

Gaussian Learning-Without-Recall in a Dynamic Social Network

Chu Wang¹ and Bernard Chazelle²

Abstract—We analyze the dynamics of the Learning-Without-Recall model with Gaussian priors in a dynamic social network. Agents seeking to learn the state of the world, the “truth”, exchange signals about their current beliefs across a changing network and update them accordingly. The agents are assumed memoryless and rational, meaning that they Bayes-update their beliefs based on current states and signals, with no other information from the past. The other assumption is that each agent hears a noisy signal from the truth at a frequency bounded away from zero. Under these conditions, we show that the system reaches truthful consensus almost surely with a convergence rate that is polynomial in expectation. Somewhat paradoxically, high outdegree can slow down the learning process. The lower-bound assumption on the truth-hearing frequency is necessary: even infinitely frequent access to the truth offers no guarantee of truthful consensus in the limit.

I. INTRODUCTION

People typically form opinions by updating their current beliefs and reasons in response to new signals from other sources (friends, colleagues, social media, newspapers, etc.) [1], [2], [3]. Suppose there were an information source that made a noisy version of the “truth” available to agents connected through a social network. Under which conditions would the agents reach consensus about their beliefs? What would ensure truthful consensus (meaning that the consensus coincided with the truth)? How fast would it take for the process to converge? To address these questions requires agreeing on a formal model of distributed learning. Fully rational agents update their beliefs by assuming a prior and using Bayes’ rule to integrate all past information available to them [4], [5], [6], [7], [8], [9]. Full rationality is intractable in practice [10], [11], so much effort has been devoted to developing computationally effective mechanisms, including non- (or partially) Bayesian methods [12], [10], [3], [13], [14]. Much of this line of work can be traced back to the seminal work of DeGroot [15] on linear opinion pooling.

This paper is no exception. Specifically, it follows the *Bayesian-Without-Recall* (BWR) model recently proposed by Rahimian and Jadbabaie in [11]; see also [16], [17], [18]. The agents are assumed to be memoryless and rational: this means that they use Bayesian updates based on current beliefs and signals with no other information from the past. The process is local in that agents can collect information only from their neighbors in a directed graph. In this work, the graph is allowed to change at each time step. The BWR

model seeks to capture the benefits of rational behavior while keeping both the computation and the information stored to a minimum [18].

A distinctive feature of our work is that the social network need *not* be fixed once and for all. The ability to modify the communication channels over time reflects the inherently changing nature of social networks as well as the reality that our contacts do not all speak to us at once. Thus even if the underlying network is fixed over long timescales, the model allows for agents to be influenced by selected subsets of their neighbors. Dynamic networks are common occurrences in opinion dynamics [19], [20], [21], [22] but, to our knowledge, somewhat new in the context of social learning.

Our working model in this paper posits a Gaussian setting: the likelihoods and initial priors of the agents are normal distributions. During the learning process, signals are generated as noisy measurements of agents’ beliefs and the noise is assumed normal and unbiased. Thus all beliefs remain Gaussian at all times [11], [23].

Our main result is that, under the assumption that each agent hears a noisy signal from the truth at a frequency bounded away from zero, the system reaches truthful consensus almost surely with a convergence rate polynomial in expectation. Specifically, we show that, as long as each agent receives a signal from the truth at least once every $1/\gamma$ steps, the convergence rate is $O(t^{-\gamma/2d})$, where d is the maximum node outdegree.

Somewhat paradoxically, high outdegree can slow down learning. The reason is that signals from peer agents are imperfect conveyors of the truth and can, on occasion, contaminate the network with erroneous information; this finding is in line with a similar phenomenon uncovered by Harel et al. [24], in which social learning system with two Bayesian agents is found to be hindered by increased interaction between the agents. We note that our lower-bound assumption on the truth-hearing frequency is necessary: even infinitely frequent access to the truth is not enough to achieve truthful consensus in the limit.

Further background. Researchers have conducted empirical evaluations of both Bayesian and non-Bayesian models [25], [26], [27], [28]. In [29], Mossel et al. analyzed a Bayesian learning system in which each agent gets signals from the truth only once at the beginning and then interact with other agents using Gaussian estimators. In [30], Moscarini et al. considered social learning in a model where the truth is not fixed but is, instead, supplied by a Markov chain. In the different but related realm of iterated learning,

¹Nokia Bell Labs, 600-700 Mountain Avenue, Murray Hill, New Jersey 07974, chu.wang@nokia.com

²Department of Computer Science, Princeton University, 35 Olden Street, Princeton, New Jersey 08540, chazelle@cs.princeton.edu

agents learn from ancestors and teach descendants. The goal is to pass on the truth through generations while seeking to prevent information loss [31], [32].

Organization. Section II introduces the model and the basic formulas for single-step belief updates. Section III investigates the dynamics of the beliefs in expectation and derive the polynomial upper bound on the convergence rate under the assumption that each agent hears a signal from the truth at a frequency bounded away from zero. We demonstrate the necessity of this assumption in Section IV and prove that the convergence occurs almost surely.

