Convergence of Rule-of-Thumb Learning Rules in Social Networks

Daron Acemoglu
Department of Economics
Massachusetts Institute of Technology
Cambridge, MA 02142
Email: daron@mit.edu

Angelia Nedić
Department of Industrial and
Enterprise Systems Engineering
University of Illinois
Urbana-Champaign, IL 61801
Email: angelia@uiuc.edu

Asuman Ozdaglar
Department of Electrical Engineering
and Computer Science
Massachusetts Institute of Technology
Cambridge, MA 02142
Email: asuman@mit.edu

***Abstract*— We study the problem of dynamic learning by a social network of agents. Each agent receives a signal about an underlying state and communicates with a subset of agents (his neighbors) in each period. The network is connected. In contrast to the majority of existing learning models, we focus on the case where the underlying state is time-varying. We consider the following class of *rule of thumb learning rules*: at each period, each agent constructs his posterior as a weighted average of his prior, his signal and the information he receives from neighbors. The weights given to signals and neighbors can vary across agents and over time. We distinguish between two subclasses: (1) constant weight rules; (2) diminishing weight rules. The latter reduces weights given to signals asymptotically to 0. Our main results characterize the asymptotic behavior of beliefs. We show that the general class of rules leads to unbiased estimates of the underlying state. When the underlying state has innovations with variance tending to zero asymptotically, we show that the diminishing weight rules ensure convergence in the mean-square sense. In contrast, when the underlying state has persistent innovations, constant weight rules enable us to characterize explicit bounds on the mean square error between an agent's belief and the underlying state as a function of the type of learning rule and signal structure.**

I. Introduction

A central question of social sciences is how a large group of agents form their beliefs about an underlying state of the world, which may correspond to economic or social opportunities or the appropriateness of different policies and actions. A variety of evidence shows that individuals form their beliefs in part on the basis of their *social network*, consisting of friends, neighbors, coworkers and family members who communicate relatively frequently. For example, Granovetter [12] and Ioannides and Loury [13] document how information obtained from an individual's social network influences employment opportunities, while Foster and Rosenzweig [10] and Munshi [15] show the importance of the information obtained from social networks for technology adoption.

These observations have motivated a large literature investigating how the structure of a social network and the pattern of communication influences the dynamics of beliefs and whether individuals will form unbiased beliefs about the underlying state. A natural approach is to formulate this problem as a dynamic game and characterize its (perfect Bayesian) equilibria. Though theoretically attractive, this strategy runs into two problems. First, the characterization of such equilibria in complex networks is generally a non-tractable problem.[1] Second, individuals in practice use somewhat coarser ways of aggregating information forming their beliefs than Bayesian updating. These considerations have motivated a large literature to investigate how individuals form their beliefs when they use "reasonable rules of thumb".[2] One difficulty is that what is "reasonable" is generally difficult to decide without reference to the structure of the network and the learning problem. Another challenge for this entire literature is that it focuses on situations in which there is a fixed underlying state, which remains constant while individuals accumulate more signals. Most real-world situations involve at least some change in the underlying state. For example, which types of jobs offer better careers, what brands of products have higher quality and which parties are more reliable provide examples of underlying states on which individuals form beliefs, but in none of these cases these states are likely to remain fixed for long periods of time.

In this paper, we contribute to this literature in three ways. First, we consider the class of rule-of-thumb learning rules that are flexible and nest various Bayesian rules (in the case of normally distributed states and signals). In particular, we look at learning rules that take time-varying averages of an agent's prior belief, new signal and observations of neighbors. Second, we allow for a time-varying underlying state. Finally, though as in the previous literature, we look at asymptotic learning (in particular, whether individuals form unbiased expectations of the underlying state and whether they learn the true value of the underlying state asymptotically), we also provide explicit bounds on the mean square error between individual beliefs and the underlying state.

This research was partially supported by the National Science Foundation under CAREER grants CMMI 07-42538 and DMI-0545910.

[1] For this reason, existing literature focuses on relatively simple and stylized environments. See, for example, Bikchandani, Hirshleifer and Welch [5], Banerjee [4], and Smith and Sorensen [20] on models where each individual takes a single action and observes all past actions. Acemoglu, Dahleh, Lobel and Ozdaglar [1] generalize these results to an arbitrary network structure (while keeping the assumption of a single decision for each agent).

[2] See, for example, Ellison and Fudenberg [8], [9], Bala and Goyal [2], [3], DeMarzo, Vayanos and Zwiebel [7] and Golub and Jackson [11].

We distinguish two cases. In the first, the underlying state changes with *diminishing innovations*, that is, the variance of the innovations to the underlying state (when scaled appropriately) tends to zero asymptotically. In the second, the underlying state changes with *persistent innovations*, in the sense that, the variance of the innovations to the underlying state need not go to zero. We also consider two subclasses of the general class of learning rules described above. These are: (1) *constant weight rules*; (2) *diminishing weight rules*. Diminishing weight rules reduce the weights given to own signals to 0 asymptotically.[3] In contrast, with constant weight rules, these weights are uniformly bounded away from zero. Throughout, we consider a general social network, represented by a direct graph, which specifies how information flows across individuals. We assume that the graph is connected (that is, there exists a path connecting each agent to every other).

