

A New Approach to Distributed Hypothesis Testing and Non-Bayesian Learning: Improved Learning Rate and Byzantine Resilience

Aritra Mitra , John A. Richards , and Shreyas Sundaram , Senior Member, IEEE

Abstract—We study a setting where a group of agents, each receiving partially informative private signals, seek to collaboratively learn the true underlying state of the world (from a finite set of hypotheses), which generates their joint observation profiles. To solve this problem, we propose a distributed learning rule that differs fundamentally from existing approaches, in that it does not employ any form of "belief-averaging." Instead, agents update their beliefs based on a min-rule. Under standard assumptions on the observation model and the network structure, we establish that each agent learns the truth asymptotically almost surely. As our main contribution, we prove that with probability 1, each false hypothesis is ruled out by every agent exponentially fast, at a network-independent rate that is strictly larger than existing rates. We then develop a computationally efficient variant of our learning rule that is provably resilient to agents who do not behave as expected (as represented by a Byzantine adversary model) and deliberately try to spread misinformation.

Index Terms—Fault-tolerant control, inference algorithms, multi-agent systems, statistical Learning.

I. INTRODUCTION

IVEN noisy data, the task of making meaningful inferences about a quantity of interest is at the heart of various complex estimation and detection problems arising in signal processing, information theory, machine learning, and control systems. When the information required to solve such problems

Manuscript received July 4, 2019; revised June 16, 2020; accepted September 25, 2020. Date of publication October 22, 2020; date of current version August 30, 2021. This work was supported in part by NSF CAREER Award 1653648 and in part by the Laboratory Directed Research and Development program at Sandia National Laboratories, a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under Contract DE-NA0003525. The views expressed in the article do not necessarily represent the views of the U.S. Department of Energy or the United States Government. Recommended by Associate Editor F. Fagnani. (Corresponding author: Aritra Mitra.)

Aritra Mitra and Śhreyas Sundaram are with the School of Electrical and Computer Engineering, Purdue University, Philadelphia, PA 19104 USA (e-mail: ariabhi1990@gmail.com; sundara2@purdue.edu).

John A. Richards is with the Pathfinder Technologies, Sandia National Laboratories, Albuquerque, NM 87185-1163 USA (e-mail: jaricha@sandia.gov).

Color versions of one or more of the figures in this article are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TAC.2020.3033126

is dispersed over a network, several interesting questions arise. How should the individual entities in the network combine their own private observations with the information received from neighbors to learn the quantity of interest? What are the minimal requirements on the information structure of the entities and the topology of the network for this to happen? How fast does information spread as a function of the diffusion rule and the structure of the network? What can be said when the underlying network changes with time and/or certain entities deviate from nominal behavior? In this article, we provide rigorous theoretical answers to such questions for the setting where a group of agents receive a stream of private signals generated by an unknown quantity known as the "true state of the world." Communication among such agents is modeled by a graph. The goal of each agent is to eventually identify the true state from a finite set of hypotheses. However, while the *collective* signals across all agents might facilitate identification of the true state, signals received by any given agent may, in general, not be rich enough for identifying the state in isolation. Thus, the problem of interest is to develop and analyze local interaction rules that facilitate inference of the true state at every agent. The setup described above serves as a common mathematical abstraction for modeling and analyzing various decision-making problems in social and economic networks (e.g., opinion formation and spreading), and classification/detection problems arising in large-scale engineered systems (e.g., object recognition by a group of aerial robots). While the former is typically studied under the moniker of non-Bayesian social learning, the latter usually goes by the name of distributed detection/hypothesis testing. In what follows, we discuss relevant literature.

A. Related Literature

Much of the earlier work on this topic of interest assumed the existence of a centralized fusion center for performing computational tasks [4]–[6]. Our work in this article, however, belongs to a more recent body of literature wherein individual agents are endowed with computational capabilities, and interactions among them are captured by a graph [1]–[3], [7]–[16]. These works are essentially inspired by the model in [1], where each agent maintains a belief vector (over the set of hypotheses) that

¹Although the model of interest to us (see Section II) has been used to study decision-making in social networks [1]–[3], we do not claim that the rules developed in this article capture human reasoning in any way.

0018-9286 © 2020 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

is sequentially updated as the convex combination of its own Bayesian posterior and the priors of its neighbors. Subsequent approaches share a common theme: They typically involve a learning rule that combines a local Bayesian update with a consensus-based opinion pooling of neighboring beliefs. The key point of distinction among such rules stems from the specific manner in which neighboring opinions are aggregated. Specifically, linear opinion pooling is studied in [1], [2], [7], and [8], whereas log-linear opinion pooling is studied in [9]–[16]. Under appropriate conditions on the observation model and the network structure, each of these approaches enables every agent to learn the true state exponentially fast, with probability 1. The rate of convergence, however, depends on the specific nature of the learning rule. Notably, finite-time concentration results are derived in [11]–[13], and a large-deviation analysis is conducted in [14] and [15] for a broad class of distributions that generate the agents' observation profiles. Extensions to different types of time-varying graphs have also been considered in [7], [10]–[13]. In a recent paper [3], Molavi et al. go beyond specific functional forms of belief-update rules and, instead, adopt an axiomatic framework that identifies the fundamental factors responsible for social learning. We point out that belief-consensus algorithms on graphs have been studied prior to Jadbabaie et al. [1] as well as in [17] and [18]. The model in [17] and [18] differs from that in [1], [2], [7]–[16] in one key aspect: While in the former each agent has access to only one observation, the latter allows for influx of new information into the network in the form of a time-series of observations at every agent.

B. Our Contributions

In light of the aforementioned developments, we now elaborate on the main contributions of this article.

1) Novel Distributed Learning Rule: In [12, Sec. III], Nedić et al. explain that the commonly studied linear and log-linear forms of belief aggregation are specific instances of a more general class of opinion pooling known as g-Quasi-Linear Opinion pools (g-QLOP), introduced in [19]. Our first contribution is the development of a novel belief update rule that deviates fundamentally from the broad family of g-QLOP learning rules. Specifically, the learning algorithm that we propose in Section III does not rely on any linear consensus-based belief aggregation protocol. Instead, each agent maintains two sets of belief vectors: a local belief vector and an actual belief vector. Each agent updates its local belief vector in a Bayesian manner based on only its private observations, i.e., without the influence of neighbors. The actual belief on each hypothesis is updated (up to normalization) as the minimum of the agent's own local belief and the actual beliefs of its neighbors on that particular hypothesis. We provide theoretical guarantees on the performance of this algorithm in Section IV. As we explain later in the article, establishing such guarantees requires proof techniques that differ substantially from those existing.

2) Strict Improvement in Rate of Learning: While data-aggregation via arithmetic or geometric averaging of neighboring beliefs allows asymptotic learning, such schemes may potentially dilute the rate at which false hypotheses are eliminated. In

particular, for the linear consensus protocol introduced in [1], the limiting rate at which a particular false hypothesis is eliminated is almost surely upper-bounded by a quantity that depends on the relative entropies and centralities of the agents [2]. The log-linear rules in [11]–[15] improve upon such a rate: With probability 1, the asymptotic rate of rejection of a false hypothesis under such rules is a convex combination of the agents' relative entropies, where the convex weights correspond to the eigenvector centralities of the agents. In contrast, based on our approach, each false hypothesis is rejected by every agent exponentially fast, at a rate that is almost surely lower-bounded by the best relative entropy (between the true state and the false hypothesis) among all agents, provided the underlying network is static and strongly connected. In Theorem 1, we show that the aforementioned result continues to hold even when the network changes with time, as long as a mild joint strong-connectivity condition is met. Thus, to the best of our knowledge, our approach leads to a strict improvement in the rate of learning over all existing approaches: This constitutes our main contribution.

3) Resilience to Adversaries: Despite the wealth of literature on distributed inference, there is limited understanding of the impact of misbehaving agents who do not follow the prescribed learning algorithm. Such agents may represent stubborn individuals or ideological extremists in the context of a social network, or model faults (either benign or malicious) in a networked control system. In the presence of such misbehaving entities, how should the remaining agents process their private observations and the beliefs of their neighbors to eventually learn the truth? To answer this question, we capture deviant behavior via the classical Byzantine adversary model [20], and develop a provably correct, resilient version of our proposed learning rule in Section V. Theorem 4 characterizes the performance of this rule and, in particular, reveals that each regular agent can infer the truth exponentially fast. Furthermore, we identify conditions on the observation model and the network structure that guarantee applicability of our Byzantine-resilient learning rule, and argue that such conditions can be checked in polynomial time. The only related work that we are aware of in this regard is [16]. As we discuss in detail in Section V, our proposed approach has various computational advantages relative to those in [16].

In addition to the main contributions discussed earlier, a minor contribution of this article is the following. For static graphs where all agents behave normally, Corollary 1 establishes consistency of our learning rule under conditions that are necessary for *any* belief update rule to work, when agents make conditionally independent observations. In particular, we show that the typical assumption of strong-connectivity on the network can be relaxed, and identify the minimal requirement for uniquely learning any state that gets realized.² Despite its various advantages, our approach cannot, in general, handle the scenario where there does not exist any single true state that generates signals consistent with those seen by every agent. The method in [12] and [13], however, is applicable to this case as

²A strongly connected graph has a path between every pair of nodes.

well, and enables each agent to identify the hypothesis that *best* explains the groups' observations.

A preliminary version of this article appeared as the work in [21]. We significantly expand upon the content in [21] by first, providing detailed convergence rate analyses of our algorithms, second, extending our results to the case of time-varying graphs, and third, elaborating on the significance of our results relative to prior work, and validating them via suitable simulation studies.

II. MODEL AND PROBLEM FORMULATION

- 1) Network Model: Let $\mathbb N$ and $\mathbb N_+$ denote the set of nonnegative integers and positive integers, respectively. We consider a group of agents $\mathcal V=\{1,2,\dots,n\}$ interacting over a timevarying, directed communication graph $\mathcal G[t]=(\mathcal V,\mathcal E[t])$, where $t\in\mathbb N$. An edge $(i,j)\in\mathcal E[t]$ indicates that agent i can directly transmit information to agent j at time-step t. If $(i,j)\in\mathcal E[t]$, then, at time t, agent i will be called a neighbor of agent j, and agent j will be called an out-neighbor of agent i. The set $\mathcal N_i[t]$ will be used to denote the neighbors of agent i (excluding itself) at time t, whereas the set $\mathcal N_i[t] \cup \{i\}$ will be referred to as the inclusive neighborhood of agent i at time t. We will use $|\mathcal C|$ to denote the cardinality of a set $\mathcal C$.
- 2) Observation Model: Let $\Theta = \{\theta_1, \theta_2, \dots, \theta_m\}$ denote m possible states of the world; each $\theta_i \in \Theta$ will be called a hypothesis. At each time-step $t \in \mathbb{N}_+$, every agent $i \in \mathcal{V}$ privately observes a signal $s_{i,t} \in \mathcal{S}_i$, where \mathcal{S}_i denotes the signal space of agent i. The joint observation profile so generated across the network is denoted $s_t = (s_{1,t}, s_{2,t}, \dots, s_{n,t})$, where $s_t \in \mathcal{S}$, and $\mathcal{S} = \mathcal{S}_1 \times \mathcal{S}_2 \times \dots \mathcal{S}_n$. The signal s_t is generated based on a conditional likelihood function $l(\cdot|\theta^*)$, governed by the true state of the world $\theta^* \in \Theta$. Let $l_i(\cdot|\theta^*)$, $i \in \mathcal{V}$ denote the ith marginal of $l(\cdot|\theta^*)$. The signal structure of each agent $i \in \mathcal{V}$ is then characterized by a family of parameterized marginals $\{l_i(w_i|\theta): \theta \in \Theta, w_i \in \mathcal{S}_i\}$.

We make the following standard assumptions [1]–[3], [7], [8], [10]–[16].

- i) The signal space of each agent i, namely S_i , is finite.
- ii) Each agent i has knowledge of its local likelihood functions $\{l_i(\cdot|\theta_p)\}_{p=1}^m$, and it holds that $l_i(w_i|\theta) > 0, \forall w_i \in \mathcal{S}_i$, and $\forall \theta \in \Theta$.
- iii) The observation sequence of each agent is described by an i.i.d. random process over time; however, at any given time-step, the observations of different agents may potentially be correlated.
- iv) There exists a fixed true state of the world $\theta^{\star} \in \Theta$ (unknown to the agents) that generates the observations of all the agents.