II. PRELIMINARIES

A. The Model

We choose the real line \mathbb{R} as the state space and we denote the agents by $1, 2, \dots, n$; for convenience, we add an extra agent, labeled 0, whose belief is a fixed number, unknown to others, called the *truth*. At time $t = 0, 1, \dots$, the belief of agent i is a probability distribution over the state space \mathbb{R} , which is denoted by $\mu_{t,i}$. We assume that the initial belief $\mu_{0,i}$ of agent i is Gaussian: $\mu_{0,i} \sim \mathcal{N}(x_{0,i}, \sigma_{0,i}^2)$. Without loss of generality, we assume the truth is a constant (single-point distribution: $\mu_{t,0} = 0$; $\sigma_{t,0} = 0$ for all t) and the standard deviation is the same for all other agents, ie, $\sigma_{0,i} = \sigma_0 > 0$ for $i > 0$.

The interactions between agents are modeled by an infinite sequence $(G_t)_{t \geq 0}$, where each G_t is a directed graph over the node set $\{0, \dots, n\}$. An edge pointing from i to j in G_t indicates that i receives data from j at time t . Typically, the sequence of graphs is specified ahead of time or it is chosen randomly: the only condition that matters is that it should be independent of the randomness used in the learning process; specifically, taking expectations and variances of the random variables that govern the dynamics will assume a fixed graph sequence (possibly random). Because agent 0 holds the truth, no edge points away from it. The adjacency matrix of G_t is denoted by A_t : it is an $(n+1) \times (n+1)$ matrix whose first row is $(1, 0, \dots, 0)$.

B. Information Transfer

At time $t \geq 0$, each agent $i > 0$ samples a state $\theta_{t,i} \in \mathbb{R}$ consistent with her own belief: $\theta_{t,i} \sim \mu_{t,i}$. A noisy measurement $a_{t,i} = \theta_{t,i} + \varepsilon_{t,i}$ is then sent to each agent j such that $(A_t)_{ji} = 1$. All the noise terms $\varepsilon_{t,i}$ are sampled *iid* from $\mathcal{N}(0, \sigma^2)$. An equivalent formulation is to say that the likelihood function $l(a|\theta)$ is drawn from $\mathcal{N}(\theta, \sigma^2)$. In our setting, agent i sends the same data to all of her neighbors; this is done for notational convenience and the same results would still hold if we were to resample independently for each neighbor. Except for the omission of explicit utilities and actions, our setting is easily identified as a variant of the BWR model [11].

C. Updating Beliefs

A single-step update for agent $i > 0$ consists of setting $\mu_{t+1,i}$ as the posterior $\mathbb{P}[\mu_{t,i}|d] \propto \mathbb{P}[d|\mu_{t,i}]\mathbb{P}[\mu_{t,i}]$, where d is the data from the neighbors of i received at time t . Plugging in the corresponding Gaussians gives us the classical update rules from Bayesian inference [23]. Updated beliefs remain Gaussian so we can use the notation $\mu_{t,i} \sim \mathcal{N}(x_{t,i}, \tau_{t,i}^{-1})$, where $\tau_{t,i}$ denotes the precision $\sigma_{t,i}^{-2}$. Writing $\tau = \sigma^{-2}$ and letting $d_{t,i}$ denote the outdegree of i in G_t , for any $i > 0$ and $t \geq 0$,

$$\begin{cases} x_{t+1,i} = (\tau_{t,i}x_{t,i} + \tau a_1 + \dots + \tau a_{d_{t,i}})/(\tau_{t,i} + d_{t,i}\tau); \\ \tau_{t+1,i} = \tau_{t,i} + d_{t,i}\tau, \end{cases} \quad (1)$$

where $a_1, \dots, a_{d_{t,i}}$ are the signals received by agent i from its neighbors at time t .

D. Expressing the Dynamics in Matrix Form

Let D_t and P_t denote the $(n+1)$ -by- $(n+1)$ diagonal matrices $\text{diag}(d_{t,i})$ and $(\tau_0/\tau)I + \sum_{k=0}^{t-1} D_k$, respectively, where I is the identity matrix and the sum is 0 for $t = 0$. It follows from (1) that $\mu_{t,i} \sim \mathcal{N}(x_{t,i}, (\tau P_t)_{ii}^{-1})$ for $i > 0$. Regrouping the means in vector form, $\mathbf{x}_t := (x_{t,0}, \dots, x_{t,n})^T$, where $x_{t,0} = 0$ and $x_{0,1}, \dots, x_{0,n}$ are given as inputs, we have

$$\mathbf{x}_{t+1} = (P_t + D_t)^{-1} (P_t \mathbf{x}_t + A_t (\mathbf{x}_t + \mathbf{u}_t + \boldsymbol{\varepsilon}_t)), \quad (2)$$