Our first result shows that any learning rule within the class we consider is unbiased. Therefore, this fairly general class of reasonable learning rules all preclude systematic biases in the formation of beliefs. Our second result characterizes bounds on the mean square error of the gap between an agent's belief and the underlying state. When there are diminishing innovations, diminishing weight rules ensure convergence in the mean-square sense (so that the second moment of the gap between beliefs and the true value of the underlying state tend to zero asymptotically) and in probability. This reflects the intuitive notion that when there are diminishing innovations, the underlying state will essentially stabilize at some level in the future and there is enough information acquired by individuals from their own signals and from the communication of others that they will be able to track this state fairly well. Diminishing weight rules ensure that in the very far future, individuals will not respond excessively to their own signal, thus enabling convergence in the mean-square sense. This result is typically not true with constant weight rules, since even when individuals have accumulated a lot of information about the underlying state, their beliefs are still highly responsive to their own signals. In contrast, when there are persistent innovations, diminishing weight rules do not perform well. In this case, each agent needs to track a moving target and diminishing weight rules cease to incorporate the new information about this target (underlying state). However, constant weight rules perform well in this case and provide us with explicit bounds on the second moment of the gap between individual beliefs and the underlying state. This bound is a function of the parameters of the learning rule and the information structure.

These results are intuitive and suggest that relatively simple learning rules can ensure learning of the underlying state or accurate tracking of the changes in this state by the agents under a wide range of network topologies. The intuition can best be understood by going back to the normal updating case. The normal updating formula for the estimation of an unknown, but fixed parameter is a special case of our diminishing weight rules, whereas the updating formula for the estimation of a time varying parameter corresponds to constant weight rules. This intuition also highlights that we may indeed expect agents to use rule of thumb rules with diminishing weights when the state is fixed and with constant weights when the state is time varying. Overall, our analysis therefore suggests that a limited amount of rationality is sufficient for learning and for efficient exchange of information over the network.

In the learning literature, our paper is most closely related to DeMarzo, Vayanos and Zwiebel [7] and Golub and Jackson [11], who also consider learning over a social network represented by a connected graph. These papers focus on the subset of our rules corresponding to constant weight rules and assume that the underlying state is fixed. Our results generalize their findings both by showing learning for a wider variety of environments and also by characterizing learning behavior when the underlying state is changing.

In addition to the papers on learning mentioned above, our paper is related to work on consensus, which is motivated by different problems, but typically leads to a similar mathematical formulation ([21], [22], [14], [6], [17], [18], [16]). In consensus problems, the focus is on whether the beliefs or the values held by different units (which might correspond to individuals, sensors or distributed processors) converge to a common value. Our analysis here does not only focus on consensus, but also whether the consensus happens around the true value of the underlying state and whether this consensus is able to track changes in the underlying state.

The rest of this paper is organized as follows: In Section II, we introduce our notation, formulate the learning problem, and describe the assumptions imposed on the weights used in forming the posterior beliefs. In Section III, we present our main convergence and rate of convergence results. In Section IV, we provide concluding remarks.

Regarding notation, for a matrix A, we write $[A]_i^j$ to denote the matrix entry in the i-th row and j-th column. We write $[A]_i$ and $[A]^j$ to denote the i-th row and the j-th column of the matrix A, respectively. A vector a is said to be a *stochastic vector* when $a_i \geq 0$ for all i and $\sum_i a_i = 1$. A square matrix A is said to be a *(row) stochastic matrix* when each row of A is a stochastic vector. The transpose of a matrix A is denoted by A'.

II. Model

We consider a set $V = \{1, \ldots, m\}$ of agents interacting over a static social network. We use the weights a_j^i, $i, j \in \{1, \ldots, m\}$, to capture the agent interactions, i.e., the nonnegative scalar a_j^i denotes the *weight or trust that agent i places on the information he receives from neighboring agent j*.

We study a model of learning where the agents are trying to learn some unknown parameter $\theta(k) \in \mathbb{R}$ varying over

[3] More generally, these rules can also reduce the weights given to neighbors to zero. In this paper, we keep the weights given to neighbors fixed to simplify notation.

time k according to[4]

$$\theta(k+1) = \theta(k) + \nu(k),$$

where $\nu(k)$ is a zero mean *innovation*, which is independent across time. The initial state $\theta(0)$ is assumed to be an independently drawn random variable.

At each time k, every agent i receives a private signal $s_i(k) \in \mathbb{R}$ which is a noisy version of the value of the parameter at time k. More specifically, we assume that at each time $k \geq 1$, the private signal $s_i(k)$ of agent i is generated according to

$$s_i(k) = \theta(k) + \nu_i(k) \qquad i \in \{1, \ldots, m\}, \tag{1}$$

where $\nu_i(k)$ is a zero-mean *observation noise*, which is independent across agents and time.