Finally, we define a probability triple $(\Omega, \mathcal{F}, \mathbb{P}^{\theta^*})$, where $\Omega \triangleq \{\omega : \omega = (s_1, s_2, \ldots), s_t \in \mathcal{S}, t \in \mathbb{N}_+\}$, \mathcal{F} is the σ -algebra generated by the observation profiles, and \mathbb{P}^{θ^*} is the probability measure induced by sample paths in Ω . Specifically, $\mathbb{P}^{\theta^*} = \prod_{t=1}^{\infty} l(\cdot|\theta^*)$. For the sake of brevity, we will say that an event

occurs almost surely to mean that it occurs almost surely w.r.t. the probability measure $\mathbb{P}^{\theta^{\star}}$.

Remark 1: We point out that the existence of a true state that generates the private signals of all agents is a critical assumption for our approach to work; the method in [12] does not require this assumption. Moreover, unlike the work in [9] and [22], our rules do not apply to continuous parameter spaces.

Note that assumptions (i) and (ii) on the observation model imply the existence of a constant $L \in (0, \infty)$ such that

$$\max_{i \in \mathcal{V}} \max_{w_i \in \mathcal{S}_i} \max_{\theta_p, \theta_q \in \Theta} \left| \log \frac{l_i(w_i | \theta_p)}{l_i(w_i | \theta_q)} \right| \le L. \tag{1}$$

We will make use of the above fact later in our analysis.

Given the previous setup, the goal of each agent in the network is to discern the true state of the world θ^{\star} . The challenge associated with such a task stems from the fact that the private signal structure of any given agent is in general only partially informative. To make this notion precise, define $\Theta_i^{\theta^{\star}} \triangleq \{\theta \in \Theta: l_i(w_i|\theta) = l_i(w_i|\theta^{\star}), \forall w_i \in \mathcal{S}_i\}$. In words, $\Theta_i^{\theta^{\star}}$ represents the set of hypotheses that are observationally equivalent to the true state θ^{\star} from the perspective of agent i. In general, for any agent $i \in \mathcal{V}$, we may have $|\Theta_i^{\theta^{\star}}| > 1$, necessitating collaboration among agents subject to the restrictions imposed by the time-varying communication topology.

Our *objective* in this article will be to design a distributed learning rule that allows each agent $i \in \mathcal{V}$ to identify the true state of the world asymptotically almost surely. To this end, we now introduce the following notion of source agents that will be useful in our subsequent developments.

Definition 1 (Source Agents): An agent i is said to be a source agent for a pair of distinct hypotheses $\theta_p, \theta_q \in \Theta$, if $K_i(\theta_p, \theta_q) > 0$, where $K_i(\theta_p, \theta_q)$ represents the Kullback-Leibler divergence between the distributions $l_i(\cdot|\theta_p)$ and $l_i(\cdot|\theta_q)$, and is given by⁴

$$K_i(\theta_p, \theta_q) = \sum_{w_i \in S_i} l_i(w_i | \theta_p) \log \frac{l_i(w_i | \theta_p)}{l_i(w_i | \theta_q)}.$$
 (2)

The set of all source agents for the pair θ_p, θ_q is denoted by $\mathcal{S}(\theta_p, \theta_q)$.

In words, a source agent for a pair $\theta_p, \theta_q \in \Theta$ is an agent that can distinguish between the pair of hypotheses θ_p, θ_q based on its private signal structure. It should be noted that $\mathcal{S}(\theta_p, \theta_q) = \mathcal{S}(\theta_q, \theta_p)$, since $K_i(\theta_p, \theta_q) > 0 \iff K_i(\theta_q, \theta_p) > 0$ [23]. In this article, we will assume that each state $\theta \in \Theta$ is *globally identifiable* w.r.t. the joint observation model of the entire network. Based on our terminology of source agents, this translates to the following.

Assumption 1 (Global Identifiability): For each pair $\theta_p, \theta_q \in \Theta$ such that $\theta_p \neq \theta_q$, the set $\mathcal{S}(\theta_p, \theta_q)$ of agents that can distinguish between the pair θ_p, θ_q is nonempty.

The previous assumption is standard in the related literature. To illustrate the concepts described earlier, let us consider the following simple example.

³While $w_i \in \mathcal{S}_i$ will be used to refer to a generic element of the signal space of agent $i, s_{i,t}$ will denote the random variable (with distribution $l_i(\cdot|\theta^\star)$) that corresponds to the observation of agent i at time-step t.

⁴Although the standard notation for the Kullback-Leibler divergence between $l_i(\cdot|\theta_p)$ and $l_i(\cdot|\theta_q)$ is $D(l_i(\cdot|\theta_p)||l_i(\cdot|\theta_q))$, we use $K_i(\theta_p,\theta_q)$ as a shorthand for the same to avoid cluttering the exposition.

	$s_1 = H$	$s_2 = T$
$l_1(\cdot heta_1)$	1/2	1/2
$l_1(\cdot \theta_2)$	1/4	3/4
$l_1(\cdot \theta_3)$	1/2	1/2

	$s_1 = H$	$s_2 = T$
$l_2(\cdot \theta_1)$	1/3	2/3
$l_2(\cdot \theta_2)$	1/3	2/3
$l_2(\cdot \theta_3)$	1/6	5/6

Fig. 1. Likelihood models for the two agents in Example 1. The model on the left is that of agent 1, whereas that on the right corresponds to agent 2.

Example 1: Consider a network of two agents with likelihood models as described in Fig. 1 . At every time-step, each agent either observes heads H, or tails T. Thus, the common signal space for both agents is $\mathcal{S}_1 = \mathcal{S}_2 = \{H, T\}$. From Fig. 1, note that at each time-step, the probability of agent 1 observing H is 0.5 if either θ_1 or θ_3 gets realized, and 0.25 if θ_2 gets realized. Observe immediately that $\Theta_1^{\theta_1} = \{\theta_1, \theta_3\}$ and $\Theta_1^{\theta_2} = \{\theta_2\}$, i.e., agent 1 cannot distinguish between the states θ_1 and θ_3 ; however, it can tell θ_2 apart from either of the other two states. Agent 2's likelihood model can be interpreted similarly. Based on our terminology, we then have $\mathcal{S}(\theta_1,\theta_2) = 1$, $\mathcal{S}(\theta_2,\theta_3) = \{1,2\}$, and $\mathcal{S}(\theta_3,\theta_1) = 2$, implying global identifiability as per Assumption 1.

In addition to Assumption 1, we will make a mild assumption on the time-varying communication topology. To this end, let the union graph over an interval $[t_1,t_2], 0 \leq t_1 < t_2$, indicate a graph with vertex set \mathcal{V} , and edge set $\bigcup_{\tau=t_1}^{t_2} \mathcal{E}[\tau]$. Based on this convention, we will assume (unless stated otherwise) that the sequence of communication graphs $\{\mathcal{G}[t]\}_{t=0}^{\infty}$ is jointly strongly connected, in the following sense.

Assumption 2 (*Joint Strong-Connectivity*): There exists $T \in \mathbb{N}_+$ such that the union graph over every interval of the form [rT, (r+1)T) is strongly connected, where $r \in \mathbb{N}$.

While the previous assumption on the network connectivity pattern is not necessary for solving the problem at hand, it is fairly standard in the analysis of distributed algorithms over time-varying networks [12], [24], [25]. Having introduced the model and the problem formulation, we now proceed to a formal description of our learning algorithm.

III. PROPOSED LEARNING RULE

In this section, we propose a novel belief update rule (see Algorithm 1) and discuss the intuition behind it. Every agent i maintains and updates (at every time-step t) two separate sets of belief vectors, namely, $\pi_{i,t}$ and $\mu_{i,t}$. Each of these vectors is a probability distribution over the hypothesis set Θ . We will refer to $\pi_{i,t}$ and $\mu_{i,t}$ as the "local" belief vector (for reasons that will soon become obvious), and the "actual" belief vector, respectively, maintained by agent i. The goal of each agent $i \in \mathcal{V}$ in the network will be to use its own private signals and the information available from its neighbors to update $\mu_{i,t}$ sequentially, so that $\lim_{t\to\infty}\mu_{i,t}(\theta^*)=1$ almost surely. To do so, at each time-step t+1 (where $t\in\mathbb{N}$), agent i does the following for each $\theta\in\Theta$. It first generates $\pi_{i,t+1}(\theta)$ via a local Bayesian update rule that incorporates the private observation $s_{i,t+1}$ using $\pi_{i,t}(\theta)$ as a prior (line 5 in Algorithm 1). Having generated $\pi_{i,t+1}(\theta)$, agent

Algorithm 1: Belief update rule for each $i \in \mathcal{V}$.

- 1: **Initialization:** $\mu_{i,0}(\theta) > 0$, $\pi_{i,0}(\theta) > 0$, $\forall \theta \in \Theta$, and $\sum_{\theta \in \Theta} \mu_{i,0}(\theta) = 1$, $\sum_{\theta \in \Theta} \pi_{i,0}(\theta) = 1$
- 2: Transmit $\mu_{i,0}$ to out-neighbors at time 0
- 3: for $t+1 \in \mathbb{N}_+$ do
- 4: for $\theta \in \Theta$ do
- 5: Update local belief on θ as

$$\pi_{i,t+1}(\theta) = \frac{l_i(s_{i,t+1}|\theta)\pi_{i,t}(\theta)}{\sum_{p=1}^{m} l_i(s_{i,t+1}|\theta_p)\pi_{i,t}(\theta_p)}$$
(3)

6: Update actual belief on θ as

$$\mu_{i,t+1}(\theta) = \frac{\min\{\{\mu_{j,t}(\theta)\}_{j \in \mathcal{N}_i[t] \cup \{i\}}, \pi_{i,t+1}(\theta)\}}{\sum_{p=1}^{m} \min\{\{\mu_{j,t}(\theta_p)\}_{j \in \mathcal{N}_i[t] \cup \{i\}}, \pi_{i,t+1}(\theta_p)\}}$$
(4)

- 7: end for
- 8: Transmit $\mu_{i,t+1}$ to out-neighbors at time t+1
- 9: end for

i updates $\mu_{i,t+1}(\theta)$ (up to normalization) by setting it to be the *minimum* of its locally generated belief $\pi_{i,t+1}(\theta)$, and the actual beliefs $\mu_{j,t}(\theta), j \in \mathcal{N}_i[t] \cup \{i\}$ of its inclusive neighborhood at the previous time-step (line 6 in Algorithm 1). It then reports $\mu_{i,t+1}$ to each of its out-neighbors at time t+1.

A. Intuition Behind the Learning Rule

At the core of our learning algorithm are two key principles: 1) preservation of the intrinsic discriminatory capabilities of the agents, and 2) propagation of low beliefs on each false hypothesis. We now elaborate on these features.

Consider the set of source agents $S(\theta^*, \theta)$ that can differentiate between a certain false hypothesis θ and the true state θ^{\star} . By definition, the signal structures of such agents are rich enough for them to be able to eliminate θ on their own, i.e., without the support of their neighbors. To achieve this, we require each agent to maintain a local belief vector that is updated (via (3)) without any network influence using only the agent's own private signals. Doing so ensures that $\pi_{i,t}(\theta) \to 0$ a.s. for each $i \in \mathcal{S}(\theta^*, \theta)$. Next, leveraging this property, we want to be able to propagate low beliefs on θ from $\mathcal{S}(\theta^*, \theta)$ to $\mathcal{V} \setminus \mathcal{S}(\theta^*, \theta)$, i.e., the agents in $\mathcal{S}(\theta^*, \theta)$ should contribute toward driving the actual beliefs of their out-neighbors (and eventually, of all the agents in the set $V \setminus S(\theta^*, \theta)$) on the hypothesis θ to zero. Using a min-rule of the form (4), with $\pi_{i,t+1}(\theta)$ featuring as an external network-independent input, facilitates such propagation without compromising the abilities of agents in $S(\theta^*, \theta)$ to eliminate θ . When set in motion, our learning rule triggers a process of belief reduction on θ originating at $\mathcal{S}(\theta^*, \theta)$ that eventually propagates to each agent in the network reachable from $S(\theta^*, \theta)$.

⁵Note that based on our algorithm, agents only exchange their actual beliefs, and not their local beliefs.

Remark 2: We emphasize that the proposed learning rule given by Algorithm 1 does not employ any form of "belief-averaging." This feature is in stark contrast with existing approaches to distributed hypothesis testing that rely either on linear opinion pooling [1], [2], [7], [8], or log-linear opinion pooling [9]–[16]. As such, the lack of linearity in our belief update rule precludes (direct or indirect) adaptation of existing analysis techniques to suit our needs.

IV. ANALYSIS OF ALGORITHM 1

A. Statement of the Results

In this section, we characterize the performance of Algorithm 1. We start with one of the main results of the article, proven in Appendix A. Before stating the result, we remind the reader that for an agent i, $K_i(\theta_p,\theta_q)$ represents the KL-divergence between the distributions $l_i(\cdot|\theta_p)$ and $l_i(\cdot|\theta_q)$, and captures agent i's ability to distinguish between the states θ_p and θ_q .

