

suggesting that $Pr(\alpha_j^{(t)} = 0|\mathcal{H}_k) = 1 - l_{j,k}^{(t)}$.

In computing $l_{j,k}^{(t)}$, A_i must know its neighbor's current social belief $\pi_j^{(t)}$, which, however, is unknown to A_i. Consequently, A_i has to make a "guess" on $\pi_j^{(t)}$. We propose two guessing methods. They are as follows:

Method 1: The agent A_i considers $\pi_j^{(t)}$ to be a point estimate and guesses that it is equal to its own social belief. Namely, in the inference procedure of A_i, $\pi_j^{(t)} = \pi_i^{(t)}$, i.e., $p(\pi_j^{(t)}) = \delta(\pi_j^{(t)} - \pi_i^{(t)})$. The motivation for this approximation is that A_i and A_j are neighboring agents in the network, and therefore there is a large chance that they have common neighbors. The shared neighbors entail that A_i and A_j get similar information from the society.

Method 2: The agent A_i draws $\pi_j^{(t)} \in [0, 1]$ from a Beta distribution, i.e., $\pi_j^{(t)} \sim \text{Beta}(a_j^{(t)}, b_j^{(t)})$ given by

$$p(\pi_j^{(t)}|a_j^{(t)}, b_j^{(t)}) = \frac{\Gamma\left(a_j^{(t)} + b_j^{(t)}\right)}{\Gamma\left(a_j^{(t)}\right)\Gamma\left(b_j^{(t)}\right)} \times \left(\pi_j^{(t)}\right)^{a_j^{(t)}-1}\left(1-\pi_j^{(t)}\right)^{b_j^{(t)}-1} \quad (14)$$

with $a_j^{(t)}$ and $b_j^{(t)}$ satisfying

$$\frac{a_j^{(t)}}{a_j^{(t)} + b_j^{(t)}} = \mathbb{E}\,\pi_j^{(t)} = \pi_i^{(t)}, \quad (15)$$
$$a_j^{(t)} + b_j^{(t)} - 2 = n_j^{(t)}, \quad (16)$$

where $n_j^{(t)}$ is the number of times A_i selects A_j up until t. In (15), we write $\mathbb{E}\,\pi_j^{(t)} = \pi_i^{(t)}$ for the same reason as in Method 1. Recall that if we have a Bernoulli random variable $S \in \text{Ber}(q)$ with n

samples, $\mathbf{s} = [s_1, \cdots, s_n]$, independently drawn from Ber(q), then given a Beta prior with parameters $a_0 = b_0 = 1$ (a uniform prior of q on $[0, 1]$), the posterior $p(q|\mathbf{s})$ is also a Beta distribution but with parameters a_n and b_n, which satisfy $a_n + b_n - 2 = n$ as in (16).

With both methods, A_i computes $l_{j,k}^{(t)}$ according to

$$l_{j,k}^{(t)} = \int_0^1 Pr(\alpha_j^{(t)} = 1|\mathcal{H}_k, \pi_j^{(t)})\, p(\pi_j^{(t)})\, \mathrm{d}\pi_j^{(t)}. \quad (17)$$

With method one, (17) simplifies to

$$l_{j,k}^{(t)} = Pr(\alpha_j^{(t)} = 1|\mathcal{H}_k, \pi_j^{(t)} = \pi_i^{(t)}). \quad (18)$$

4. ANALYSIS

In this section, we analyze the opinion dynamics with the proposed models. In the following proposition, we show that the social belief of A_i is non-decreasing if it receives a neighbor's decision of one. By contrast, the social belief of A_i is non-increasing if it receives a neighbor's decision of zero.

Proposition 1 *In the proposed model with both methods, let A_i be selected at time t and let it choose A_j for learning. Then $\pi_i^{(t+1)} \geq \pi_i^{(t)}$ if $\alpha_j^{(t)} = 1$ and $\pi_i^{(t+1)} \leq \pi_i^{(t)}$, otherwise.*

Proof: We first show the case when $\alpha_j^{(t)} = 1$. From (7), it is sufficient to show that $\log\left(\frac{p(\alpha_j^{(t)}|\mathcal{H}_0)}{p(\alpha_j^{(t)}|\mathcal{H}_1)}\right) \leq 0$, or equivalently, we need to show that $w_1 \geq w_0$, where $w_k = Pr(\alpha_j^{(t)} = 1|\mathcal{H}_k)$.