We use $x_i(k) \in \mathbb{R}$ to denote the belief of agent i about the value of parameter $\theta(k)$ at time k. Each agent i starts with some initial belief $x_i(0)$. At each time, the agents exchange their beliefs with their neighbors. At time k, agent i updates his belief according to the following rule:

$$\begin{aligned} x_i(k+1) &= x_i(k) + \alpha(k)(s_i(k) - x_i(k)) \\ &\quad + \sum_{j \neq i} a_j^i (x_j(k) - x_i(k)), \end{aligned} \tag{2}$$

where $\alpha(k)$ is a positive weight given to the information due to the private signal, which we refer to as the *signal weight*.

This rule captures the idea that in updating his belief, each agent uses the new information in his signal relative to his current belief and the new information provided by its neighbors relative to his own belief. While this update rule does not correspond to Bayesian updating of the posterior beliefs, it is flexible enough to nest a wide range of behavior. In particular, it corresponds to Bayesian updating if the state, the innovations, the signals and priors were all normally distributed.

III. Convergence of Beliefs

In this section, we study the convergence of beliefs under different assumptions on the signal weight and on the statistical properties of the innovation and the observation noise.

We assume that in updating their beliefs at time k, agents take a convex combination of the agent beliefs available at that time. This amounts to assuming that the weights $a_1^i, \ldots, a_m^i$ sum to 1 for all i, hence implicitly defining the weight a_i^i as $a_i^i = 1 - \sum_{j \neq i} a_j^i$. We use the $m \times m$ matrix A defined by $[A]_j^i = a_j^i$ for all $i, j \in \{1, \ldots, m\}$ to capture the agent interactions in the social network.

Assumption 1: (Stochastic Weights) *For each i, the weights a_j^i are nonnegative for all j and they sum to 1, i.e.,*

$$a_j^i \geq 0 \quad \text{for all } i, j, \qquad \sum_{j=1}^m a_j^i = 1 \quad \text{for all } i.$$

[4]Our model is applicable to a more general case where the parameter θ, the agent beliefs x_i, the innovation ν and the private signals ν_i are vectors, in which case our analysis can be applied componentwise.

Under the Stochastic Weights assumption, the belief update rule of agent i [cf. Eq. (2)] can be equivalently represented as

$$x_i(k+1) = \sum_{j=1}^m a_j^i x_j(k) + \alpha(k)(s_i(k) - x_i(k)). \tag{3}$$

In addition, we use another assumption on the weights.

Assumption 2: (Weights) *We have:*

(a) *There exists a scalar η with $0 < \eta < 1$ such that $a_i^i \geq \eta$ for all $i \in \{1, \ldots, m\}$, and if $a_j^i > 0$, then $a_j^i \geq \eta$.*
(b) *The directed graph (V, E), where E denotes the set of directed edges (j, i) such that $a_j^i > 0$, is strongly connected.*

Assumption 2(a) guarantees that each agent gives significant weights to his belief and the beliefs of his neighbors. Assumption 2(b) imposes a connectivity assumption on the social network.

For the observation noise $\nu_i(k)$ we assume independence across agents and time. Similarly, we assume that the innovation $\nu(k)$ is independent across time and independent of the observation noise $\nu_i(k)$ at all times. We also assume that the initial state $\theta(0)$ and the initial beliefs $x_i(0)$ are independently drawn random variables with finite second moments. We state these in the following.

Assumption 3: We have:

(a) *The random variables $\nu_i(k),\ i = 1, \ldots, m,\ k \geq 0$, are zero mean, and independent across agents and time, i.e.,*

$$E[\nu_i(k)] = 0 \quad \text{for all } i, k,$$

$$E[\nu_i(s)\nu_i(k)] = 0 \qquad \text{for all } i \ \text{ and all } \ k \neq s,$$

$$E[\nu_i(s)\nu_j(k)] = 0 \qquad \text{for all } i \neq j \ \text{ and all } \ k, s.$$

Furthermore, there exists a scalar σ_a^2 such that

$$E[\nu_i^2(k)] \leq \sigma_a^2 \qquad \text{for all } i, k.$$

(b) *The random variables $\nu(k),\ k \geq 0$, are zero mean, independent across time, and independent of $\nu_i(k)$ for all i and k, i.e., $E[\nu(k)] = 0$ for all k and*

$$E[\nu(s)\nu(k)] = 0 \qquad \text{when } s \neq k,$$

$$E[\nu(s)\nu_i(k)] = 0 \qquad \text{for all } i \ \text{ and all } k, s.$$

(c) *The random variables $\theta(0)$ and $x_i(0)$ are independent, i.e.,*

$$E[\theta(0)x_i(0)] = 0 \qquad \text{for all } i,$$

and there exist scalars σ_θ^2 and σ_x^2 such that

$$E[\theta^2(0)] \leq \sigma_\theta^2, \qquad E[x_i^2(0)] \leq \sigma_x^2 \quad \text{for all } i.$$

The next lemma provides a key relation in the tracking error $x_i(k) - \theta(k)$, which will be key in our subsequent analysis.

Lemma 1: Let Assumptions 1 and 2 hold. Let the sequences $\{x^i(k)\}$, $i = 1, \ldots, m$, be generated by Eq. (3). Assume that the signal weight $\alpha(k)$ satisfies

$$0 < \alpha(k) \leq a_i^i \qquad \text{for all } i,\ k.$$

Then, we have for all i and $k \geq 0$,

$$x_i(k+1) - \theta(k+1) = (1-\alpha(k)) \sum_{j=1}^{m} [B(k)]_j^i \,(x_j(k) - \theta(k)) + \alpha(k)\nu_i(k) - \nu(k), \quad (4)$$

where the matrix $B(k)'$ is a stochastic matrix for all $k \geq 0$.