where \mathbf{u}_t is such that $u_{t,0} \sim \mathcal{N}(0, 0)$ and, for $i > 0$, $u_{t,i} \sim \mathcal{N}(0, (\tau(P_t)_{ii})^{-1})$; and $\boldsymbol{\varepsilon}_t$ is such that $\varepsilon_{t,0} \sim \mathcal{N}(0, 0)$ and, for $i > 0$, $\varepsilon_{t,i} \sim \mathcal{N}(0, 1/\tau)$. We refer to the vectors \mathbf{x}_t and $\mathbf{y}_t := \mathbb{E} \mathbf{x}_t$ as the *mean process* and the *expected mean process*, respectively. Taking expectations on both sides of (2) with respect to the random vectors \mathbf{u}_t and $\boldsymbol{\varepsilon}_t$ yields the update rule for the expected mean process: $\mathbf{y}_0 = \mathbf{x}_0$ and, for $t > 0$,

$$\mathbf{y}_{t+1} = (P_t + D_t)^{-1} (P_t + A_t) \mathbf{y}_t. \quad (3)$$

A key observation is that $(P_t + D_t)^{-1} (P_t + A_t)$ is a stochastic matrix, so the expected mean process \mathbf{y}_t forms a diffusive influence system [22]: the vector evolves by taking convex combinations of its own coordinates. What makes the analysis different from standard multiagent agreement systems is that the weights vary over time. In fact, some weights typically tend to 0, which violates one of the cardinal assumptions used in the analysis of averaging systems [22], [33]. This leads us to the use of arguments, such as fourth-order moment bounds, that are not commonly encountered in this area.

E. Our Results

The belief vector $\boldsymbol{\mu}_t$ is Gaussian with mean \mathbf{x}_t and covariance matrix Σ_t formed by zeroing out the top-left element of $(\tau P_t)^{-1}$. We say that the system reaches *truthful consensus* if both the mean process \mathbf{x}_t and the covariance matrix tend to zero as t goes to infinity. This indicates that all the agents' beliefs share a common mean equal to the truth and the "error bars" vanish over time. In view of (1), the

covariance matrix indeed tends to 0 as long as the degrees are nonzero infinitely often, a trivial condition. To establish truthful consensus, therefore, boils down to studying the mean process \mathbf{x}_t . We do this in two parts: first, we show that the expected mean process converges to the truth; then we prove that fluctuations around it eventually vanish almost surely.¹

Truth-hearing assumption: Given any interval of length $\kappa := \lfloor 1/\gamma \rfloor$, every agent $i > 0$ has an edge $(i, 0)$ in G_t for at least one value of t in that interval.

THEOREM 2.1: Under the truth-hearing assumption, the system reaches truthful consensus with a convergence rate bounded by $O(t^{-\gamma/2d})$, where d is the maximum outdegree over all the networks.

We prove the theorem in the next two sections. It will follow directly from Lemmas 3.1 and 4.1 below. The convergence rate can be improved to the order of $t^{-(1-\varepsilon)\gamma/d}$, for arbitrarily small $\varepsilon > 0$. The inverse dependency on γ is not surprising: the more access to the truth the stronger the attraction to it. On the other hand, it might seem counterintuitive that a larger outdegree should slow down convergence. This illustrates the risk of groupthink. It pays to follow the crowds when the crowds are right. When they are not, however, this distracts from the lonely voice that happens to be right.

How essential is the truth-hearing assumption? We show that it is necessary. Simply having access to the truth infinitely often is not enough to achieve truthful consensus.

F. Useful Matrix Inequalities

We highlight certain matrix inequalities to be used throughout. We use the standard element-wise notation $R \leq S$ to indicate that $R_{ij} \leq S_{ij}$ for all i, j . The infinity norm $\|R\|_\infty = \max_i \sum_j |r_{ij}|$ is submultiplicative: $\|RS\|_\infty \leq \|R\|_\infty \|S\|_\infty$, for any matching rectangular matrices. On the other hand, the max-norm $\|R\|_{\max} := \max_{i,j} |r_{ij}|$ is not, but it is transpose-invariant and also satisfies: $\|RS\|_{\max} \leq \|R\|_\infty \|S\|_{\max}$. It follows that

$$\begin{aligned} \|RSR^T\|_{\max} &\leq \|R\|_\infty \|SR^T\|_{\max} = \|R\|_\infty \|RS^T\|_{\max} \\ &\leq \|R\|_\infty^2 \|S^T\|_{\max} = \|R\|_\infty^2 \|S\|_{\max}. \end{aligned} \quad (4)$$

III. THE EXPECTED MEAN PROCESS DYNAMICS

We analyze the convergence of the mean process in expectation. The expected mean $\mathbf{y}_t = \mathbb{E} \mathbf{x}_t$ evolves through an averaging process entirely determined by the initial value $\mathbf{y}_0 = (0, x_{0,1}, \dots, x_{0,n})^T$ and the graph sequence G_t . Intuitively, if an agent communicates repeatedly with a holder of the truth, the weight of the latter should accumulate and

¹The Kullback-Leibler divergence [12] is not suitable here because the estimator is Gaussian, hence continuous, whereas the truth is a single-point distribution.

increasingly influence the belief of the agent in question. Our goal in this section is to prove the following result:

LEMMA 3.1: Under the truth-hearing assumption, the expected mean process \mathbf{y}_t converges to the truth asymptotically. If, at each step, no agent receives information from more than d agents, then the convergence rate is bounded by $Ct^{-\gamma/2d}$, where C is a constant that depends on $\mathbf{x}_0, \gamma, d, \sigma_0/\sigma$.