By the decision making policy in (6), the decision region for $\alpha_j^{(t)} = 1$, denoted by $\mathcal{S}_1$, can be written as

$$\mathcal{S}_1 = \left\{ y_j \Big| \frac{\phi_1(y_j)}{\phi_0(y_j)} > \frac{1 - \pi_j^{(t)}}{\pi_j^{(t)}} \right\}. \quad (19)$$

Also we can define $\mathcal{S}_0 = \mathcal{S} \backslash \mathcal{S}_1$ as the decision region for $\alpha_j^{(t)} = 0$, where $\mathcal{S}$ is the support of y_j.

Therefore, we have that

$$w_k = \int_{y_j \in \mathcal{S}_1} \phi_k(y_j) \mathrm{d}y_j. \quad (20)$$

Next, we show that $w_1 - w_0 \geq 0$ by considering the following two cases. If $\pi_j^{(t)} \leq 0.5$, we have that

$$w_1 - w_0 = \int_{y_j \in \mathcal{S}_1} \big(\phi_1(y_j) - \phi_0(y_j)\big) \mathrm{d}y_j \geq 0, \quad (21)$$

because $\forall y_j \in \mathcal{S}_1$, $\phi_1(y_j) \geq \phi_0(y_j)$. Otherwise, $\pi_j^{(t)} > 0.5$ implies that $\forall y_j \in \mathcal{S}_0$, $\phi_1(y_j) < \phi_0(y_j)$. Then we have

$$\begin{aligned} w_1 - w_0 &= (1 - w_0) - (1 - w_1) \\ &= \int_{y_j \in \mathcal{S}_0} \big(\phi_0(y_j) - \phi_1(y_j)\big) \mathrm{d}y_j > 0. \quad (22) \end{aligned}$$

With (21) and (22), the proof of the case when $\alpha_j^{(t)} = 1$ is completed. When $\alpha_j^{(t)} = 0$, $\pi_i^{(t+1)} \leq \pi_i^{(t)}$ can be proved by the same procedure with just notational changes. □

With the next proposition, we show that when the LLR is bounded, if the opinion of an agent evolves to an interval close to zero or one, it will become a "stubborn" agent. Namely, the agent will stop learning from its neighbors and will keep its social belief unchanged forever.

Proposition 2 *In the proposed model with Method 1, if the LLR is bounded, then there exist two real numbers u and U with $0 < u < U < 1$, such that once $\pi_i^{(t)}$ evolves to the intervals $[0, u)$ and $(U, 1]$, $\pi_i^{(\tau)} = \pi_i^{(t)}, \forall \tau > t$.*

Proof: Since the LLR is bounded, there exist $z \in \mathbb{R}$ and $Z \in \mathbb{R}$, such that $z < \log\left(\frac{p(y_i|\mathcal{H}_1)}{p(y_i|\mathcal{H}_0)}\right) < Z$. From (9), we have that if $\log\left(\frac{1-\pi_i^{(t)}}{\pi_i^{(t)}}\right) > Z$, the agent A_i will make its decision $\alpha_i^{(t)} = 0$ regardless of y_i. Similarly, if $\log\left(\frac{1-\pi_i^{(t)}}{\pi_i^{(t)}}\right) < z$, $\alpha_i^{(t)} = 1$ for any y_i. Define $u = \frac{1}{1+e^{-z}}$ and $U = \frac{1}{1+e^{-Z}}$. Then we get that if $\pi_i^{(t)} < u$ or $U < \pi_i^{(t)}$, A_i makes its decision independently of y_i. In both scenarios, we have

$$\log\left(\frac{p(\alpha_i^{(t)}|\mathcal{H}_0)}{p(\alpha_i^{(t)}|\mathcal{H}_1)}\right) = 0. \quad (23)$$