Proof: By the definition of the update rule [cf. Eq. (3)] and the evolution of $\theta(k)$, we have

$$x_i(k+1) - \theta(k+1) = \sum_{j=1}^{m} a_j^i \, x_j(k) + \alpha(k)(s_i(k) - x_i(k)) - \theta(k) - \nu(k).$$

Since the vector $a^i = (a_1^i, \ldots, a_m^i)$ is stochastic (cf. Assumption 1), the preceding relation can be written as

$$x_i(k+1) - \theta(k+1) = \sum_{j=1}^{m} a_j^i \,(x_j(k) - \theta(k)) + \alpha(k)(s_i(k) - x_i(k)) - \nu(k).$$

By the definition of the private signal $s_i(k)$ in Eq. (1), we have

$$s_i(k) - x_i(k) = \theta(k) + \nu_i(k) - x_i(k).$$

Combining the preceding two relations, we obtain

$$x_i(k+1) - \theta(k+1) = \sum_{j=1}^{m} a_j^i \,(x_j(k) - \theta(k)) - \alpha(k)\,(x_i(k) - \theta(k)) + \alpha(k)\nu_i(k) - \nu(k).$$

By letting I denote the identity matrix, the first two terms on the right hand side can be compactly written, yielding

$$x_i(k+1) - \theta(k+1) = \sum_{j=1}^{m} [A - \alpha(k) I]_j^i \,(x_j(k) - \theta(k)) + \alpha(k)\nu_i(k) - \nu(k). \quad (5)$$

We introduce the matrix $B(k)$ as

$$B(k) = \frac{1}{1-\alpha(k)}\,(A - \alpha(k) I).$$

Since by assumption the signal weight is such that $\alpha(k) \leq a_i^i$ for all k, the matrix $A - \alpha(k)I$ has nonnegative entries. Moreover by the assumption that the agents are connected [cf. Assumption 2(b)], for each i, there exists some $j^* \neq i$ such that $a_{j^*}^i > 0$. This and the assumption that all positive weights are at least η [cf. Assumption 2(a)] imply that for each i, there exists some $j^* \neq i$ with $a_{j^*}^i \geq \eta$. Therefore,

$$\sum_{j \neq i} a_j^i \geq a_{j^*}^i \geq \eta \qquad \text{for all } i.$$

Since the vector a^i is stochastic, it follows that

$$a_i^i = 1 - \sum_{j \neq i} a_j^i \leq 1 - \eta < 1.$$

Combined with the assumption that $\alpha(k) \leq a_i^i$, we have

$$\alpha(k) \leq a_i^i < 1 \qquad \text{for all } k,$$

implying that the matrix $B(k)$ has nonnegative entries.

We finally show that $B(k)'$ is a stochastic matrix, i.e., the columns of $B(k)$ are stochastic. We have for every i and k,

$$\sum_{j=1}^{m} [B(k)]_j^i = \frac{1}{1-\alpha(k)} \sum_{j=1}^{m} [A - \alpha(k) I]_j^i = \frac{1}{1-\alpha(k)} \left(\sum_{j=1}^{m} [A]_j^i - \alpha(k) \right).$$

Since A has stochastic columns, it follows that

$$\sum_{j=1}^{m} [B(k)]_j^i = \frac{1}{1-\alpha(k)}\,(1-\alpha(k)) = 1.$$

Using the matrix $B(k)$, we can re-write Eq. (5) as

$$x_i(k+1) - \theta(k+1) = (1-\alpha(k)) \sum_{j=1}^{m} [B(k)]_j^i \,(x_j(k) - \theta(k)) + \alpha(k)\nu_i(k) - \nu(k),$$

showing the desired relation. ■

Our convergence analysis relies on the following lemma. It is based on Lemma 2 of Chapter 2 in Polyak [19], and therefore the proof is omitted.

Lemma 2: *Let $\{u_k\}$ be a scalar sequence such that $u_k \geq 0$ for all k and*

$$u_{k+1} \leq (1-\alpha_k) u_k + \beta_k \qquad \text{for all } k \geq 0.$$

(a) *If $\alpha_k = \alpha$ and $\beta_k = \beta$ for all k and some scalars α and β such that $0 < \alpha \leq 1$ and $\beta \geq 0$, then for all k,*

$$u_k \leq \frac{\beta}{\alpha} + (1-\alpha)^k \left(u_0 - \frac{\beta}{\alpha} \right).$$

(b) *If α_k and β_k satisfy $0 < \alpha_k \leq 1$ and $\beta_k \geq 0$ for all $k \geq 0$, and*

$$\lim_{k \to \infty} \frac{\beta_k}{\alpha_k} = 0, \qquad \sum_{k=0}^{\infty} \alpha_k = \infty,$$

then

$$\lim_{k \to \infty} u_k = 0.$$

In the following proposition, we show that under our assumptions on the weights a^i and the assumption that the signal weights $\alpha(k)$ satisfy $\sum_{k=0}^{\infty} \alpha(k) = \infty$, agent beliefs asymptotically provide unbiased estimates of the underlying state $\theta(k)$.