Proof. We define B_t as the matrix formed by removing the first row and the first column from the stochastic $P_{t+1}^{-1}(P_t + A_t)$. If we write \mathbf{y}_t as $(0, \mathbf{z}_t)$ then, by (3),

$$\begin{pmatrix} 0 \\ \mathbf{z}_{t+1} \end{pmatrix} = \begin{pmatrix} 1 & \mathbf{0} \\ \boldsymbol{\alpha}_t & B_t \end{pmatrix} \begin{pmatrix} 0 \\ \mathbf{z}_t \end{pmatrix}, \quad (5)$$

where $\boldsymbol{\alpha}_{t,i} = (P_{t+1}^{-1})_{ii}$ if there is an edge $(i, 0)$ at time t and $\boldsymbol{\alpha}_{t,i} = 0$ otherwise. This further simplifies to

$$\mathbf{z}_{t+1} = B_t \mathbf{z}_t. \quad (6)$$

Let $\mathbf{1}$ be the all-one column vector of length n . Since $P_{t+1}^{-1}(P_t + A_t)$ is stochastic,

$$\boldsymbol{\alpha}_t + B_t \mathbf{1} = \mathbf{1} \quad (7)$$

In matrix terms, the truth-hearing assumption means that, for any $t \geq 0$,

$$\boldsymbol{\alpha}_t + \boldsymbol{\alpha}_{t+1} + \dots + \boldsymbol{\alpha}_{t+\kappa-1} \geq Q_{t+\kappa}^{-1} \mathbf{1}, \quad (8)$$

where Q_t is the matrix derived from P_t by removing the first row and the last column; the inequality relies on the fact that P_t is monotonically nondecreasing. For any $t > s \geq 0$, we define the product matrix $B_{t:s}$ defined as

$$B_{t:s} := B_{t-1} B_{t-2} \dots B_s, \quad (9)$$

with $B_{t:t} = I$. By (6), for any $t > s \geq 0$,

$$\mathbf{z}_t = B_{t:s} \mathbf{z}_s. \quad (10)$$

To bound the infinity norm of $B_{t:0}$, we observe that, for any $0 \leq l < \kappa - 1$, the i -th diagonal element of $B_{s+\kappa:s+l+1}$ is lower-bounded by

$$\begin{aligned} \prod_{j=l+1}^{\kappa-1} (B_{s+j})_{ii} &= \prod_{j=l+1}^{\kappa-1} \frac{(P_{s+j} + A_{s+j})_{ii}}{(P_{s+j+1})_{ii}} \\ &\geq \prod_{j=l+1}^{\kappa-1} \frac{(P_{s+j})_{ii}}{(P_{s+j+1})_{ii}} = \frac{(P_{s+l+1})_{ii}}{(P_{s+\kappa})_{ii}} \geq \frac{(P_s)_{ii}}{(P_{s+\kappa})_{ii}}. \end{aligned} \quad (11)$$

The inequalities follow from the nonnegativity of the entries and the monotonicity of $(P_t)_{ii}$. Note that (11) also holds for $l = \kappa - 1$ since $(B_{s+\kappa:s+\kappa})_{ii} = 1$.

Since $P_{t+1}^{-1}(P_t + A_t)$ is stochastic, the row-sum of B_t does not exceed 1; therefore, by premultiplying B_{s+1}, B_{s+1}, \dots on both sides of (7), we obtain:

$$B_{s+\kappa:s} \mathbf{1} \leq \mathbf{1} - \sum_{l=0}^{\kappa-1} B_{s+\kappa:s+l+1} \boldsymbol{\alpha}_{s+l}. \quad (12)$$

Noting that $\|B_t\|_\infty = \|B_t \mathbf{1}\|_\infty$ for any t , as B_t is non-negative, we combine (8), (11), and (12) together to derive:

$$\|B_{s+\kappa:s}\|_\infty \leq 1 - \min_{i>0} \frac{(P_s)_{ii}}{(P_{s+\kappa})_{ii}}. \quad (13)$$

Let $d := \max_{t \geq 0} \max_{1 \leq i \leq n} d_{t,i}$ denote the maximum outdegree in all the networks, and define $\delta = \min\{\tau_0/\tau, 1\}$. For any $i > 0$ and $s \geq \kappa$,

$$\frac{s\delta}{\kappa} \leq (P_s)_{ii} \leq ds + \frac{\tau_0}{\tau}; \quad (14)$$

hence,

$$\max_i (P_{s+\kappa})_{ii} \leq d(s + \kappa) + \frac{\tau_0}{\tau}. \quad (15)$$