Theorem 1: Suppose the observation model satisfies the global identifiability condition (see Assumption 1), and the sequence of communication graphs $\{\mathcal{G}[t]\}_{t=0}^{\infty}$ is jointly strongly connected (see Assumption 2). Then, Algorithm 1 provides the following guarantees.

- 1) (Consistency): For each agent $i \in \mathcal{V}$, $\mu_{i,t}(\theta^*) \to 1$ a.s.
- 2) (Asymptotic Rate of Rejection of False Hypotheses): Consider any false hypothesis $\theta \in \Theta \setminus \{\theta^*\}$. Then, the following holds for each agent $i \in \mathcal{V}$:

$$\liminf_{t \to \infty} -\frac{\log \mu_{i,t}(\theta)}{t} \ge \max_{v \in \mathcal{S}(\theta^*, \theta)} K_v(\theta^*, \theta) \ a.s. \quad (5)$$

The above result tells us that with probability 1, every agent i will be able to rule out each false hypothesis θ exponentially fast, at a rate that is eventually lower-bounded by the best KL-divergence across the network between the pair of hypotheses θ^* and θ . In particular, this implies that given any $\epsilon>0$, the probability that agent i's instantaneous rate of rejection of θ , namely $-\log \mu_{i,t}(\theta)/t$, is lower than the quantity $\max_{v\in\mathcal{S}(\theta^*,\theta)} K_v(\theta^*,\theta)$ by an additive factor of ϵ , decays to zero. The next result, proven in Appendix B, sheds some light on the rate of decay of this probability.

Theorem 2: Suppose the conditions in Theorem 1 hold. Fix $\theta \in \Theta \setminus \{\theta^*\}$, and let $\bar{K}(\theta^*,\theta) = \max_{v \in \mathcal{S}(\theta^*,\theta)} K_v(\theta^*,\theta)$. Then, for every $\epsilon > 0$ and $\delta \in (0,1)$, there exists a set $\Omega'(\delta) \subseteq \Omega$ with $\mathbb{P}^{\theta^*}(\Omega'(\delta)) \geq 1 - \delta$, such that the following holds for each agent $i \in \mathcal{V}$:

$$\liminf_{t \to \infty} -\frac{1}{t} \log \mathbb{P}^{\theta^*} \left(\left\{ -\frac{\log \mu_{i,t}(\theta)}{t} \right\} \right) \leq \bar{K}(\theta^*, \theta) - \epsilon \left\{ \cap \Omega'(\delta) \right\} \geq \frac{\epsilon^2}{8L^2}.$$
(6)

Our next result pertains to the special case when the communication graph does not change over time, i.e., when $\mathcal{G}[t] = \mathcal{G}, \forall t \in \mathbb{N}$. To state the result, we will employ the following terminology. Given two disjoint sets $\mathcal{C}_1, \mathcal{C}_2 \subseteq \mathcal{V}$, we say \mathcal{C}_2 is reachable from \mathcal{C}_1 if for every $i \in \mathcal{C}_2$, there exists a directed path

in \mathcal{G} from some $j \in \mathcal{C}_1$ to agent i (note that j will in general be a function of i).

Corollary 1: Let the communication graph be time-invariant and be denoted by \mathcal{G} . Suppose the following conditions hold. (i) The observation model satisfies the global identifiability condition (see Assumption 1). (ii) For every pair of hypotheses $\theta_p \neq \theta_q \in \Theta$, the set $\mathcal{V} \setminus \mathcal{S}(\theta_p, \theta_q)$ is reachable from the set $\mathcal{S}(\theta_p, \theta_q)$ in \mathcal{G} . Then, Algorithm 1 guarantees consistency as in Theorem 1. Furthermore, for every $\theta \in \Theta \setminus \{\theta^\star\}$, the following holds for each agent $i \in \mathcal{V}$:

$$\liminf_{t \to \infty} -\frac{\log \mu_{i,t}(\theta)}{t} \ge \max_{v \in \mathcal{S}_i(\theta^*, \theta)} K_v(\theta^*, \theta) \ a.s. \tag{7}$$

where $S_i(\theta^*, \theta) \subseteq S(\theta^*, \theta)$ are those source agents from which there exists a directed path to i in G.

Proof: Fix $\theta \in \Theta \setminus \{\theta^*\}$, and consider an agent $i \in \mathcal{V} \setminus \mathcal{S}(\theta^*, \theta)$. The sets $\mathcal{S}(\theta^*, \theta)$ and $\mathcal{S}_i(\theta^*, \theta)$ are nonempty based on conditions (i) and (ii) of the theorem, respectively. Following a similar line of argument as in the proof of Theorem 1, one can establish the following for each $v \in \mathcal{S}_i(\theta^*, \theta)$:

$$\liminf_{t \to \infty} -\frac{\log \mu_{i,t}(\theta)}{t} \ge K_v(\theta^*, \theta) \ a.s.$$
(8)

The assertion regarding equation (7) then follows readily. Consistency follows by noting that since $S_i(\theta^*, \theta) \subseteq S(\theta^*, \theta)$, $K_v(\theta^*, \theta) > 0, \forall v \in S_i(\theta^*, \theta)$.

Our next result reveals that the combination of conditions (i) and (ii) in Corollary 1 constitutes *minimal* requirements on the observation model and the network structure for *any* learning algorithm to guarantee consistency, when the observations of the agents are conditionally independent.

Theorem 3: Let the communication graph be time-invariant and be denoted by \mathcal{G} . Then, the following assertions hold.

- i) Conditions (i) and (ii) in Corollary 1, taken together, are equivalent to global identifiability of each source component of \mathcal{G} .
- ii) Suppose the observations of the agents are independent conditional on the realization of any state, i.e., $l(\cdot|\theta) = \prod_{i=1}^n l_i(\cdot|\theta), \forall \theta \in \Theta$. Then, global identifiability of each source component of \mathcal{G} is necessary and sufficient for unique identification of any true state that gets realized, at every agent, with probability 1.

The proof of the aforementioned result is fairly straightforward and hence omitted here. We now leverage the previous results to quantify the rate at which the overall network uncertainty about the true state decays to zero. To measure such uncertainty, we employ the following metric from [2], which captures the total variation distance between the agents' beliefs at time-step t, and the probability distribution that is concentrated entirely on the true state of the world, namely $\mathbf{1}_{\theta^*}(\cdot)$:

$$e_t(\theta^*) \triangleq \frac{1}{2} \sum_{i=1}^n \|\boldsymbol{\mu}_{i,t}(\cdot) - \mathbf{1}_{\theta^*}(\cdot)\|_1 = \sum_{i=1}^n \sum_{\theta \neq \theta^*} \mu_{i,t}(\theta). \quad (9)$$

 $^{^6}$ A source component of a time-invariant graph $\mathcal G$ is a strongly connected component with no incoming edges.

Given that θ^* gets realized, the *rate of social learning* is then defined as [2], [14]:

$$\rho_L(\theta^*) \triangleq \liminf_{t \to \infty} -\frac{1}{t} \log e_t(\theta^*). \tag{10}$$

Notice that the previous expression depends on the state being realized; to account for the realization of any state, one can simply look at the quantity $\min_{\theta^* \in \Theta} \rho_L(\theta^*)$ that provides a sense for the least rate of learning one can expect given a certain observation model, a network, and a consistent learning algorithm. We have the following simple results; their proofs are trivial and hence omitted.

Corollary 2: Suppose the conditions stated in Theorem 1 are met. Then, Algorithm 1 guarantees:

$$\rho_L(\theta^*) \ge \min_{\theta \ne \theta^*} \max_{v \in \mathcal{S}(\theta^*, \theta)} K_v(\theta^*, \theta) \ a.s. \tag{11}$$

Corollary 3: Suppose the conditions stated in Corollary 1 are met. Then, Algorithm 1 guarantees:

$$\rho_L(\theta^*) \ge \min_{\theta \ne \theta^*} \min_{i \in \mathcal{V}} \max_{v \in \mathcal{S}_i(\theta^*, \theta)} K_v(\theta^*, \theta) \ a.s.$$
 (12)

B. Discussion of the Results

1) Comments on Theorem 1: Let us compare the rate of learning based on our method to those existing in the literature. Under identical assumptions of global identifiability of the observation model, and strong-connectivity (or joint strong-connectivity as in [12]) of the underlying communication graph, both linear [1], [2] and log-linear [11], [12], [14] opinion pooling lead to an asymptotic rate of rejection of the form $\sum_{i\in\mathcal{V}} \nu_i K_i(\theta^*,\theta)$ for each false hypothesis $\theta\in\Theta\setminus\{\theta^*\}$, for each agent $i \in \mathcal{V}$. Here, ν_i represents the eigenvector centrality of agent $i \in \mathcal{V}$, which is strictly positive for a strongly connected graph. Thus, referring to (5) reveals that the asymptotic rate of rejection of each false hypothesis (and hence, the rate of social learning) resulting from our algorithm (see (11)), is a strict improvement over all existing rates—This constitutes a significant contribution of our paper. Furthermore, observe from Corollary 2 that the lower bound on the rate of social learning is independent of both the size and structure of the network. A key implication of this result is the fact that as long as the total information content of the network remains the same, the specific manner in which signals are allocated to agents does not impact the long-run learning rate of our approach. In sharp contrast, existing learning rates that depend on the agents' eigenvector centralities may suffer under poor signal allocations (see [2] for a discussion on this topic).

It should, however, be noted that the network independence aspect of our approach concerns *asymptotic* learning rates. The

 $^7 \mathrm{In}$ [12], the consensus weights are chosen to obtain a network-structure independent (albeit network-size dependent) rate of rejection of θ of the form $1/n\sum_{i\in\mathcal{V}}K_i(\theta^\star,\theta).$ The same rate is obtained with static, undirected networks when the consensus weight matrix is symmetric, since the eigenvector centralities are simply 1/n in such a case.

dependence on the network structure (presumably, on the diameter) is bound to manifest itself in the transients generated by our rule. Given the nonlinear structure of our update rule (4), characterizing such a dependence is quite nontrivial.

2) Comments on Theorem 2: At any given time t, for some $i \in \mathcal{V}$ and $\theta \neq \theta^*$, let us consider the set of all sample paths where agent i's instantaneous rate of rejection of θ is lower than its asymptotic lower bound by a constant additive factor of ϵ . Theorem 2 complements Theorem 1 by telling us that an arbitrarily accurate approximation of the measure of such "bad" sample paths eventually decays to zero at an exponential rate no smaller than $\epsilon^2/8L^2$ (the approximation is arbitrarily accurate since the set $\Omega'(\delta)$ can be chosen to have measure arbitrarily close to 1). It is instructive to compare the concentration result of Theorem 2 with [12, Th. 2], [14, Th. 2], and [11, Lemma 3]. The analogous results in these papers are more elegant relative to ours, since they do not involve a set of the form $\Omega'(\delta)$ that shows up in our analysis. A refinement of Theorem 2 to obtain a cleaner nonasymptotic result would require a precise characterization of the transient dynamics generated by our learning rule: We reserve investigations along this line as future work.

3) Comments on Corollary 1: While Theorem 3 identifies an algorithm-independent necessary condition for ensuring unique identifiability of any realized state at every agent (when the communication graph is time-invariant and agents receive conditionally independent signals), Corollary 1 reveals that such a condition is also sufficient for our proposed learning algorithm to work. We believe that a result of this flavor is missing in the existing literature on distributed hypothesis testing, where strong-connectivity is a standard assumption. Molavi et al. [26] do relax the strong-connectivity assumption, but require every strongly connected component of \mathcal{G} to be globally identifiable for learning to take place [26, Prop. 4]. In contrast, Corollary 1 requires only the source components of \mathcal{G} to satisfy the global identifiability requirement. Interestingly, our conclusions in this context align with an analogous result that identifies joint detectability of each source component as the minimal requirement for solving the related problem of distributed state estimation [27], [28].

The more general network condition in Corollary 1 (as opposed to strong-connectivity) comes at the cost of a potential reduction in the rate of social learning, as reflected in Corollary 3. When the underlying graph is strongly connected, $S_i(\theta^\star,\theta) = S(\theta^\star,\theta)$. Consequently, the min w.r.t. the agent set \mathcal{V} in (12) goes away, and we recover Corollary 2.