Proposition 1: Let Assumptions 1, 2, and 3 hold. Let the sequences $\{x^i(k)\}$, $i = 1, \ldots, m$, be generated by Eq. (3). Assume that the signal weight $\alpha(k)$ satisfies

$$0 < \alpha(k) \leq a_i^i \quad \text{for all } i,\ k, \qquad \sum_{k=0}^{\infty} \alpha(k) = \infty.$$

Then, we have

$$\lim_{k\to\infty} E[x_i(k)-\theta(k)] = 0 \qquad \textit{for all } i.$$

Proof: For all $k \geq 0$, we define $M(k) \in \mathbb{R}$ and $\mu(k) \in \mathbb{R}$ as follows:

$$M(k) = \max_{1\leq i\leq m} E[x_i(k)-\theta(k)],$$

$$\mu(k) = \min_{1\leq j\leq m} E[x_j(k)-\theta(k)].$$

By taking the expectation in Eq. (4) and using the assumption that the innovation $\nu(k)$ and the observation noise $\nu_i(k)$ are zero mean for all i and k (cf. Assumption 3), we obtain

$$\begin{aligned} &E[x_i(k+1)-\theta(k+1)] \\ &= (1-\alpha(k)) \sum_{j=1}^{m} [B(k)]_j^i \, E[x_j(k)-\theta(k)] \\ &\leq (1-\alpha(k))M(k), \end{aligned}$$

where the inequality holds since the columns of the matrix $B(k)$ are stochastic (cf. Lemma 1). Similarly, we have for all $k \geq 0$ and all i,

$$(1-\alpha(k))\mu(k) \leq E[x_i(k+1)-\theta(k+1)].$$

From the preceding two relations it follows

$$M(k+1) \leq (1-\alpha(k))M(k), \tag{6}$$

$$(1-\alpha(k))\mu(k) \leq \mu(k+1),$$

implying that

$$M(k+1)-\mu(k+1) \leq (1-\alpha(k))(M(k)-\mu(k)).$$

In view of the assumptions on the signal weight $\alpha(k)$, we can use Lemma 2(b) with the identifications $\beta_k = 0$ and $u_k = M(k)-\mu(k)$ (and therefore $u_k \geq 0$ for all k). This shows that the difference sequence $\{M(k)-\mu(k)\}$ converges to 0 as $k \to \infty$. By Eq. (6), we can use Lemma 2(b) once again with the identification $u_k = |M(k)|$, implying that the sequence $\{M(k)\}$ converges to 0. Thus, the sequence $\{\mu(k)\}$ also converges to 0. Since for all i and $k \geq 0$, we have

$$\mu(k) \leq E[x_i(k)-\theta(k)] \leq M(k),$$

it follows

$$\lim_{k\to\infty} E[x_i(k)-\theta(k)] = 0 \qquad \text{for all } i.$$

■

In the following proposition, we study the mean-square gap between the agent beliefs and the underlying state. In particular, for *persistent innovations* and a constant signal weight (or a constant weight rule), we provide an upper bound on the mean-square gap for each k as a function of the innovation and noise variances, and signal weight used in the learning rule. For *diminishing innovations*, we show that for all i, the gap between agent beliefs and the state $x_i(k)-\theta(k)$ converges to 0 as $k \to \infty$ in the mean-square sense and in probability with diminishing signal weights.

Proposition 2: Let Assumptions 1, 2, and 3 hold. Let the sequences $\{x^i(k)\}$, $i = 1,\ldots,m$, be generated by Eq. (3).

(a) *(Persistent Innovation) Assume that the signal weight $\alpha(k)$ satisfies*

$$\alpha_k = \alpha \quad \textit{for all } k \textit{ and some } \alpha \textit{ such that } 0 < \alpha \leq a_i^i.$$

Assume also that $E[\nu^2(k)] \leq \sigma_n^2$ for all k, i.e., the innovation variance is uniformly bounded for all k. Then for all k and all i,

$$\begin{aligned} &E[(x_i(k)-\theta(k))^2] \\ &\leq \frac{\alpha^2\sigma_a^2+\sigma_n^2}{\alpha} + (1-\alpha)^k \left((\sigma_x^2+\sigma_\theta^2) - \frac{\alpha^2\sigma_a^2+\sigma_n^2}{\alpha} \right), \end{aligned}$$

where σ_θ^2, σ_x^2, and σ_a^2 are given in Assumption 3.