It follows that

$$\frac{(P_{s+\kappa})_{ii} - (P_s)_{ii}}{(P_{s+\kappa})_{ii}} = \frac{\sum_{l=0}^{\kappa-1} d_{s+l,i}}{(P_{s+\kappa})_{ii}} \leq \frac{d\kappa^2\delta^{-1}}{s + \kappa}. \quad (16)$$

Thus, we have

$$\begin{aligned} \min_{i>0} \frac{(P_s)_{ii}}{(P_{s+\kappa})_{ii}} &= 1 - \max_{i>0} \frac{(P_{s+\kappa})_{ii} - (P_s)_{ii}}{(P_{s+\kappa})_{ii}} \\ &\geq 1 - \frac{d\kappa^2\delta^{-1}}{s + \kappa}. \end{aligned} \quad (17)$$

We can replace the upper bound of (13) by

$$1 - \frac{1}{\max_{i>0} (P_{s+\kappa})_{ii}} \min_{i>0} \frac{(P_s)_{ii}}{(P_{s+\kappa})_{ii}},$$

which, together with (15) and (17) gives us

$$\begin{aligned} \|B_{s+\kappa:s}\|_\infty &\leq 1 - \frac{1}{d(s + \kappa) + \tau_0/\tau} \left(1 - \frac{d\kappa^2\delta^{-1}}{s + \kappa}\right) \\ &\leq 1 - \frac{1}{2d\kappa(m + 2)}. \end{aligned} \quad (18)$$

The latter inequality holds as long as $s = m\kappa > 0$ and

$$m \geq m^* := \frac{2d\kappa}{\delta} + \frac{\tau_0}{d\kappa\tau}.$$

It follows that, for $m_0 \geq m^*$,

$$\begin{aligned} \|B_{(m_0+m)\kappa:m_0\kappa}\|_\infty &\leq \prod_{j=2}^{m+1} \left(1 - \frac{1}{2d\kappa(m_0 + j)}\right) \\ &\leq \exp \left\{ -\frac{1}{2d\kappa} \sum_{j=2}^{m+1} \frac{1}{m_0 + j} \right\}. \end{aligned} \quad (19)$$

The matrices B_t are sub-stochastic so that

$$\|B_t \mathbf{z}\|_\infty \leq \|B_t\|_\infty \|\mathbf{z}\|_\infty \leq \|\mathbf{z}\|_\infty.$$

By (10), for any $t \geq (m_0 + m)\kappa$,

$$\mathbf{z}_t = B_{t:(m_0+m)\kappa} B_{(m_0+m)\kappa:m_0\kappa} \mathbf{z}_{m_0},$$

so that, by using standard bounds for the harmonic series, $\ln(k+1) < 1 + \frac{1}{2} + \dots + \frac{1}{k} \leq 1 + \ln k$, we find that

$$\begin{aligned} \|\mathbf{z}_t\|_\infty &\leq \|B_{(m_0+m)\kappa:m_0\kappa}\|_\infty \|\mathbf{z}_{m_0}\|_\infty \\ &\leq \|B_{(m_0+m)\kappa:m_0\kappa}\|_\infty \|\mathbf{z}_0\|_\infty \\ &\leq C t^{-1/(2d\kappa)}, \end{aligned}$$

where $C > 0$ depends on $\mathbf{z}_0, \kappa, d, \tau_0/\tau$. We note that the convergence rate can be improved to the order of $t^{-(1-\varepsilon)\tau/d}$, for arbitrarily small $\varepsilon > 0$, by working a little harder with (18). \square

IV. THE MEAN PROCESS DYNAMICS

Recall that $\mu_{t,i} \sim \mathcal{N}(x_{t,i}, \tau_{t,i}^{-1})$, where $\tau_{t,i}$ denotes the precision $\sigma_{t,i}^{-2}$. A key observation about the updating rule in (1) is that the precision $\tau_{t,i}$ is entirely determined by the graph sequence G_t and is independent of the actual dynamics. Adding to this the connectivity property implied by the truth-hearing assumption, we find immediately that $\tau_{t,i} \rightarrow \infty$ for any agent i . This ensures that the covariance matrix Σ_t tends to 0 as t goes to infinity, which satisfies the second criterion for truthful consensus. The first criterion requires that the mean process \mathbf{x}_t should converge to the truth $\mathbf{0}$. Take the vector $\mathbf{x}_t - \mathbf{y}_t$ and remove the first coordinate $(\mathbf{x}_t - \mathbf{y}_t)_0$ to form the vector $\Delta_t \in \mathbb{R}^n$. Under the truth-hearing assumption, we have seen that $\mathbf{y}_t \rightarrow \mathbf{0}$ (Lemma 3.1), so it suffices to prove the following:

LEMMA 4.1: Under the truth-hearing assumption, the deviation Δ_t vanishes almost surely.