V. LEARNING DESPITE MISINFORMATION

In this section, we will address the problem of learning the true state of the world despite the presence of certain agents who do not behave as expected and deliberately try to spread misinformation. In order to isolate the challenges introduced by such malicious entities, we will consider a time-invariant communication graph $\mathcal G$ for our subsequent discussion; we anticipate that our proposed approach will extend to the time-varying

case with suitable modifications. We now describe the model of agent-misbehavior that we consider.⁸

Adversary Model: We assume that a certain subset of the agents are adversarial, and model their behavior based on the Byzantine fault model [30]. Specifically, Byzantine agents possess complete knowledge of the observation model, the network model, the algorithms being used, the information being exchanged, and the true state of the world. Leveraging such information, adversarial agents can behave arbitrarily and in a coordinated manner, and can in particular, send incorrect, potentially inconsistent information to their out-neighbors. In return for allowing such worst case adversarial behavior, we will restrict the number of adversaries; in particular, we will consider an f-local adversarial model, i.e., we assume that there are at most f adversaries in the neighborhood of any nonadversarial agent, where $f \in \mathbb{N}$. Finally, we emphasize that the nonadversarial agents are unaware of the identities of the adversaries in their neighborhood. As is fairly standard in the distributed fault-tolerant literature [31]-[38], we only assume that nonadversarial agents know the upper bound f on the number of adversaries in their neighborhood. The adversarial set will be denoted by $A \subset V$, and the remaining agents $R = V \setminus A$ will be called the regular agents.

Our immediate goals are as follows. First, devise an algorithm that enables each regular agent to asymptotically identify the true state with probability 1, despite the presence of an f-local Byzantine adversarial set. Second, identify conditions on the observation model and the network structure that guarantee correctness of such an algorithm. Prior to addressing these goals, we briefly motivate the need for a novel Byzantine-resilient learning algorithm.

Motivation: A standard way to analyze the impact of adversarial agents while designing resilient distributed consensusbased protocols (for applications in consensus [31], [32], optimization [34], [35], hypothesis testing [16], and multiagent rendezvous [39]) is to construct an equivalent matrix representation of the linear update rule that involves only the regular agents [40]. In particular, this requires expressing the iterates of a regular agent as a convex combination of the iterates of its regular neighbors, based on appropriate filtering techniques, and under certain assumptions on the network structure. While this can indeed be achieved efficiently for scalar consensus problems, for problems requiring consensus on vectors (like the belief vectors in our setting), such an approach typically requires the computation of sets known as *Tverberg partitions*. However, there is no known algorithm that can compute an exact Tverberg partition in polynomial time for a general d-dimensional finite point set [41]. Consequently, since the filtering approach developed in [16] requires each regular agent to compute a Tverberg partition at every iteration, the resulting computations are forbiddingly high. Su and Vaidya [16] do briefly discuss an alternate pairwise learning rule that requires agents to perform scalar consensus on relative confidence levels (instead of beliefs)

of one hypothesis over another. Under such a rule, for each regular agent, its relative confidence on the true state over every false hypothesis approaches infinity—a condition that is difficult to verify in practice. Moreover, the pairwise learning rule in [16] requires each agent to maintain and update at each time-step a vector of dimension $O(m^2)$. In contrast, we propose a simple, lightweight Byzantine-resilient learning rule that avoids the computation of Tverberg partitions, and requires agents to update two m-dimensional belief vectors.

Algorithm 2: Belief update rule for each $i \in \mathcal{R}$.

- 1: **Initialization:** $\mu_{i,0}(\theta) > 0$, $\pi_{i,0}(\theta) > 0$, $\forall \theta \in \Theta$, and $\sum_{\theta \in \Theta} \mu_{i,0}(\theta) = 1$, $\sum_{\theta \in \Theta} \pi_{i,0}(\theta) = 1$
- 2: Transmit $\mu_{i,0}$ to out-neighbors
- 3: for $t+1 \in \mathbb{N}_+$ do
- 4: for $\theta \in \Theta$ do
- 5: Update local belief on θ as per (3)
- 6: **if** $|\mathcal{N}_i| \ge (2f+1)$ **then**
- 7: Sort $\mu_{j,t}(\theta)$, $j \in \mathcal{N}_i$ from highest to lowest, and reject the highest f and the lowest f of such beliefs.
- 8: Let $\mathcal{M}_{i,t}^{\theta}$ be the set of agents whose beliefs are not rejected in the previous step. Update $\mu_{i,t+1}(\theta)$ as

$$\mu_{i,t+1}(\theta) = \frac{\min\{\{\mu_{j,t}(\theta)\}_{j \in \mathcal{M}_{i,t}^{\theta}}, \pi_{i,t+1}(\theta)\}}{\sum_{p=1}^{m} \min\{\{\mu_{j,t}(\theta_{p})\}_{j \in \mathcal{M}_{i,t}^{\theta_{p}}}, \pi_{i,t+1}(\theta_{p})\}}$$
(13)

- 9: **els**
- 10: Update $\mu_{i,t+1}(\theta)$ as

$$\mu_{i,t+1}(\theta) = \pi_{i,t+1}(\theta) \tag{14}$$

- 11: **end if**
- 12: end for
- 13: Transmit $\mu_{i,t+1}$ to out-neighbors
- 14: **end for**

A. Byzantine-Resilient-Distributed Learning Rule

In this section, we develop an easy to implement and computationally efficient extension of Algorithm 1 that guarantees learning despite the presence of Byzantine adversaries. We call it the local-filtering-based resilient hypothesis elimination (LFRHE) algorithm (see Algorithm 2). Like Algorithm 1, the LFRHE algorithm requires every regular agent i to maintain and update (at every time-step t) a local belief vector $\boldsymbol{\pi}_{i,t}$, and an actual belief vector $\mu_{i,t}$. While $\pi_{i,t}$ is updated as before via (3), the update of $\mu_{i,t}$ is the key feature of Algorithm 2. To update $\mu_{i,t+1}(\theta)$, agent $i \in \mathcal{R}$ first checks whether it has at least 2f+1neighbors. If it does, then it rejects the highest f and the lowest f neighboring beliefs $\mu_{i,t}(\theta), j \in \mathcal{N}_i$ (line 7 in Algorithm 2), and employs a min-rule as before, but using only the remaining beliefs (line 8 in Algorithm 2). Thus, agent i filters out the most extreme neighboring beliefs on each hypothesis, and retains only the moderate ones to update its own actual belief. If agent ihas strictly fewer than 2f + 1 neighbors, then it decides against using neighboring information and, instead, updates its actual

⁸Different from our setting, the *forceful* agents in [29] do not behave arbitrarily and, in fact, update their beliefs (even if infrequently) by interacting with their neighbors; our adversary model makes no such assumptions.

belief vector to be equal to its local belief vector (line 10 in Algorithm 2).

To state our main result concerning the correctness of Algorithm 2, we require the following definitions.

Definition 2 (*r-Reachable Set* [32]): For a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, a set $\mathcal{C} \subseteq \mathcal{V}$, and an integer $r \in \mathbb{N}_+$, \mathcal{C} is an *r-reachable set* if there exists an $i \in \mathcal{C}$ such that $|\mathcal{N}_i \setminus \mathcal{C}| \geq r$.

Definition 3 (Strongly r-robust graph: w.r.t. $\mathcal{S}(\theta_p, \theta_q)$) For $r \in \mathbb{N}_+$ and $\theta_p, \theta_q \in \Theta$, a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is strongly r-robust w.r.t. the set of source agents $\mathcal{S}(\theta_p, \theta_q)$, if for every nonempty subset $\mathcal{C} \subseteq \mathcal{V} \setminus \mathcal{S}(\theta_p, \theta_q)$, \mathcal{C} is r-reachable.

Theorem 4: Suppose that for every pair of hypotheses $\theta_p, \theta_q \in \Theta$, the graph \mathcal{G} is strongly (2f+1)-robust w.r.t. the source set $\mathcal{S}(\theta_p, \theta_q)$. Then, Algorithm 2 guarantees the following despite the actions of any f-local set of Byzantine adversaries.

- 1) (Consistency): For each agent $i \in \mathcal{R}$, $\mu_{i,t}(\theta^*) \to 1$ a.s.
- 2) (Asymptotic Rate of Rejection of False Hypotheses): Consider any false hypothesis $\theta \in \Theta \setminus \{\theta^*\}$. Then, the following holds for each agent $i \in \mathcal{R}$.

$$\liminf_{t \to \infty} -\frac{\log \mu_{i,t}(\theta)}{t} \ge \min_{v \in \mathcal{S}(\theta^*,\theta) \cap \mathcal{R}} K_v(\theta^*,\theta) \ a.s.$$
(15)

Proof: See Appendix C.

Remark 3: For any pair $\theta_p, \theta_q \in \Theta$, notice that the strongrobustness condition in Theorem 4 (together with Definition 3) requires $|\mathcal{S}(\theta_p, \theta_q)| \geq (2f + 1)$, if $\mathcal{V} \setminus \mathcal{S}(\theta_p, \theta_q)$ is nonempty. In particular, it blends requirements on the signal structures of the agents with those on the communication graph. To gain intuition about this condition, suppose $\Theta = \{\theta_1, \theta_2\}$, and consider an agent $i \in \mathcal{V} \setminus \mathcal{S}(\theta_1, \theta_2)$. To enable i to learn the truth despite potential adversaries in its neighborhood, one requires (i) redundancy in the signal structures of the agents, and (ii) redundancy in the network structure to ensure reliable information flow from $S(\theta_1, \theta_2)$ to agent i. These requirements are encapsulated by Theorem 4. For a fixed source set $S(\theta_p, \theta_q)$, checking whether \mathcal{G} is strongly (2f+1)-robust w.r.t. $\mathcal{S}(\theta_p,\theta_q)$ can be done in polynomial time by drawing connections to the process of bootstrap percolation on networks [36, Prop. 5]. Since the source sets for each pair $\theta_p, \theta_q \in \Theta$ can also be computed in polynomial time via a simple inspection of the agents' signal structures, it follows that the strong-robustness condition in Theorem 4 can be checked in polynomial time.

Leveraging Theorem 4, we can characterize the rate of decay of the collective uncertainty of the regular agents regarding the true state. To do so, we employ the following modification of the metric (9):

$$e_t^{\mathcal{R}}(\theta^*) \triangleq \frac{1}{2} \sum_{i \in \mathcal{R}} \|\boldsymbol{\mu}_{i,t}(\cdot) - \mathbf{1}_{\theta^*}(\cdot)\|_1 = \sum_{i \in \mathcal{R}} \sum_{\theta \neq \theta^*} \mu_{i,t}(\theta).$$
 (16)

Note that this metric only considers the beliefs of the regular agents as the Byzantine agents can update their beliefs however they wish. With θ^* as the true state, we define the rate of social learning in the presence of Byzantine adversaries as:

$$\rho_L^{\mathcal{R}}(\theta^*) \triangleq \liminf_{t \to \infty} -\frac{1}{t} \log e_t^{\mathcal{R}}(\theta^*). \tag{17}$$

We have the following immediate corollary of Theorem 4.

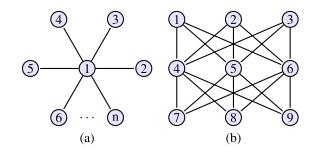


Fig. 2. (a) and (b) Network models for simulation examples 1 and 2, respectively.

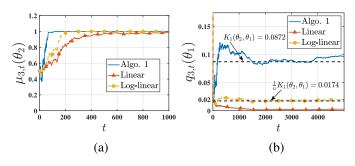


Fig. 3. Consider the setup of simulation example 1 with n=5 agents. Image (a) depicts the evolution of agent 3's belief on the true state θ_2 . Image (b) depicts the evolution of the instantaneous rate of rejection of θ_1 for agent 3, namely $q_{3,t}(\theta_1) = -\log \mu_{3,t}(\theta_1)/t$.

Corollary 4: Suppose the conditions stated in Theorem 4 are met. Then, Algorithm 2 guarantees:

$$\rho_L^{\mathcal{R}}(\theta^*) \ge \min_{\theta \ne \theta^*} \min_{v \in \mathcal{S}(\theta^*, \theta) \cap \mathcal{R}} K_v(\theta^*, \theta) \ a.s.$$
 (18)

VI. SIMULATIONS

Example 1 (Impact of Network Size on Rate of Convergence): For our first simulation study, we consider a binary hypothesis testing problem, i.e., $\Theta = \{\theta_1, \theta_2\}$, where the signal space for each agent is identical and comprises signals w_1 and w_2 . The (time-invariant) undirected network for this example is depicted in Fig. 2(a). The likelihood models of the agents are as follows: $l_1(w_1|\theta_1)=0.7, l_1(w_1|\theta_2)=0.5$, and $l_i(w_1|\theta_1) = l_i(w_1|\theta_2) = 0.5, \forall i \in V \setminus \{1\}, \text{ i.e., agent 1 is the}$ only informative agent. In order to compare the performance of Algorithm 1 to the linear and log-linear belief update rules in [1] and [12], we implement the latter assuming consensus weights are assigned based on the lazy Metropolis scheme (see [12] for details). Based on this weight assignment, it is easy to verify that the eigenvector centrality of each agent is 1/n. All agents start out with uniform priors. With $\theta^* = \theta_2$, and n = 5, Fig. 3 illustrates the performance of the three algorithms w.r.t. agent 3. In particular, Fig. 3(a) reveals that based on our approach, agent 3's belief on the true state θ_2 converges to 1 faster than the other algorithms. Fig. 3(b) makes this observation precise by plotting the instantaneous rate of rejection of θ_1 for agent 3, namely $q_{3,t}(\theta_1) = -\log \mu_{3,t}(\theta_1)/t$. Consistent with the respective theoretical findings, $q_{3,t}(\theta_1)$ is eventually lower-bounded

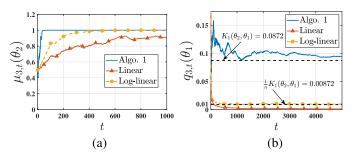


Fig. 4. Consider the setup of simulation example 1 with n=10 agents. It illustrates the dilution in the rates of social learning for the linear and log-linear rules with an increase in the number of uninformative agents. Images (a) and (b) are analogous to those in Fig. 3.

by $K_1(\theta_2,\theta_1)$ for our algorithm (see Theorem 1), approaches $K_1(\theta_2,\theta_1)/n$ for the log-linear rule in [12], and is eventually upper-bounded by $K_1(\theta_2,\theta_1)/n$ for the linear rule in [1]. Similar conclusions hold for the other agents.