Consider that in Method 1, A_i uses its own social belief $\pi_i^{(t)}$ as $\pi_j^{(t)}$ to make inference on the decision of A_j. Thus, once $\pi_i^{(t)}$ evolves to the intervals $[0, u)$ and $(U, 1]$, A_i will get $\log\left(\frac{p(\alpha_j^{(t)}|\mathcal{H}_0)}{p(\alpha_j^{(t)}|\mathcal{H}_1)}\right) = 0$, implying that $\pi_i^{(t+1)} = \pi_i^{(t)}$. Therefore, $\pi_i^{(t)}$ will become unchanged hereafter. □

Here we remark that Proposition 2 indicates that A_i will also stop updating its private belief once it becomes stubborn. This is because of the private belief formulation in (3). From the above two propositions, it is tempting to conclude that with Method 1, there exists a finite time $\tilde{t}$ such that beyond $\tilde{t}$, every agent's belief evolves to the stubborn region and no agent changes its belief anymore. In the next section we will show with our simulations that we observed this phenomenon in every trial. Nevertheless, we do not have a rigorous proof of this conjecture yet.

5. SIMULATION RESULTS

In this section, we provide two simulation experiments to demonstrate the properties of the proposed models. In both experiments, the multi-agent system was modeled by a random geometric graph $G(\mathcal{N}_A, \mathcal{E})$, where the N agents were chosen uniformly and independently on a square of size 1×1. Each pair of agents was connected if the Euclidian distance between them was smaller than $\sqrt{\log(N)/N}$, which is due to the connectivity requirement. Once the graph was formed, its connectivity was checked.

We considered the following Binomial hypothesis testing where the agents had to make a decision between the two hypotheses,

$$\mathcal{H}_1: \quad y_n \sim \text{Bin}\,(n, p_1), \quad (24)$$
$$\mathcal{H}_0: \quad y_n \sim \text{Bin}\,(n, p_0), \quad (25)$$

where n, p_0 and p_1 were known by all the agents. For the prior probabilities of the hypotheses, we let $P(\mathcal{H}_0) = P(\mathcal{H}_1) = 1/2$. Without loss of generality, we assumed that the data were generated from $\mathcal{H}_1$. In both experiments we set $n = 7$, $p_1 = 0.51$ and $p_0 = 0.49$. The number of agents was 20.

In the first experiment, we modeled the behavior of the agents when they adopted Method 1. In Fig. 1, the top-left and bottom-left plots show two realizations of the evolution of the number of agents at time t, $K^{(t)}$, making decision 1. We can see that the number of agents making decision 1 eventually became stable. In the plots on the right, we displayed the evolution of the social beliefs of the agents over time in a same realization. The abscissa represents the time index and the ordinate the social belief of all the 20 agents at a given time t. In the top-right plot, the social belief of the agents evolved to the same interval, while in the bottom plot, the belief diverged into two intervals. From the figures we can see that after reaching one of the thresholds, the beliefs of the agents stayed stable.

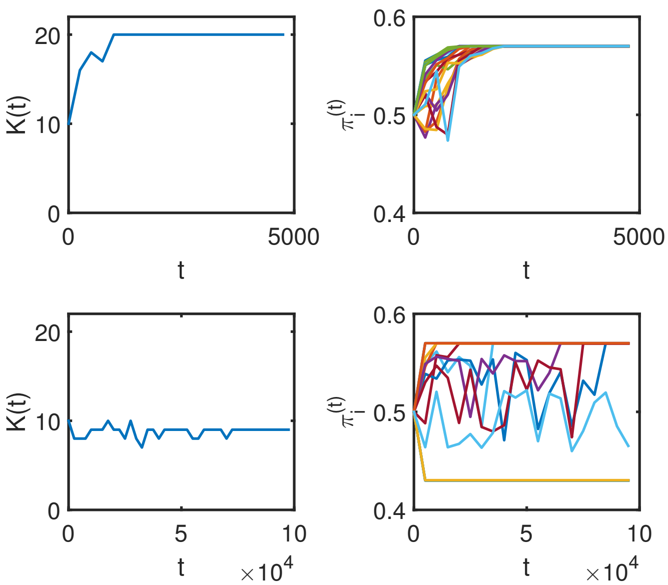

Fig. 1. Top-left and bottom-left: evolution of number of agents with time that make decision one using Method 1; Top-right and bottom-right: evolution of agents' social beliefs with time using Method 1. The two figures on top depict the results from one and the bottom figures the results from another realization.