Proof. We use a fourth-moment argument. The justification for the high order is technical: it is necessary to make a certain ‘‘deviation power’’ series converge. By (2), \mathbf{x}_t is a linear combination of independent Gaussian random vectors \mathbf{u}_s and $\boldsymbol{\varepsilon}_s$ for $0 \leq s \leq t-1$, and thus \mathbf{x}_t itself is a Gaussian random vector. Therefore Δ_t is also Gaussian and its mean is zero. From Markov’s inequality, for any $c > 0$,

$$\sum_{t \geq 0} \mathbb{P}[|\Delta_{t,i}| \geq c] \leq \sum_{t \geq 0} \frac{\mathbb{E} \Delta_{t,i}^4}{c^4}. \quad (20)$$

If we are able to show the right hand side of (20) is finite for any $c > 0$, then, by the Borel-Cantelli lemma, with probability one, the event $|\Delta_{t,i}| \geq c$ occurs only a finite number of times, and so $\Delta_{t,i}$ goes to zero almost surely. Therefore, we only need to analyze the order of the fourth moment $\mathbb{E} \Delta_{t,i}^4$. By subtracting (3) from (2), we have:

$$\Delta_{t+1} = B_t \Delta_t + M_t \mathbf{v}_t, \quad (21)$$

where $\mathbf{v}_t := \mathbf{u}_t + \boldsymbol{\varepsilon}_t$ and $M_t := P_{t+1}^{-1} A_t$; actually, for dimensions to match, we remove the top coordinate of \mathbf{v}_t and the first row and first column of M_t (see previous section for definition of B_t). Transforming the previous identity into a telescoping sum, it follows from $\Delta_0 = \mathbf{x}_0 - \mathbf{y}_0 = \mathbf{0}$ and the definition $B_{t:s} = B_{t-1} B_{t-2} \dots B_s$ that

$$\Delta_t = \sum_{s=0}^{t-1} B_{t:s+1} M_s \mathbf{v}_s = \sum_{s=0}^{t-1} R_{t,s} \mathbf{v}_s, \quad (22)$$

where $R_{t,s} := B_{t:s+1} M_s$. We denote by C_1, C_2, \dots suitably large constants (possibly depending on $\kappa, d, n, \tau, \tau_0$).

By (14), $\|M_s\|_\infty \leq C_1/(s+1)$ and, by (19), for sufficiently large s ,

$$\|B_{t:s+1}\|_\infty \leq C_2(s+1)^\beta(t+1)^{-\beta},$$

where $\beta = 1/2d\kappa < 1$. Combining the above inequalities, we obtain the following estimate of $R_{t,s}$ as

$$\|R_{t,s}\|_\infty \leq C_3(s+1)^{-1+\beta}(t+1)^{-\beta}. \quad (23)$$

In the remainder of the proof, the power of a vector is understood element-wise. We use the fact that \mathbf{v}_s and $\mathbf{v}_{s'}$ are independent if $s \neq s'$ and that the expectation of an odd power of an unbiased Gaussian is always zero. By Cauchy-Schwarz and Jensen's inequalities,

$$\begin{aligned} \mathbb{E} \Delta_t^4 &= \left(\sum_{s=0}^{t-1} R_{t,s} \mathbf{v}_s \right)^4 \\ &= \sum_{s=0}^{t-1} \mathbb{E}(R_{t,s} \mathbf{v}_s)^4 + \sum_{0 \leq s \neq s' < t} 3 \mathbb{E}(R_{t,s} \mathbf{v}_s)^2 \mathbb{E}(R_{t,s'} \mathbf{v}_{s'})^2 \\ &\leq \sum_{s=0}^{t-1} \mathbb{E}(R_{t,s} \mathbf{v}_s)^4 + 3 \left(\sum_{s=0}^{t-1} \mathbb{E}(R_{t,s} \mathbf{v}_s)^2 \right)^2 \\ &\leq \sum_{s=0}^{t-1} \mathbb{E}(R_{t,s} \mathbf{v}_s)^4 + 3t \sum_{s=0}^{t-1} \mathbb{E}^2(R_{t,s} \mathbf{v}_s)^2 \\ &\leq (3t+1) \sum_{s=0}^{t-1} \mathbb{E}(R_{t,s} \mathbf{v}_s)^4. \end{aligned} \quad (24)$$

Notice that since the variance of $\mathbf{v}_t = (v_{t,1}, \dots, v_{t,n})^T$ is nonincreasing, there exists a constant C_4 such that $\mathbb{E} v_{t,i}^4 \leq C_4$. By Jensen's inequality and the fact that the variables $v_{t,i}$ are independent for different values of i , we have, for any i, j, k, l ,

$$|\mathbb{E} v_{t,i} v_{t,j} v_{t,k} v_{t,l}| \leq \max_k \mathbb{E} v_{t,k}^4.$$