Suppose we now double the number of agents in the network. Agent 1 continues to remain the only informative agent. Fig. 4 compares the performances of the three algorithms for this case. Notably, the convergence rate for our approach remains unaffected, whereas that for the linear and log-linear rules gets diluted. This observation can be attributed to the fact that while the rate provided by our algorithm is both network-structure and network-size independent for strongly connected networks (see Section IV-B), the rates of the linear and log-linear rules depend crucially on the eigenvector centralities of the agents, which, in this case, correspond to 1/n. Thus, the gap between the performance of our algorithm, and that of the linear and log-linear update rules (as measured by convergence rates), becomes more pronounced as the number of uninformative agents increase (i.e., as n increases, but the total information content of the network remains the same).

Example 2 (Impact of Adversaries): While the previous example highlighted the benefits of Algorithm 1, we now focus on an example that demonstrates the resilience of its variant, namely the LFRHE algorithm (see Algorithm 2), to the presence of Byzantine adversaries. To this end, consider the undirected network in Fig. 2(b). For this example, $\Theta = \{\theta_1, \theta_2, \theta_3\}$, and $S_i = \{w_1, w_2\}, \forall i \in \mathcal{V}$. Suppose the agent likelihood models are given by $l_i(w_1|\theta_1) = 3/4, l_i(w_1|\theta_2) = l_i(w_1|\theta_3) = 1/3, \forall i \in \{1, 2, 3\},\$ $l_i(w_1|\theta_1) = l_i(w_1|\theta_2) = 2/5, l_i(w_1|\theta_3) = 1/7, \forall i \in \{4, 5, 6\},\$ $l_i(w_1|\theta_1) = l_i(w_1|\theta_2) = 1/2, l_i(w_1|\theta_3) = 5/6, \forall i \in$ $\{7, 8, 9\}$. Suppose f = 1 and agent 5 is the only adversarial agent. It is easy to see that condition (i) in Theorem 4 is met. We will compare the performance of Algorithm 2 with the linear rule in [1], and the log-linear rule in [12]. For implementing the latter, we again assign consensus weights based on the lazy Metropolis scheme. All agents start out with uniform priors. The adversary, agent 5, maintains a belief of 0.1 on the true state, and 0.45 on each of the false hypotheses, for all $t \ge 20$. Figures 5(a) and (b) illustrates the repercussions of this action on agent 7, when $\theta^* = \theta_1$ and $\theta^* = \theta_2$, respectively: while the linear and log-linear rules fail to recover from the

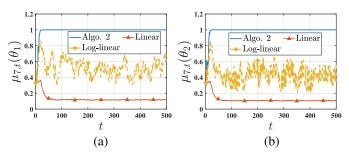


Fig. 5. Consider the setup of simulation example 2, where agent 5 acts as an adversary. Images (a) and (b) depict the evolution of agent 7's belief on the true state, when $\theta^* = \theta_1$, and $\theta^* = \theta_2$, respectively.

attack, Algorithm 2 enables agent 7 to infer the truth. Similar conclusions hold for the other regular agents.

VII. CONCLUSION

We proposed and analyzed a novel algorithm for addressing the problem of distributed hypothesis testing. The key distinguishing feature of our learning algorithm is that it does not employ any linear consensus-based data aggregation protocol. Instead, it relies on a "min-rule" to spread beliefs through the network. Under mild assumptions of global identifiability and joint strong-connectivity, we established consistency of our learning rule. In particular, we showed that the rate of learning resulting from our approach strictly improves upon all existing rates. For static networks, we established consistency of our algorithm under minimal requirements on the observation model and the network structure. Finally, we proposed a simple and computationally efficient version of our learning rule that accounts for worst case adversarial behavior on the part of certain agents in the network.

As future work, we plan to investigate the transient dynamics generated by our rule with the goal of deriving finite-time concentration results. Moreover, as we stated in Remark 1, our approach cannot handle scenarios where there does not exist a true state; it would be interesting to see if there are variations of our rule, which are applicable to these cases as well. Finally, as part of our ongoing research, we are interested in analyzing the impacts of communication constraints on the performance of distributed learning algorithms; preliminary results along this line are reported in [42].

APPENDIX A PROOF OF THEOREM 1

The proof of Theorem 1 is based on several intermediate results. We start with the following simple lemma that characterizes the asymptotic behavior of the local belief sequences generated based on (3); we provide a proof (adapted to our notation) to keep the article self-contained, and to introduce certain quantities that will be referenced later in our analysis.

Lemma 1: Consider a false hypothesis $\theta \in \Theta \setminus \{\theta^*\}$, and an agent $i \in \mathcal{S}(\theta^*, \theta)$. Suppose $\pi_{i,0}(\theta_p) > 0, \forall \theta_p \in \Theta$. Then, the update rule (3) ensures that (i) $\pi_{i,t}(\theta) \to 0$ a.s., (ii) $\pi_{i,\infty}(\theta^*) \triangleq \lim_{t \to \infty} \pi_{i,t}(\theta^*)$ exists a.s. and satisfies $\pi_{i,\infty}(\theta^*) \geq \pi_{i,0}(\theta^*)$,

and (iii) the following holds:

$$\lim_{t \to \infty} \frac{1}{t} \log \frac{\pi_{i,t}(\theta)}{\pi_{i,t}(\theta^*)} = -K_i(\theta^*, \theta) \ a.s.$$
 (19)

Proof: Consider any agent $i \in \mathcal{S}(\theta^*, \theta)$, and define

$$\rho_{i,t}(\theta) \triangleq \log \frac{\pi_{i,t}(\theta)}{\pi_{i,t}(\theta^{\star})}, \quad \lambda_{i,t}(\theta) \triangleq \log \frac{l_i(s_{i,t}|\theta)}{l_i(s_{i,t}|\theta^{\star})}. \tag{20}$$

Then, based on (3), we obtain the following recursion:

$$\rho_{i,t+1}(\theta) = \rho_{i,t}(\theta) + \lambda_{i,t+1}(\theta), \quad \forall t \in \mathbb{N}.$$
 (21)

Rolling out the aforementioned equation over time yields

$$\rho_{i,t}(\theta) = \rho_{i,0}(\theta) + \sum_{k=1}^{t} \lambda_{i,k}(\theta), \quad \forall t \in \mathbb{N}_{+}.$$
 (22)

Notice that $\{\lambda_{i,t}(\theta)\}$ is a sequence of i.i.d. random variables with finite means (see equation (1)). In particular, it is easy to verify that each random variable $\lambda_{i,t}(\theta)$ has mean given by $-K_i(\theta^\star,\theta)$. Thus, based on the strong law of large numbers, we have $\frac{1}{t}\sum_{k=1}^t \lambda_{i,k}(\theta) \to -K_i(\theta^\star,\theta)$ almost surely. Dividing both sides of (22) by t, and taking the limit as t goes to infinity, we then obtain:

$$\lim_{t \to \infty} \frac{1}{t} \rho_{i,t}(\theta) = -K_i(\theta^*, \theta) \ a.s., \tag{23}$$

establishing part (iii) of the lemma. Now note that based on the definition of the set $\mathcal{S}(\theta^\star,\theta), K_i(\theta^\star,\theta)>0$. It then follows from (23) that $\rho_{i,t}(\theta)\to -\infty$ almost surely, and hence $\pi_{i,t}(\theta)\to 0$ almost surely. This establishes part (i) of the lemma. For any $\theta\in\Theta_i^{\theta^\star}$, observe that $\lambda_{i,t}(\theta)=0, \forall t\in\mathbb{N}_+$. It then follows from (21) that for each $\theta\in\Theta_i^{\theta^\star}, \ \rho_{i,t}(\theta)=\rho_{i,0}(\theta), \forall t\in\mathbb{N}_+$. From the earlier discussion, we conclude that a limiting belief vector $\pi_{i,\infty}$ exists almost surely, with non-zero entries corresponding to each $\theta\in\Theta_i^{\theta^\star}$. Part (ii) of the lemma then follows readily.

While our proposed learning rule is tailored to facilitate propagation of low beliefs on false hypotheses, it is crucial to also ensure that the beliefs of all agents on the true state remain bounded away from zero. In particular, consider the following scenario. During a transient phase, certain agents see private signals that cause them to temporarily lower their local beliefs on the true state. This effect manifests itself in the actual beliefs of the agents via the min-rule (4). We ask: Can such a transient phenomenon trigger a cascade of progressively lower beliefs on the true state? The next important result asserts that this will almost surely never be the case.

Lemma 2: Suppose the conditions stated in Theorem 1 hold, and Algorithm 1 is employed by each agent. Then, there exists a set $\bar{\Omega} \subseteq \Omega$ with the following properties: (i) $\mathbb{P}^{\theta^*}(\bar{\Omega}) = 1$, and (ii) for each $\omega \in \bar{\Omega}$, there exist constants $\eta(\omega) \in (0,1)$ and $t'(\omega) \in (0,\infty)$ such that on the sample path ω ,

$$\pi_{i,t}(\theta^*) \ge \eta(\omega), \mu_{i,t}(\theta^*) \ge \eta(\omega), \quad \forall t \ge t'(\omega), \quad \forall i \in \mathcal{V}.$$
(24)

 9 More precisely, the mean here is obtained by using the expectation operator $\mathbb{E}^{\theta^*}[\cdot]$ associated with the measure \mathbb{P}^{θ^*} .

Proof: Let $\bar{\Omega} \subseteq \Omega$ denote the set of sample paths for which assertions (i)-(iii) in Lemma 1 hold for each false hypothesis $\theta \in \Theta \setminus \{\theta^{\star}\}$. Based on Lemma 1, we note that $\mathbb{P}^{\theta^{\star}}(\bar{\Omega}) = 1$. Consequently, to prove the result, it suffices to establish the existence of $\eta(\omega) \in (0,1)$, and $t'(\omega) \in (0,\infty)$ for each sample path $\omega \in \Omega$, such that (24) holds. To this end, fix an arbitrary sample path $\omega \in \bar{\Omega}$. We first argue that the local beliefs of every agent on the true state θ^* are bounded away from 0 on ω . To see this, pick any agent $i \in \mathcal{V}$. Suppose there exists some $\theta \in \Theta \setminus \{\theta^*\}$ for which $i \in \mathcal{S}(\theta^*, \theta)$. Then, based on our choice of ω , Lemma 1 implies that $\pi_{i,\infty}(\theta^*) \geq \pi_{i,0}(\theta^*) > 0$, where the last inequality follows from the requirement of non-zero priors in line 1 of Algorithm 1. In particular, given the structure of the update rule (3), it follows that $\pi_{i,t}(\theta^*) > 0$ for all time. This is true since if $\pi_{i,t}(\theta^*) = 0$ at any instant, then the corresponding belief would remain at 0 for all subsequent time-steps, thereby violating the fact that $\pi_{i,\infty}(\theta^*) \geq \pi_{i,0}(\theta^*) > 0$. Now consider the scenario where there exists no $\theta \in \Theta \setminus \{\theta^*\}$ for which $i \in \mathcal{S}(\theta^*, \theta)$, i.e., every hypothesis in Θ is observationally equivalent to θ^* from the point of view of agent i. In this case, it is easy to see that based on (3), $\pi_{i,t} = \pi_{i,0}, \forall t \in \mathbb{N}_+$. In particular, this implies $\pi_{i,t}(\theta^*) = \pi_{i,0}(\theta^*) > 0, \forall t \in \mathbb{N}_+$. This establishes our claim that on ω , $\pi_{i,t}(\theta^*)$ remains bounded away from zero $\forall i \in \mathcal{V}$.