(b) *(Diminishing Innovation) Assume that the signal weight $\alpha(k)$ satisfies*

$$0 < \alpha(k) \leq a_i^i, \quad \alpha(k) \to 0, \quad \sum_{k=0}^{\infty} \alpha(k) = \infty.$$

Assume also that

$$\lim_{k\to\infty} \frac{E[\nu^2(k)]}{\alpha(k)} = 0,$$

i.e., as $k \to \infty$, the innovation variance converges to 0 faster than the signal weight $\alpha(k)$. Then,

$$\lim_{k\to\infty} E[(x_i(k)-\theta(k))^2] = 0 \qquad \textit{for all } i.$$

Moreover, the sequence $\{x_i(k)-\theta(k)\}$ converges to 0 in probability for all i, i.e., for any $\epsilon > 0$, we have

$$\lim_{k\to\infty} \mathrm{P}\{\,|x_i(k)-\theta(k)| \geq \epsilon\,\} = 0 \qquad \textit{for all } i.$$

Proof: By taking the square of Eq. (4) (cf. Lemma 1), we obtain

$$\begin{aligned} &(x_i(k+1)-\theta(k+1))^2 \\ &= (1-\alpha(k))^2 \left(\sum_{j=1}^{m} [B(k)]_j^i \,(x_j(k)-\theta(k)) \right)^2 \\ &\quad +\alpha^2(k) \left(\nu_i(k) - \frac{\nu(k)}{\alpha(k)} \right)^2 + 2\alpha(k)(1-\alpha(k)) \\ &\quad \times \left(\nu_i(k) - \frac{\nu(k)}{\alpha(k)} \right) \sum_{j=1}^{m} [B(k)]_j^i \,(x_j(k)-\theta(k)). \end{aligned}$$

Since the vector $[B(k)]^i$ is stochastic and the function $(\cdot)^2$ is convex, it follows that

$$\left(\sum_{j=1}^{m} [B(k)]_j^i \,(x_j(k)-\theta(k)) \right)^2 \leq \sum_{j=1}^{m} [B(k)]_j^i (x_j(k)-\theta(k))^2.$$

Thus, we have

$$\begin{aligned} &(x_i(k+1)-\theta(k+1))^2 \\ &\leq (1-\alpha(k))^2 \sum_{j=1}^{m} [B(k)]_j^i (x_j(k)-\theta(k))^2 \\ &\quad +\alpha^2(k) \left(\nu_i(k) - \frac{\nu(k)}{\alpha(k)} \right)^2 + 2\alpha(k)(1-\alpha(k)) \end{aligned}$$

$$\times \left(\nu_i(k) - \frac{\nu(k)}{\alpha(k)}\right) \sum_{j=1}^m [B(k)]^i_j \,(x_j(k) - \theta(k)).$$

By taking the expectations of both sides in the preceding relation, we obtain

$$\begin{aligned} &E[(x_i(k+1) - \theta(k+1))^2] \\ &\le (1-\alpha(k))^2 \sum_{j=1}^m [B(k)]^i_j E[(x_j(k) - \theta(k))^2] \\ &\quad + \alpha^2(k) E\left[\left(\nu_i(k) - \frac{\nu(k)}{\alpha(k)}\right)^2\right] \\ &\quad + 2\alpha(k)(1-\alpha(k)) \\ &\quad \times E\left[\left(\nu_i(k) - \frac{\nu(k)}{\alpha(k)}\right) \sum_{j=1}^m [B(k)]^i_j \,(x_j(k) - \theta(k))\right]. \end{aligned}$$

We now compute each of the expectations. Let us introduce the notation

$$E[\nu^2(k)] = \sigma^2(k), \qquad E[\nu_i^2(k)] = \sigma_i^2(k) \qquad \text{for all } k.$$

Using the independence of $\nu(k)$ and $\nu_i(k)$, we have

$$E\left[\left(\nu_i(k) - \frac{\nu(k)}{\alpha(k)}\right)^2\right] = \sigma_i^2(k) + \frac{\sigma^2(k)}{\alpha^2(k)}.$$

The random variables $\nu_i(k)$ and $x_j(k)-\theta(k)$ are independent for all j. Similarly, the random variables $\nu(k)$ and $x_j(k) - \theta(k)$ are independent for all j. Hence,

$$E\left[\left(\nu_i(k) - \frac{\nu(k)}{\alpha(k)}\right) \sum_{j=1}^m [B(k)]^i_j \,(x_j(k) - \theta(k))\right] = 0.$$

Combining the preceding four relations, we obtain

$$\begin{aligned} &E[(x_i(k+1) - \theta(k+1))^2] \\ &\le (1-\alpha(k))^2 \sum_{j=1}^m [B(k)]^i_j E[(x_j(k) - \theta(k))^2] \\ &\quad + \alpha^2(k)\sigma_i^2(k) + \sigma^2(k). \end{aligned}$$

By using the assumption that $E[\nu_i^2(k)] = \sigma_i^2(k) \le \sigma_a^2$ [cf. Assumption 3(a)], this implies

$$\begin{aligned} &E[(x_i(k+1) - \theta(k+1))^2] \\ &\le (1-\alpha(k))^2 \sum_{j=1}^m [B(k)]^i_j E[(x_j(k) - \theta(k))^2] \\ &\quad + \alpha^2(k)\sigma_a^2 + \sigma^2(k) \\ &\le (1-\alpha(k)) \max_{1\le j\le m} E[(x_j(k) - \theta(k))^2] \\ &\quad + \alpha^2(k)\sigma_a^2 + \sigma^2(k), \end{aligned}$$

where the second inequality follows since $\alpha(k) \le 1$ for both parts of the proposition. Taking the maximum over all $i = 1, \ldots, m$ in the preceding relation yields

$$\begin{aligned} &\max_{1\le i\le m} E[(x_i(k+1) - \theta(k+1))^2 \\ &\le (1-\alpha(k)) \max_{1\le j\le m} E[(x_j(k) - \theta(k))^2] \\ &\quad + \alpha^2(k)\sigma_a^2 + \sigma^2(k). \end{aligned}$$