By direct calculation, it then follows that

$$\begin{aligned} \max_i \mathbb{E}(R_{t,s} \mathbf{v}_s)_i^4 &= \max_i \mathbb{E} \left(\sum_{j=1}^n (R_{t,s})_{ij} v_{s,j} \right)^4 \\ &\leq \max_i \left(\sum_{j=1}^n (R_{t,s})_{ij} \right)^4 \max_k \mathbb{E} v_{s,k}^4 \\ &= \|R_{t,s}\|_\infty^4 \max_k \mathbb{E} v_{s,k}^4 \\ &\leq C_5(s+1)^{-4+4\beta}(t+1)^{-4\beta}. \end{aligned} \quad (25)$$

Summing (25) over $0 \leq s \leq t-1$, we conclude from (24) that $\mathbb{E} \Delta_t^4 \leq C_6 t^{-2}$, and thus

$$\sum_{t \geq 0} \mathbb{E} \Delta_t^4 \leq C_6 \sum_{t \geq 1} t^{-2} \leq C_7. \quad (26)$$

By the Borel-Cantelli lemma, it follows that Δ_t vanishes almost surely. \square

Theorem 2.1 follows directly from Lemmas 3.1 and 4.1. \square

We now show why the truth-hearing assumption is necessary. We describe a sequence of graphs G_t that allows every agent infinite access to the truth and yet does not lead to truthful consensus. For this, it suffices to ensure that the expected mean process \mathbf{y}_t does not converge. Consider a system with two learning agents with priors $\mu_{0,1}$ and $\mu_{0,2}$ from the same distribution $\mathcal{N}(2, 1)$. We have $x_{0,1} = x_{0,2} = y_{0,1} = y_{0,2} = 2$ and, as usual, the truth is assumed to be 0; the noise variance is $\sigma^2 = 1$. The graph sequence is defined as follows: set $t_1 = 0$; for $k = 1, 2, \dots$, agent 1 links to the truth agent at time t_k and to agent 2 at times $t_k + 1, \dots, s_k - 1$; then at time s_k , agent 2 links to the truth agent, and then to agent 1 at times $s_k + 1, \dots, t_{k+1} - 1$. The time points s_k and t_k are defined recursively to ensure that

$$y_{s_k,1} \geq 1 + 2^{-2k+1} \quad \text{and} \quad y_{t_k,2} \geq 1 + 2^{-2k}. \quad (27)$$

In this way, the expected mean processes of the two agents alternate while possibly sliding down toward 1 but never lower. The existence of these time points can be proved by induction. Since $y_{0,2} = 2$, the inequality $y_{t_k,2} \geq 1 + 2^{-2k}$ holds for $k = 1$, so let's assume it holds up to $k > 0$. The key to the proof is that, by (3), as agent 1 repeatedly links to agent 2, she is pulled arbitrarily close to it. Indeed, the transition rule gives us

$$y_{t+1,1} = \frac{(P_t)_{11}}{(P_{t+1})_{11}} y_{t,1} + \frac{1}{(P_{t+1})_{11}} y_{t,2},$$

where $(P_{t+1})_{11} = (P_t)_{11} + 1$, which implies that $y_{t,1}$ can be brought arbitrarily close to $y_{t,2}$ while the latter does not move: this follows from the fact that any product of the form $\prod_{t > t_a}^t \frac{t}{t+1}$ tends to 0 as t_b grows.²

It follows that a suitably increasing sequence of s_k, t_k ensures the two conditions (27). The beliefs of the two agents do not converge to the truth even though they link to the truth agent infinitely often.

REFERENCES

- [1] Alireza Tabbaz-Salehi, Alvaro Sandroni, and Ali Jadbabaie. Learning under social influence. In *Decision and Control, 2009 held jointly with the 2009 28th Chinese Control Conference. CDC/CCC 2009. Proceedings of the 48th IEEE Conference on*, pages 1513–1519. IEEE, 2009.
- [2] Daron Acemoglu and Asuman Ozdaglar. Opinion dynamics and learning in social networks. *Dynamic Games and Applications*, 1(1):3–49, 2011.
- [3] Benjamin Golub and Matthew O. Jackson. Naive learning in social networks and the wisdom of crowds. *American Economic Journal: Microeconomics*, 2(1):112–149, 2010.
- [4] Daron Acemoglu, Munther A Dahleh, Ilan Lobel, and Asuman Ozdaglar. Bayesian learning in social networks. *The Review of Economic Studies*, 78(4):1201–1236, 2011.
- [5] Manuel Mueller-Frank. A general framework for rational learning in social networks. *Theoretical Economics*, 8(1):1–40, 2013.
- [6] Ilan Lobel and Evan Sadler. Preferences, homophily, and social learning. *Operations Research*, 64(3):564–584, 2015.

²We note that the construction shares a family resemblance with one used by Moreau [33] to show the non-consensual dynamics of certain multiagent averaging systems. The difference here is that the weights of the averaging change at each step by increasing the agent's self-confidence.