To proceed, define $\gamma_1 \triangleq \min_{i \in \mathcal{V}} \pi_{i,0}(\theta^*) > 0$, where the inequality follows from line 1 in Algorithm 1. Pick a small number $\delta > 0$ such that $\delta < \gamma_1$, and notice that our discussion concerning the evolution of the local beliefs readily implies the existence of a time-step $t'(\omega)$, such that for all $t \geq t'(\omega)$, $\pi_{i,t}(\theta^*) \geq \gamma_1 - 1$ $\delta > 0, \forall i \in \mathcal{V}.$ With $\gamma_2(\omega) \triangleq \min_{i \in \mathcal{V}} \{\mu_{i,t'(\omega)}(\theta^*)\}$, we claim that $\gamma_2(\omega) > 0$. The claim follows by noting that given the structure of the update rule (4), and the requirement of non-zero priors in Algorithm 1, $\gamma_2(\omega)$ can equal 0 if and only if some agent in the network sets its local belief on θ^* to 0 at some time-step prior to $t'(\omega)$. However, this possibility is ruled out in view of the previously established fact that on ω , $\pi_{i,t}(\theta^*) > 0$, $\forall t \in \mathbb{N}$, $\forall i \in$ \mathcal{V} . Let $\eta(\omega) = \min\{\gamma_1 - \delta, \gamma_2(\omega)\} > 0$. In words, $\eta(\omega)$ lowerbounds the lowest belief (considering both local and actual beliefs) on the true state θ^* held by an agent at time-step $t'(\omega)$. It is apparent from the preceding discussion that $\pi_{i,t}(\theta^{\star}) \geq$ $\eta(\omega), \forall t \geq t'(\omega), \forall i \in \mathcal{V}$. Thus, to complete the proof, it remains to establish that $\mu_{i,t}(\theta^*) \geq \eta(\omega), \forall t \geq t'(\omega), \forall i \in \mathcal{V}$. To this end, let us fix an agent i and observe the following:

$$\mu_{i,t'(\omega)+1}(\theta^{\star}) \stackrel{(a)}{=} \frac{\min\{\{\mu_{j,t'(\omega)}(\theta^{\star})\}_{j\in\mathcal{N}_{i}[t'(\omega)]\cup\{i\}}, \pi_{i,t'(\omega)+1}(\theta^{\star})\}\}}{\sum_{p=1}^{m} \min\{\{\mu_{j,t'(\omega)}(\theta_{p})\}_{j\in\mathcal{N}_{i}[t'(\omega)]\cup\{i\}}, \pi_{i,t'(\omega)+1}(\theta_{p})\}}$$

$$\stackrel{(b)}{\geq} \frac{\eta(\omega)}{\sum_{p=1}^{m} \min\{\{\mu_{j,t'(\omega)}(\theta_{p})\}_{j\in\mathcal{N}_{i}[t'(\omega)]\cup\{i\}}, \pi_{i,t'(\omega)+1}(\theta_{p})\}}}$$

$$\geq \frac{\eta(\omega)}{\sum_{p=1}^{m} \pi_{i,t'(\omega)+1}(\theta_{p})} \stackrel{(c)}{=} \eta(\omega), \tag{25}$$

where (a) is given by (4), (b) follows from the way $\eta(\omega)$ is defined and by noting that $\pi_{i,t}(\theta^*) \geq \eta(\omega), \forall t \geq t'(\omega), \forall i \in \mathcal{V}$, and (c) follows by noting that the local belief vectors generated via (3) are valid probability distributions over the hypothesis set Θ at each time-step, and hence, $\sum_{p=1}^m \pi_{i,t'(\omega)+1}(\theta_p) = 1$. The previous reasoning applies to every agent in the network, and can be repeated to establish (24) via induction.

The next result establishes that the intrinsic discriminatory capabilities of an agent are preserved under our learning rule.

Lemma 3: Suppose the conditions stated in Theorem 1 hold, and Algorithm 1 is employed by each agent. Consider any false hypothesis $\theta \in \Theta \setminus \{\theta^{\star}\}$, and an agent $i \in \mathcal{S}(\theta^{\star}, \theta)$. Then,

$$\liminf_{t \to \infty} -\frac{\log \mu_{i,t}(\theta)}{t} \ge K_i(\theta^*, \theta) \ a.s. \tag{26}$$

Proof: With $\bar{\Omega}$ defined as in Lemma 2, recall that $\mathbb{P}^{\theta^*}(\bar{\Omega}) = 1$, and pick any $\omega \in \bar{\Omega}$. Now consider any false hypothesis $\theta \in \Theta \setminus \{\theta^*\}$, and an agent $i \in \mathcal{S}(\theta^*, \theta)$. Fix any $\epsilon > 0$, and notice that since $i \in \mathcal{S}(\theta^*, \theta)$, Eq. (19) in Lemma 1 implies that there exists $t_i(\omega, \theta, \epsilon)$, such that

$$\pi_{i,t}(\theta) < e^{-(K_i(\theta^*,\theta)-\epsilon)t} \quad \forall t \ge t_i(\omega,\theta,\epsilon).$$
 (27)

Furthermore, since $\omega \in \bar{\Omega}$, Lemma 2 guarantees the existence of a time-step $t'(\omega) \in (0,\infty)$, and a constant $\eta(\omega) \in (0,1)$, such that on ω , $\pi_{i,t}(\theta^*) \geq \eta(\omega)$, $\mu_{i,t}(\theta^*) \geq \eta(\omega)$, $\forall t \geq t'(\omega)$, $\forall i \in \mathcal{V}$. Let $\bar{t}_i(\omega,\theta,\epsilon) = \max\{t'(\omega),t_i(\omega,\theta,\epsilon)\}$. Let us suppress the dependence of $\bar{t}_i(\omega,\theta,\epsilon)$ on i,ω,θ and ϵ for simplicity of notation, and observe the following inequalities:

$$\mu_{i,\bar{t}+1}(\theta) \stackrel{(a)}{\leq} \frac{\pi_{i,\bar{t}+1}(\theta)}{\sum_{p=1}^{m} \min\{\{\mu_{j,\bar{t}}(\theta_p)\}_{j \in \mathcal{N}_i[\bar{t}] \cup \{i\}}, \pi_{i,\bar{t}+1}(\theta_p)\}}$$

$$\leq \frac{\pi_{i,\bar{t}+1}(\theta)}{\min\{\{\mu_{j,\bar{t}}(\theta^*)\}_{j \in \mathcal{N}_i[\bar{t}] \cup \{i\}}, \pi_{i,\bar{t}+1}(\theta^*)\}}$$

$$\stackrel{(b)}{\leq} \frac{e^{-(K_i(\theta^*,\theta) - \epsilon)(\bar{t}+1)}}{\eta(\omega)}.$$

$$(28)$$

In the above inequalities, (a) follows from (4), whereas (b) follows from (27) and by noting that all agents have both their local and actual beliefs lower bounded by $\eta(\omega)$ beyond time-step \bar{t} . In particular, it is easy to see that the arguments used to arrive at (28) apply to each time-step $t \geq \bar{t} + 1$. Based on (28), we then obtain that $\forall t \geq \bar{t} + 1$,

$$-\frac{\log \mu_{i,t}(\theta)}{t} > (K_i(\theta^*, \theta) - \epsilon) + \frac{\log \eta(\omega)}{t}.$$
 (29)

Taking the limit inferior on both sides of (29), and noting that ϵ can be made arbitrarily small, readily leads to (26).

For the subsequent discussion, let us fix a particular false hypothesis $\theta \in \Theta \setminus \{\theta^*\}$, and assume that global identifiability holds. Let $v_{\theta} \in \operatorname{argmax}_{l \in \mathcal{S}(\theta^*, \theta)} K_l(\theta^*, \theta)$ represent any agent with the best discriminatory power w.r.t. the false hypothesis θ ,

given that θ^* gets realized. Based on Lemma 3, we have

$$\liminf_{t \to \infty} -\frac{\log \mu_{v_{\theta}, t}(\theta)}{t} \ge K_{v_{\theta}}(\theta^*, \theta) \text{ a.s.}$$
(30)

Our goal is to now establish that each agent $i \in \mathcal{V} \setminus \{v_{\theta}\}$ inherits the same asymptotic rate of rejection of θ as that of agent v_{θ} in (30). Roughly speaking, we will achieve this by showing that under the assumption of joint strong-connectivity, the belief of any agent $i \in \mathcal{V} \setminus \{v_{\theta}\}$ on θ is "not too far off" from the belief of agent v_{θ} on θ . In what follows, we make this idea precise. First, we require some additional notation: With each agent $i \in \mathcal{V}$, we associate a non-negative scalar $c_{i,t}(\theta) \in [0,\infty]$. These parameters evolve based on the following rules. ¹⁰

- i) $c_{v_{\theta},t}(\theta) = 0, \forall t \in \mathbb{N}.$
- ii) $c_{i,0}(\theta) = \infty, \forall i \in \mathcal{V} \setminus \{v_{\theta}\}.$
- iii) For each $i \in \mathcal{V} \setminus \{v_{\theta}\}$ and $t \in \mathbb{N}$, define $\tau_{i,t}(\theta) \triangleq \min_{j \in \mathcal{N}_i[t] \cup \{i\}} c_{j,t}(\theta)$, and

$$c_{i,t+1}(\theta) \triangleq \tau_{i,t}(\theta) + 1.$$
 (31)

To explain the purpose of the above rules, we will adhere to the following terminology. We say that there exists a path of length $m \in \mathbb{N}_+$ from v_{θ} to $i \in \mathcal{V} \setminus \{v_{\theta}\}$ over [t-m, t-m]1], if there exist agents $x(t-m+1), \ldots, x(t) \in \mathcal{V} \setminus \{v_{\theta}\},$ such that $(x(\tau-1), x(\tau)) \in \mathcal{E}[\tau-1]$, where $\tau \in \{t-m+1\}$ $1, \ldots, t\}, x(t-m) = v_{\theta}$, and x(t) = i. Note that the agents appearing in the path need not be distinct, and that we have assumed the presence of self-loops in each graph $\mathcal{G}[t], t \in \mathbb{N}$. Rules (i)–(iii) have been designed in a manner such that if $c_{i,t}(\theta)$ is finite at any time-step $t \in \mathbb{N}$ for any agent $i \in \mathcal{V} \setminus \{v_{\theta}\}$, then there exists a path of length $c_{i,t}(\theta)$ from v_{θ} to i over $[t-c_{i,t}(\theta),t-1]$, in the sense described above. Analyzing the time-evolution of $c_{i,t}(\theta)$ enables us to then relate the belief $\mu_{i,t}(\theta)$ of agent i to a delayed version of the belief $\mu_{v_{\theta},t}(\theta)$ of agent v_{θ} , where the delay is precisely $c_{i,t}(\theta)$ (the above statements are formalized and proven in Lemma 5). Since agent v_{θ} is the reference agent here, its delay w.r.t. its own belief on θ is set to 0 for all time, thus explaining rule (i). Initially, all agents in $V \setminus \{v_{\theta}\}$ start out with an "infinite-delay" w.r.t. the belief of agent v_{θ} ; this is captured by rule (ii). Finally, the rationale behind updating $c_{i,t}(\theta)$ via rule (iii) is to formalize the intuition that under the assumption of joint strong-connectivity, the lengths of paths linking v_{θ} to agents in $\mathcal{V} \setminus \{v_{\theta}\}$ (and hence, the corresponding delays) should eventually remain uniformly bounded; we begin by establishing this fact in the following

Lemma 4: Consider any $\theta \in \Theta \setminus \{\theta^*\}$ and suppose the joint strong-connectivity assumption (Assumption 2) holds. Then, the following is true:

$$c_{i,t}(\theta) < 2(n-1)T, \quad \forall i \in \mathcal{V}, \quad \forall t > (n-1)T, \quad (32)$$

where T is the constant appearing in Assumption 2.