In Fig. 2, we plotted the histograms of the social beliefs of the agents in the first experiment at time slots t = 1, 20, 50, 100, 200, 500, 1000, 1500, 2000, respectively. From the histograms, we can see that these beliefs started from 0.5, which corresponds to the uninformative prior we have set. After iterations of belief updating, the beliefs gradually moved into the stable intervals and stopped evolving.

In the second experiment, we modeled the behavior of the agents when they use Beta distribution to approximate the social beliefs of the other agents. This is our Method 2. The results are shown in Fig. 3. In contrast to the first model, we did not observe that the social beliefs evolve to two clusters. Instead, we see that all the social beliefs converge to a very small interval.

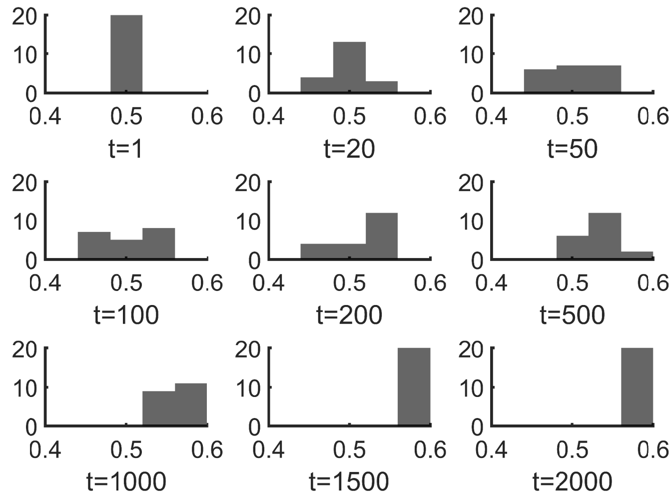

Fig. 2. The histogram of the agents' social belief at time slots $t = 1, 20, 50, 100, 200, 500, 1000, 1500, 2000$ when they use Method 1.

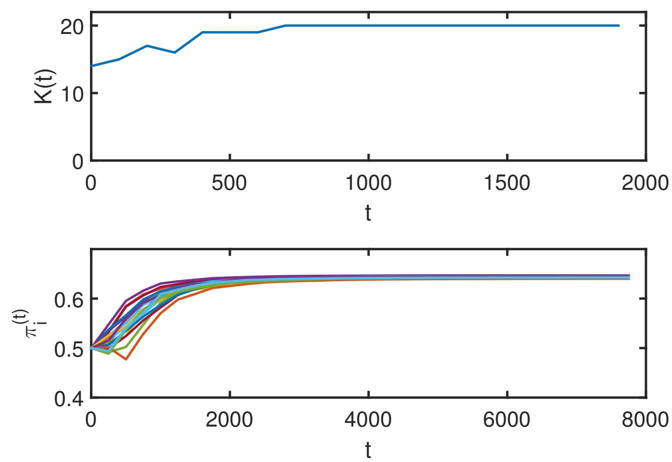

Fig. 3. Top: The evolution of the number of agents making decision one over time using Method 2; Bottom: The evolution of agents' social beliefs over time using Method 2.

6. CONCLUSION

In this paper, a Bayesian learning model was introduced for studying opinion dynamics in systems with decision exchanges. We presented analysis of the asymptotic properties of the proposed models and verified them by simulations. It was shown that with a bounded LLR every agent eventually becomes stubborn if for inference of the opinions of their neighbors the agents use Method 1. By contrast, if they use Method 2, the social beliefs of the agents will not stop evolving. This work has several intriguing directions. They include (i) finding the relationship between the number of formed clusters and the observation model and (ii) determining the critical time after which the social belief stops evolving.

7. REFERENCES

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

[2] 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.

[3] 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.

[4] 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.

[5] 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.

[6] 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.

[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] 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.

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

[10] 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.

[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] 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.

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

[14] 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.

[15] 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.

[16] J. B. Tenenbaum, T. L. Griffiths, and C. Kemp, "Theory-based Bayesian models of inductive learning and reasoning," *Trends in cognitive sciences*, vol. 10, no. 7, pp. 309–318, 2006.

[17] R. A. Holley and T. M. Liggett, "Ergodic theorems for weakly interacting infinite systems and the voter model," *The Annals of Probability*, pp. 643–663, 1975.