- [7] Elchanan Mossel, Allan Sly, and Omer Tamuz. From agreement to asymptotic learning. *Arxiv preprint arXiv*, 1105, 2011.
- [8] Abhijit V. Banerjee. A simple model of herd behavior. *The Quarterly Journal of Economics*, pages 797–817, 1992.
- [9] Venkatesh Bala and Sanjeev Goyal. Learning from neighbours. *The review of economic studies*, 65(3):595–621, 1998.
- [10] Pooya Molavi, Alireza Tahbaz-Salehi, and Ali Jadbabaie. Foundations of non-bayesian social learning. *Columbia Business School Research Paper*, 2015.
- [11] Mohammad Amin Rahimian and Ali Jadbabaie. Learning without recall from actions of neighbors. In *2016 American Control Conference (ACC)*, pages 1060–1065. IEEE, 2016.
- [12] Ali Jadbabaie, Pooya Molavi, Alvaro Sandroni, and Alireza Tahbaz-Salehi. Non-bayesian social learning. *Games and Economic Behavior*, 76(1):210–225, 2012.
- [13] Benjamin Golub and Matthew O. Jackson. How homophily affects the speed of learning and best response dynamics. 2012.
- [14] Ali Jadbabaie, Pooya Molavi, and Alireza Tahbaz-Salehi. Information heterogeneity and the speed of learning in social networks. *Columbia Business School Research Paper*, (13-28), 2013.
- [15] Morris H. DeGroot. Reaching a consensus. *Journal of the American Statistical Association*, 69(345):118–121, 1974.
- [16] Mohammad Amin Rahimian, Shahin Shahrampour, and Ali Jadbabaie. Learning without recall by random walks on directed graphs. In *2015 54th IEEE Conference on Decision and Control (CDC)*, pages 5538–5543. IEEE, 2015.
- [17] Mohammad Amin Rahimian et al. Learning without recall: A case for log-linear learning. *IFAC-PapersOnLine*, 48(22):46–51, 2015.
- [18] Mohammad Amin Rahimian and Ali Jadbabaie. Naive social learning in ising networks. In *2016 American Control Conference (ACC)*, pages 1088–1093. IEEE, 2016.
- [19] Rainer Hegselmann and Ulrich Krause. Opinion dynamics and bounded confidence models, analysis, and simulation. *Journal of Artificial Societies and Social Simulation*, 5(3), 2002.
- [20] Soheil Mohajer and Behrouz Touri. On convergence rate of scalar hegselmann-krause dynamics. In *Proceedings of the IEEE American Control Conference (ACC)*, 2013.
- [21] Bernard Chazelle and Chu Wang. Inertial Hegselmann-Krause systems. In *Proceedings of the IEEE American Control Conference (ACC)*, pages 1936–1941, 2016.
- [22] Bernard Chazelle. Diffusive influence systems. *SIAM Journal on Computing*, 44(5):1403–1442, 2015.
- [23] George E.P. Box and George C. Tiao. *Bayesian inference in statistical analysis*, volume 40. John Wiley & Sons, 2011.
- [24] Matan Harel, Elchanan Mossel, Philipp Strack, and Omer Tamuz. The speed of social learning. *arXiv preprint arXiv:1412.7172*, 2014.
- [25] Veronika Grimm and Friederike Mengel. An experiment on belief formation in networks. *Available at SSRN 2361007*, 2014.
- [26] Arun G. Chandrasekhar, Horacio Larreguy, and Juan Pablo Xandri. Testing models of social learning on networks: Evidence from a lab experiment in the field. Technical report, National Bureau of Economic Research, 2015.
- [27] Abhimanyu Das, Sreenivas Gollapudi, and Kamesh Munagala. Modeling opinion dynamics in social networks. In *Proceedings of the 7th ACM international conference on Web search and data mining*, pages 403–412. ACM, 2014.
- [28] Eytan Bakshy, Itamar Rosenn, Cameron Marlow, and Lada Adamic. The role of social networks in information diffusion. In *Proceedings of the 21st international conference on World Wide Web*, pages 519–528. ACM, 2012.
- [29] Elchanan Mossel and Omer Tamuz. Efficient bayesian learning in social networks with gaussian estimators. *arXiv preprint arXiv:1002.0747*, 2010.
- [30] Giuseppe Moscarini, Marco Ottaviani, and Lones Smith. Social learning in a changing world. *Economic Theory*, 11(3):657–665, 1998.
- [31] Thomas L. Griffiths and Michael L. Kalish. Language evolution by iterated learning with bayesian agents. *Cognitive Science*, 31(3):441–480, 2007.
- [32] Kenny Smith. Iterated learning in populations of bayesian agents. In *Proceedings of the 31st annual conference of the cognitive science society*, pages 697–702. Citeseer, 2009.
- [33] Luc Moreau. Stability of multiagent systems with time-dependent communication links. *IEEE Transactions on Automatic Control*, 50:169–182, 2005.