Proof: Observe that the conclusion in (32) is trivially true for agent v_{θ} since $c_{v_{\theta},t}(\theta) = 0, \forall t \in \mathbb{N}$. To prove the result for

¹⁰Note that the agents do not actually maintain or update the parameters $c_{i,t}(\theta)$. Instead, they have been introduced solely for the purpose of analysis.

agents in the set $V \setminus \{v_{\theta}\}$, we begin by claiming that

$$c_{i,(n-1)T}(\theta) \le (n-1)T, \quad \forall i \in \mathcal{V}.$$
 (33)

To prove this claim, let $\mathcal{L}_0(\theta^\star,\theta)=\{v_\theta\}$, and define

$$\mathcal{L}_{1}(\theta^{\star},\theta) \triangleq \left\{ i \in \mathcal{V} \setminus \mathcal{L}_{0}(\theta^{\star},\theta) : \left\{ \bigcup_{\tau=0}^{T-1} \mathcal{N}_{i}[\tau] \right\} \right.$$
$$\left. \cap \mathcal{L}_{0}(\theta^{\star},\theta) \neq \emptyset \right\}$$
(34)

as the set of agents in $\mathcal{V}\setminus\{v_{\theta}\}$ that have a direct edge from agent v_{θ} at least once over the interval [0,T). Assumption 2 implies that $\mathcal{L}_1(\theta^{\star},\theta)$ is non-empty (barring the trivial case when $\mathcal{V}=\{v_{\theta}\}$). Now pick any agent $i\in\mathcal{L}_1(\theta^{\star},\theta)$, and notice that since $v_{\theta}\in\mathcal{N}_i[\tau]$ for some $\tau\in[0,T)$, update rule (31) implies $c_{i,\tau+1}(\theta)=1$. In particular, based on (31),

$$c_{i,t+1}(\theta) \le c_{i,t}(\theta) + 1. \tag{35}$$

Based on the previous discussion, it follows that for each agent $i \in \mathcal{L}_1(\theta^\star,\theta), c_{i,T}(\theta) \leq T$. The claim in (33) follows readily for each agent $i \in \mathcal{L}_1(\theta^\star,\theta)$ by appealing to (35). Let us now recursively define the sets $\mathcal{L}_r(\theta^\star,\theta), 1 \leq r \leq (n-1)$, as

$$\mathcal{L}_{r}(\theta^{\star},\theta) \triangleq \left\{ i \in \mathcal{V} \setminus \bigcup_{q=0}^{(r-1)} \mathcal{L}_{q}(\theta^{\star},\theta) : \left\{ \bigcup_{q=(r-1)T}^{rT-1} \mathcal{N}_{i}[\tau] \right\} \right.$$

$$\left. \cap \left\{ \bigcup_{q=0}^{(r-1)} \mathcal{L}_{q}(\theta^{\star},\theta) \right\} \neq \emptyset \right\}. \tag{36}$$

In words, $\mathcal{L}_r(\theta^*, \theta)$ are those agents belonging to $\mathcal{V} \setminus$ $\bigcup_{q=0}^{(r-1)} \mathcal{L}_q(\theta^\star,\theta)$ that each have at least one neighbor from the set $\bigcup_{q=0}^{(r-1)} \mathcal{L}_q(\theta^*, \theta)$ over the interval [(r-1)T, rT-1]. We complete the proof of the claim by inducting on r. The base case with r=1 has already been proven above. Now suppose the following is true: $c_{i,rT}(\theta) \leq rT, \forall i \in \mathcal{L}_r(\theta^*, \theta)$, where $r \in \{1, ..., m-1\}$, and $m \in \{2, ..., n-1\}$. Let r = m. If $\mathcal{V}\setminus\bigcup_{q=0}^{(m-1)}\mathcal{L}_q(\theta^\star,\theta)$ is empty, then we are done. Else, based on Assumption 2, it must be that $\mathcal{L}_m(\theta^*, \theta)$ is non-empty. Pick any agent $i \in \mathcal{L}_m(\theta^*, \theta)$, and notice that it has a neighbor j (say) from the set $\bigcup_{q=0}^{(m-1)} \mathcal{L}_q(\theta^*, \theta)$ at some time-step $\tau \in [(m-1)]$ 1)T, mT). The induction hypothesis coupled with (35) implies that $c_{j,\tau}(\theta) \leq \tau$, and hence, $c_{i,\tau+1}(\theta) \leq c_{j,\tau}(\theta) + 1 \leq \tau + 1$ based on (31). Appealing to (35) then reveals that $c_{i,mT}(\theta) \leq$ mT, thus completing the induction step. Finally, noting that $\bigcup_{q=0}^{(n-1)} \mathcal{L}_q(\theta^{\star}, \theta) = \mathcal{V}$ completes our proof of the claim (33). An identical line of argument as above can be employed to show that $c_{i,2(n-1)T} \leq (n-1)T, \forall i \in \mathcal{V}$. In particular, this can be done by first taking $C_0(\theta^*, \theta) = \{v_\theta\}$, and recursively defining the sets $C_r(\theta^*, \theta), 1 \le r \le (n-1)$ as

$$C_{r}(\theta^{\star},\theta) \triangleq \left\{ i \in \mathcal{V} \setminus \bigcup_{q=0}^{(r-1)} C_{q}(\theta^{\star},\theta) : \left\{ \bigcup_{\tau=(n+r-2)T}^{(n+r-1)T-1} \mathcal{N}_{i}[\tau] \right\} \right.$$

$$\left. \cap \left\{ \bigcup_{q=0}^{(r-1)} C_{q}(\theta^{\star},\theta) \right\} \neq \emptyset \right\}. \tag{37}$$

One can then easily prove via induction that $c_{i,(n-1+r)T}(\theta) \le rT, \forall i \in \mathcal{C}_r(\theta^*, \theta)$, where $1 \le r \le (n-1)$. The rest then follows from (35).

We can keep repeating the previous argument to establish that $c_{i,m(n-1)T}(\theta) \leq (n-1)T, \forall i \in \mathcal{V}, \forall m \in \mathbb{N}_+$. Finally, based on the previous bound and (35), it follows that for each agent $i \in \mathcal{V}$, $c_{i,t}(\theta)$ is upper bounded by 2(n-1)T at any time-step $t \in (m(n-1)T, (m+1)(n-1)T)$, where $m \in \mathbb{N}_+$. This establishes (32), and completes the proof.

The next lemma relates $\mu_{i,t}(\theta), i \in \mathcal{V} \setminus \{v_{\theta}\}$ to $\mu_{v_{\theta},t}(\theta)$ in terms of the parameter $c_{i,t}(\theta)$ and, in turn, provides the final ingredient required to prove Theorem 1.

Lemma 5: Consider any $\theta \in \Theta \setminus \{\theta^*\}$. Suppose the joint strong-connectivity assumption holds (Assumption 2), and each agent applies Algorithm 1. Suppose $c_{i,t}(\theta)$ is finite, where $i \in \mathcal{V} \setminus \{v_{\theta}\}$, and $t \in \mathbb{N}$. Then, the following are true.

- i) There exists a path of length $c_{i,t}(\theta)$ from v_{θ} to i over $[t c_{i,t}(\theta), t 1]$.
- ii) Let the path linking v_{θ} to i over $[t-c_{i,t}(\theta),t-1]$ in part (i) be denoted $x(t-c_{i,t}(\theta)),x(t-c_{i,t}(\theta)+1),\ldots,x(t)$, where $x(t-c_{i,t}(\theta))=v_{\theta}$ and x(t)=i. Then

$$\mu_{i,t}(\theta) \le \frac{\mu_{v_{\theta}, a_{i,t}(\theta)}(\theta)}{\prod_{\tau=a_{i,t}(\theta)+1}^{t} \eta_{x(\tau),\tau}(\theta^{\star})}, \tag{38}$$

where $a_{i,t}(\theta) = t - c_{i,t}(\theta)$, and

$$\eta_{i,t}(\theta^{\star}) \triangleq \min\{\{\mu_{j,t-1}(\theta^{\star})\}_{j \in \mathcal{N}_i[t-1] \cup \{i\}}, \pi_{i,t}(\theta^{\star})\},$$
$$\forall i \in \mathcal{V}. \tag{39}$$

Proof: We prove part (i) by inducting on the value of $c_{i,t}(\theta)$. For the base case, suppose $c_{i,t}(\theta)=1$ for some agent $i\in\mathcal{V}\setminus\{v_{\theta}\}$ at some time-step t. Based on (31), notice that this can happen if and only if $v_{\theta}\in\mathcal{N}_i[t-1]$; the claim in part (i) then follows readily for the base case. Fix an integer $m\geq 2$, and suppose that the assertion of part (i) holds for any agent $i\in\mathcal{V}\setminus\{v_{\theta}\}$ and at any time-step t, whenever $c_{i,t}(\theta)\in\{1,\ldots,m-1\}$. Now suppose that at some time-step t, $c_{i,t}(\theta)=m$ for some agent $i\in\mathcal{V}\setminus\{v_{\theta}\}$. Referring to (31), this is true only if $c_{l,t-1}(\theta)=m-1$ for some $l\in\mathcal{N}_i[t-1]\cup\{i\}$. Since $m\geq 2$, we have $c_{l,t-1}(\theta)\geq 1$, and hence $l\in\mathcal{V}\setminus\{v_{\theta}\}$. The induction hypothesis thus applies to agent l, implying the existence of a path of length m-1 from v_{θ} to l over $[(t-1)-c_{l,t-1}(\theta),t-2]$, i.e., over [t-m,t-2]. Appending this path with the edge $(l,i)\in\mathcal{E}[t-1]$ immediately leads to the desired conclusion.

For part (ii), consider the path $x(t-c_{i,t}(\theta)), x(t-c_{i,t}(\theta)+1), \ldots, x(t)$ from v_{θ} to i over $[t-c_{i,t}(\theta), t-1]$, where $x(t-c_{i,t}(\theta)) = v_{\theta}$ and x(t) = i. By definition of this path, $x(\tau - c_{i,t}(\theta)) = v_{\theta}$

 $^{^{11}}$ Notice that based on the update rule (31), $c_{i,t}(\theta) \geq 1, \forall i \in \mathcal{V} \setminus \{v_{\theta}\}.$ Thus, $\operatorname{argmin}_{j \in \mathcal{N}_i[t] \cup \{i\}} c_{j,t}(\theta) = v_{\theta},$ whenever $v_{\theta} \in \mathcal{N}_i[t],$ since $c_{v_{\theta},t}(\theta) = 0, \forall t \in \mathbb{N}$.

1) $\in \mathcal{N}_{x(\tau)}[\tau - 1] \cup \{x(\tau)\}$, for all $\tau \in \{a_{i,t}(\theta) + 1, \dots, t\}$. Thus, referring to (4), we obtain

$$\mu_{x(\tau),\tau}(\theta) \\ \leq \frac{\mu_{x(\tau-1),\tau-1}(\theta)}{\sum_{\tau=1}^{m} \min\{\{\mu_{j,\tau-1}(\theta_p)\}_{j \in \mathcal{N}_{x(\tau)}[\tau-1] \cup \{x(\tau)\}}, \pi_{x(\tau),\tau}(\theta_p)\}}$$

$$\leq \frac{\mu_{x(\tau-1),\tau-1}(\theta)}{\eta_{x(\tau),\tau}(\theta^*)}.\tag{40}$$

Using the above inequality recursively with $\tau \in \{a_{i,t}(\theta) + 1, \dots, t\}$ immediately leads to (38).

Proof: (**Theorem 1**): Fix a false hypothesis $\theta \in \Theta \setminus \{\theta^*\}$. Based on the assumption of global identifiability, note that the set $\mathcal{S}(\theta^*,\theta)$ is nonempty. Recall that v_θ is any agent for which $K_i(\theta^*,\theta), i \in \mathcal{S}(\theta^*,\theta)$ is maximum, and note that we have already established that the assertion of Theorem 1, namely inequality (5), holds for agent v_θ in Lemma 3. Now consider an agent $i \in \mathcal{V} \setminus \{v_\theta\}$, and notice that if $t \geq (n-1)T$, then $c_{i,t}(\theta)$ is uniformly bounded based on Lemma 4. Thus, the assertions in Lemma 5 hold for all $t \geq (n-1)T$. Taking the natural log on both sides of (38), dividing throughout by t, and simplifying, we obtain the following for all $t \geq (n-1)T$:

$$-\frac{\log \mu_{i,t}(\theta)}{t} \ge -\frac{\log \mu_{v_{\theta},a_{i,t}(\theta)}(\theta)}{t} + \sum_{\tau=a_{i,t}(\theta)+1}^{t} \frac{\log \eta_{x(\tau),\tau}(\theta^{\star})}{t},$$
(41)

where $a_{i,t}(\theta) = t - c_{i,t}(\theta)$, $\eta_{i,t}(\theta^*)$ is as defined in (39), and $x(\tau), \tau \in \{a_{i,t}(\theta) + 1, \dots, t\}$, are agents in the path linking v_{θ} to i over $[a_{i,t}(\theta), t-1]$. For the remainder of the proof, to lighten the notation, let us drop the subscript on v_{θ} , and let $a(t) = a_{i,t}(\theta)$. Based on (4), we then have:

$$\mu_{v,a(t)}(\theta) \le \frac{\pi_{v,a(t)}(\theta)}{\eta_{v,a(t)}(\theta^*)}.$$
(42)