We can now use Lemma 2 with the identifications

$$u_k = \max_{1\le j\le m} E[(x_j(k) - \theta(k))^2], \quad \alpha_k = \alpha(k),$$

$$\beta_k = \alpha^2(k)\sigma_a^2 + \sigma^2(k).$$

Under the assumption that $\alpha(k) = \alpha$ and $\sigma^2(k) \le \sigma_n^2$, we can take $\beta_k = \alpha^2\sigma_a^2 + \sigma_n^2$ for all k. It follows from Lemma 2(a) that

$$\begin{aligned} &\max_{1\le j\le m} E[(x_j(k) - \theta(k))^2] \\ &\le \frac{\alpha^2\sigma_a^2 + \sigma_n^2}{\alpha} \\ &\quad + (1-\alpha)^k \left(\max_{1\le j\le m} E[(x_j(0) - \theta(0))^2] - \frac{\alpha^2\sigma_a^2 + \sigma_n^2}{\alpha}\right). \end{aligned}$$

Using the assumption that $\theta(0)$ and $x_j(0)$ are independent, and $E[\theta^2(0)] \le \sigma_\theta^2$ and $E[x_j^2(0)] \le \sigma_x^2$ for all j [cf. Assumption 3(c)], we have

$$\max_{1\le j\le m} E[(x_j(0) - \theta(0))^2] \le \sigma_x^2 + \sigma_\theta^2.$$

Combining the preceding two relations, we obtain

$$\begin{aligned} E[(x_i(k) - \theta(k))^2] \;\le\; & \frac{\alpha^2\sigma_a^2 + \sigma_n^2}{\alpha} \\ & + (1-\alpha)^k \left((\sigma_x^2 + \sigma_\theta^2) - \frac{\alpha^2\sigma_a^2 + \sigma_n^2}{\alpha}\right), \end{aligned}$$

for all i, establishing part (a).

Under the assumptions of part (b) on the signal weight and the innovation variance $E[\nu^2(k)] = \sigma^2(k)$, we have

$$\lim_{k\to\infty} \frac{\beta_k}{\alpha_k} = \frac{\alpha^2(k)\sigma_a^2 + \sigma^2(k)}{\alpha(k)} = 0.$$

Therefore, part (b) of Lemma 2 applies and shows that

$$\lim_{k\to\infty} \max_{1\le j\le m} E[(x_j(k) - \theta(k))^2] = 0,$$

establishing the first result of part (b).

We finally show that $x_i(k) - \theta(k)$ converges to 0 in probability for all i. Since $\sqrt{t}$ is a concave function for $t \ge 0$, we have for all i and k,

$$\begin{aligned} E[\,|x_i(k) - \theta(k)|\,] &= E[\sqrt{(x_i(k) - \theta(k))^2}] \\ &\le \sqrt{E[(x_i(k) - \theta(k))^2]}, \end{aligned}$$

where the inequality follows by Jensen's inequality. By Proposition 2(b), it follows that

$$\lim_{k\to\infty} E[\,|x_i(k) - \theta(k)|\,] = 0 \qquad \text{for all } i. \tag{7}$$

The Markov Inequality states that for any nonnegative random variable Y with a finite mean $E[Y]$, the probability that the random variable Y exceeds any given scalar $\epsilon > 0$ satisfies

$$\mathrm{P}\{Y \ge \epsilon\} \le \frac{E[Y]}{\epsilon},$$

where $\mathrm{P}\{\mathcal{A}\}$ denotes the probability of a random event $\mathcal{A}$. By applying the Markov inequality to the nonnegative random variable $|x_i(k) - \theta(k)|$, which has a finite expectation by Eq. (7), we obtain

$$\mathrm{P}\{|x_i(k) - \theta(k)| \geq \epsilon\} \leq \frac{E[\,|x_i(k) - \theta(k)|\,]}{\epsilon} \qquad \text{for all } i \text{ and } k,$$

which in view of the relation (7) yields

$$\lim_{k\to\infty} \mathrm{P}\{|x_i(k) - \theta(k)| \geq \epsilon\} = 0 \qquad \text{for all } i.$$

The preceding relation holds for any $\epsilon > 0$, thus showing that $|x_i(k) - \theta(k)|$ converges to 0 in probability for each i. ■

IV. Conclusions

In this paper, we studied the dynamics of rule-of-thumb learning over a social network. The social network determines the communication structure among agents with heterogeneous information and beliefs. Agents update their beliefs from signals and communication with their neighbors. Motivated by a range of real-world social learning problems, we study an environment in which the underlying state changes over time. We assume that agents update their beliefs as a time-varying weighted average of their prior beliefs, the signal they receive and the communication of their neighbors. We show that beliefs generated according to this general class of learning rules are unbiased and we study the convergence of these beliefs to the true underlying state. When the underlying state changes with diminishing innovations (asymptotically to zero), update rules with diminishing weights, which reduce the weight given to signals ensure mean square convergence and convergence in probability. In contrast, update rules with constant weights typically fail to do so. However, when agents are trying to learn an underlying state with persistent innovations, the asymptotic behavior of constant weight rules is superior. For these rules, we provide explicit bounds on the mean square (variance) on the gap between individual beliefs and the true underlying state.

Though both our focus on learning a time-varying state and our method of analysis are novel, the results are intuitive and suggests clear directions for future research. In particular, our results suggest that relatively uncomplicated learning rules can ensure learning by the agents under a wide range of network topologies. The intuition for why different types of rules perform better depending on the nature of the innovations to the true underlying state can best be understood by going back to the normal updating case. The normal updating formula for the estimation of an unknown, but fixed parameter is a special case of our diminishing weight rules, whereas the updating formula for the estimation of a time varying parameter corresponds to constant weight rules. This intuition also highlights that we may indeed expect agents to use rule-of-thumb rules with diminishing weights when the state is fixed and with constant weights when the state is time varying. Our analysis therefore suggests that a limited amount of rationality is sufficient for learning and for efficient exchange of information over the network.

Our analysis also opens the way for a more systematic study of simple learning rules in social network settings. In particular, our results can be extended to provide convergence rates and can be used to develop estimates over learning in non-asymptotic environments. Other important areas for investigation, which are part of our ongoing work, include developing an explicit characterization of how learning behavior depends on network topology and the connection structure and also investigating the conditions under which beliefs converge to the underlying state (with diminishing innovations) almost surely.

References

[1] D. Acemoglu, Munther Dahleh, Ilan Lobel, and A. Ozdaglar, *Bayesian learning in social networks*, Preprint, 2008.

[2] V. Bala and S. Goyal, *Learning from neighbours*, Review of Economic Studies **65** (1998), no. 3, 595–621.

[3] ———, *Conformism and diversity under social learning*, Economic Theory **17** (2001), 101–120.

[4] A. Banerjee, *A simple model of herd behavior*, Quarterly Journal of Economics **107** (1992), 797–817.

[5] S. Bikchandani, D. Hirshleifer, and I. Welch, *A theory of fads, fashion, custom, and cultural change as information cascades*, Journal of Political Economy **100** (1992), 992–1026.

[6] V.D. Blondel, J.M. Hendrickx, A. Olshevsky, and J.N. Tsitsiklis, *Convergence in multiagent coordination, consensus, and flocking*, Proceedings of IEEE CDC, 2005.

[7] P.M. DeMarzo, D. Vayanos, and J. Zwiebel, *Persuasion bias, social influence, and unidimensional opinions*, The Quarterly Journal of Economics **118** (2003), no. 3, 909–968.

[8] G. Ellison and D. Fudenberg, *Rules of thumb for social learning*, Journal of Political Economy **101** (1993), 612–643.

[9] ———, *Word-of-mouth communication and social learning*, The Quarterly Journal of Economics **110** (1995), 93–126.

[10] A. Foster and M. Rosenzweig, *Learning by doing and learning from others: Human capital and technical change in agriculture*, The Journal of Political Economy **103** (1995), no. 6, 1176–1209.

[11] B. Golub and M. O. Jackson, *Naive learning in social networks: Convergence, influence, and the wisdom of crowds*, Preprint, 2007.

[12] M. Granovetter, *The strength of weak ties*, The American Journal of Sociology **78** (1973), no. 6, 1360–1380.

[13] M. Ioannides and L. Loury, *Job information networks, neighborhood effects, and inequality*, The Journal of Economic Literature **42** (2004), no. 2, 1056–1093.

[14] A. Jadbabaie, J. Lin, and S. Morse, *Coordination of groups of mobile autonomous agents using nearest neighbor rules*, IEEE Transactions on Automatic Control **48** (2003), no. 6, 988–1001.

[15] K. Munshi, *Social learning in a heterogeneous population: Technology diffusion in the indian green revolution*, The Journal of Development Economics **73** (2004), no. 1, 185–215.

[16] A. Nedić and A. Ozdaglar, *Distributed subgradient methods for multi-agent optimization*, forthcoming in IEEE Transactions on Automatic Control, 2008.

[17] R. Olfati-Saber and R.M. Murray, *Consensus problems in networks of agents with switching topology and time-delays*, IEEE Transactions on Automatic Control **49** (2004), no. 9, 1520–1533.

[18] A. Olshevsky and J.N. Tsitsiklis, *Convergence speed in distributed consensus and averaging*, forthcoming in SIAM Journal on Control and Optimization, 2008.

[19] B.T. Polyak, *Introduction to optimisation*, Optimization Software, Inc., New York, 1987.

[20] L. Smith and P. Sorensen, *Pathological outcomes of observational learning*, Econometrica **68** (2000), no. 2, 371–398.

[21] J.N. Tsitsiklis, *Problems in decentralized decision making and computation*, Ph.D. thesis, Dept. of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, 1984.

[22] J.N. Tsitsiklis, D.P. Bertsekas, and M. Athans, *Distributed asynchronous deterministic and stochastic gradient optimization algorithms*, IEEE Transactions on Automatic Control **31** (1986), no. 9, 803–812.