A bit of straightforward algebra then yields:

$$-\frac{\log \mu_{v,a(t)}(\theta)}{t} \ge -\frac{\log \pi_{v,t}(\theta)}{t} + \frac{\log \frac{\pi_{v,t}(\theta)}{\pi_{v,a(t)}(\theta)}}{t} + \frac{\log \eta_{v,a(t)}(\theta^*)}{t}.$$

$$(43)$$

Combining (41) and (43), we obtain for $t \ge (n-1)T$:

$$-\frac{\log \mu_{i,t}(\theta)}{t} \ge -\frac{\log \pi_{v,t}(\theta)}{t} + b(t), \tag{44}$$

where $b(t) = b_1(t) + b_2(t) + b_3(t)$,

$$b_1(t) = \sum_{\tau=a(t)+1}^{t} \frac{\log \eta_{x(\tau),\tau}(\theta^*)}{t}, \ b_2(t) = \frac{\log \frac{\pi_{v,t}(\theta)}{\pi_{v,a(t)}(\theta)}}{t}, \quad (45)$$

and

$$b_3(t) = \frac{\log \eta_{v,a(t)}(\theta^*)}{t}.$$
 (46)

We now argue that each of the terms $b_1(t), b_2(t)$, and $b_3(t)$ converge to 0 almost surely as $t \to \infty$. To do so, recall that the set $\bar{\Omega} \subseteq \Omega$ in Lemma 2 has measure 1. In what follows, we

prove that $b_1(t), b_2(t)$, and $b_3(t)$ converge to 0 for each sample path $\omega \in \bar{\Omega}$. Accordingly, fix $\omega \in \bar{\Omega}$, and recall $\eta(\omega) \in (0,1)$ and $t'(\omega) \in (0,\infty)$ from Lemma 2. Suppose $t > t'(\omega) + 2\bar{T}$, where $\bar{T} = (n-1)T$. We then claim the following:

$$\pi_{l,\tau}(\theta^*) \ge \eta(\omega), \mu_{l,\tau}(\theta^*) \ge \eta(\omega), \quad \forall l \in \mathcal{V}, \quad \forall \tau \ge a(t).$$
(47)

To see why this is true, notice that based on Lemma 4, the following holds when $t > t'(\omega) + 2\bar{T}$:

$$a(t) = t - c_{i,t}(\theta) \ge t - 2\bar{T} > t'(\omega).$$
 (48)

The claim regarding (47) then follows readily from (24) in Lemma 2. Based on the previous discussion, and referring to (39), we immediately note that when $t > t'(\omega) + 2\bar{T}$,

$$\eta_{l,\tau}(\theta^*) > \eta(\omega), \quad \forall l \in \mathcal{V}, \quad \forall \tau > a(t).$$
(49)

For establishing the convergence of $b_1(t)$, $b_2(t)$, and $b_3(t)$, suppose $t > t'(\omega) + 2\bar{T}$. Regarding $b_1(t)$, we then observe:

$$|b_{1}(t)| = \left| \sum_{\tau=a(t)+1}^{t} \frac{\log \eta_{x(\tau),\tau}(\theta^{\star})}{t} \right|$$

$$\stackrel{(a)}{\leq} \sum_{\tau=a(t)+1}^{t} \frac{\left| \log \eta_{x(\tau),\tau}(\theta^{\star}) \right|}{t}$$

$$\stackrel{(b)}{\leq} \frac{(t-a(t))}{t} \log \frac{1}{\eta(\omega)}$$

$$\stackrel{(c)}{\leq} \frac{2\bar{T}}{t} \log \frac{1}{\eta(\omega)},$$

$$(50)$$

where (a) follows from the triangle inequality, (b) follows from (49), and (c) follows from (48). From (50), we immediately note that $b_1(t) \to 0$ along ω . Let us now turn our attention to $b_2(t)$, and take note of the following:

$$|b_{2}(t)| \stackrel{(a)}{=} \frac{1}{t} \left| \log \frac{\pi_{v,t}(\theta^{\star})}{\pi_{v,a(t)}(\theta^{\star})} + \sum_{\tau=a(t)+1}^{t} \log \frac{l_{v}(s_{v,\tau}|\theta)}{l_{v}(s_{v,\tau}|\theta^{\star})} \right|$$

$$\stackrel{(b)}{\leq} \frac{1}{t} \left| \log \frac{\pi_{v,t}(\theta^{\star})}{\pi_{v,a(t)}(\theta^{\star})} \right| + \frac{1}{t} \sum_{\tau=a(t)+1}^{t} \left| \log \frac{l_{v}(s_{v,\tau}|\theta)}{l_{v}(s_{v,\tau}|\theta^{\star})} \right|$$

$$\stackrel{(c)}{\leq} \frac{1}{t} \log \frac{1}{\eta(\omega)} + \frac{(t-a(t))L}{t}$$

$$\stackrel{(d)}{\leq} \frac{1}{t} \left(\log \frac{1}{\eta(\omega)} + L\bar{T} \right),$$

where (a) follows from (22) and some simple manipulations, (b) is a consequence of the triangle inequality, (c) follows from (1) and (47), and (d) follows from (48). Based on (51), we then note that $b_2(t) \to 0$ along ω . Finally, the fact that $b_3(t)$ converges to 0 along ω follows immediately by appealing to (49). We have thus established that $b(t) \to 0$ almost surely. The desired conclusion then follows by taking the limit inferior on both sides of (44), and noting that

$$\lim_{t \to \infty} -\frac{\log \pi_{v,t}(\theta)}{t} = \lim_{t \to \infty} -\frac{1}{t} \rho_{v,t}(\theta) = K_v(\theta^*, \theta) \text{ a.s.}, \quad (52)$$

where $\rho_{v,t}(\theta)$ is as defined in Lemma 1. The fact that $\mu_{i,t}(\theta) \to 0$ is immediate, since $K_v(\theta^*,\theta) > 0$ based on global identifiability. The aforementioned analysis applies identically to each $\theta \in \Theta \setminus \{\theta^*\}$. This establishes consistency of our rule.

APPENDIX B PROOF OF THEOREM 2

To prove Theorem 2, we will make use of one of Littlewood's three principles: Every pointwise convergent sequence of measurable functions is nearly uniformly convergent.

Theorem 5 (*Egoroff's Theorem*): [43, Ch. 18] Let (X, \mathcal{M}, μ) be a finite measure space and $\{f_n\}$ a sequence of measurable functions on X that converge pointwise a.e. (almost everywhere) on X to a function f that is finite a.e. on X. Then, for each $\epsilon > 0$, there is a measurable subset X_{ϵ} of X for which $f_n \to f$ uniformly on X_{ϵ} , and $\mu(X_{\epsilon}) \geq 1 - \epsilon$.

Proof: (**Theorem 2**): Consider a $\theta \in \Theta \setminus \{\theta^*\}$, and recall that $K_{v_{\theta}}(\theta^*, \theta) = \max_{l \in S(\theta^*, \theta)} K_l(\theta^*, \theta) = \bar{K}(\theta^*, \theta)$. We only prove the result for $i \in \mathcal{V} \setminus \{v_{\theta}\}$, since the argument for agent v_{θ} will be similar. To this end, let us fix an agent $i \in \mathcal{V} \setminus \{v_{\theta}\}$. We adhere to the notation used in the proof of Lemma 1, and for simplicity assume that the initial local belief vectors $\pi_{i,0}, i \in \mathcal{V}$ are uniform distributions over the hypothesis set Θ ; our subsequent arguments will continue to hold (with simple modifications) under the more general assumption on priors in line 1 of Algorithm 1. We immediately note that based on the assumption of uniform priors, $\rho_{i,0}(\theta) = 0, \forall i \in \mathcal{V}$. Now referring to inequality (44) in the proof of Theorem 1, we obtain the following for $t \geq (n-1)T$:

$$\mathbb{P}^{\theta^{\star}} \left(-\frac{\log \mu_{i,t}(\theta)}{t} \leq \bar{K}(\theta^{\star}, \theta) - \frac{\epsilon}{2} + b(t) \right) \\
\stackrel{(a)}{\leq} \mathbb{P}^{\theta^{\star}} \left(-\frac{\log \pi_{v_{\theta},t}(\theta)}{t} \leq \bar{K}(\theta^{\star}, \theta) - \frac{\epsilon}{2} \right) \\
\stackrel{(b)}{\leq} \mathbb{P}^{\theta^{\star}} \left(-\frac{\rho_{v_{\theta},t}(\theta)}{t} \leq \bar{K}(\theta^{\star}, \theta) - \frac{\epsilon}{2} \right) \\
\stackrel{(c)}{=} \mathbb{P}^{\theta^{\star}} \left(\frac{1}{t} \sum_{k=1}^{t} \lambda_{v_{\theta},k}(\theta) - (-K_{v_{\theta}}(\theta^{\star}, \theta)) \geq \frac{\epsilon}{2} \right) \\
\stackrel{(d)}{\leq} \exp\left(-\frac{\epsilon^{2}t}{8L^{2}} \right). \tag{53}$$

In the above steps, (a) follows directly from (44), and (b) follows by noting that based on the definition of $\rho_{v_{\theta},t}(\theta)$,

$$\frac{\log \pi_{v_{\theta},t}(\theta)}{t} \le \frac{\rho_{v_{\theta},t}(\theta)}{t}, \quad \forall t \in \mathbb{N}.$$
 (54)

Step (c) follows directly from (22) with $\rho_{v_{\theta},0}(\theta)=0$. Finally, noting that $\frac{1}{t}\sum_{k=1}^{t}\lambda_{v_{\theta},k}(\theta)\to -K_{v_{\theta}}(\theta^{\star},\theta)$ a.s. (as argued in the proof of Lemma 1), using the fact that $|\lambda_{v_{\theta},t}(\theta)|\leq L, \forall t\in\mathbb{N}_+$ based on (1), and applying Hoeffding's inequality [44, Th. 2], leads to (d). Now recall from the proof of Theorem 1 that $b(t)\to 0$ almost surely. Appealing to Egoroff's theorem, we then infer that given any arbitrarily small $\delta\in(0,1)$, there exists a set $\Omega'(\delta)\subseteq\Omega$ of $\mathbb{P}^{\theta^{\star}}$ -measure at least $(1-\delta)$, such that b(t) converges to 0 uniformly on $\Omega'(\delta)$. Thus, given any $\epsilon>0$,

there exists a ω -independent constant $t(\epsilon,\delta) \in (0,\infty)$, such that $|b(t)| \leq \frac{\epsilon}{2}, \forall t \geq t(\epsilon,\delta)$, along each sample path $\omega \in \Omega'(\delta)$. Setting $t'(\epsilon,\delta,n,T) = \max\{t(\epsilon,\delta),(n-1)T\}$, and referring to (53), we immediately obtain that $\forall t \geq t'(\epsilon,\delta,n,T)$,

$$\mathbb{P}^{\theta^{\star}} \left(\left\{ -\frac{\log \mu_{i,t}(\theta)}{t} \leq \bar{K}(\theta^{\star}, \theta) - \epsilon \right\} \cap \Omega'(\delta) \right) \\
\leq \mathbb{P}^{\theta^{\star}} \left(\left\{ -\frac{\log \mu_{i,t}(\theta)}{t} \leq \bar{K}(\theta^{\star}, \theta) - \frac{\epsilon}{2} + b(t) \right\} \cap \Omega'(\delta) \right) \\
\leq \mathbb{P}^{\theta^{\star}} \left(-\frac{\log \mu_{i,t}(\theta)}{t} \leq \bar{K}(\theta^{\star}, \theta) - \frac{\epsilon}{2} + b(t) \right) \\
\leq \exp \left(-\frac{\epsilon^{2}t}{8L^{2}} \right).$$
(55)

Taking the natural log on both sides of the resulting inequality, dividing throughout by t, simplifying, and then taking the limit inferior on both sides, leads to the desired result.

APPENDIX C PROOF OF THEOREM 4

Proof: Consider an f-local adversarial set $A \subset V$, and let $\mathcal{R} = \mathcal{V} \setminus A$. We study two separate cases.

<u>Case 1</u>: Consider a regular agent $i \in \mathcal{R}$ such that $|\mathcal{N}_i| < (2f+1)$. Based on the hypothesis of the theorem, we claim that $i \in \mathcal{S}(\theta_p,\theta_q)$, for every pair $\theta_p,\theta_q \in \Theta$. We prove this claim via contradiction. To do so, suppose there exists a pair $\theta_p,\theta_q \in \Theta$, such that $i \in \mathcal{V} \setminus \mathcal{S}(\theta_p,\theta_q)$. As $|\mathcal{N}_i| < (2f+1)$, the set $\{i\}$ is clearly not (2f+1)-reachable (see Definition 2). Thus, \mathcal{G} is not strongly (2f+1)-robust w.r.t. the source set $\mathcal{S}(\theta_p,\theta_q)$, a fact that contradicts the hypothesis of the theorem. Thus, we have established that if the graph-theoretic condition identified in the theorem is met, then regular agents with fewer than (2f+1) neighbors can distinguish between every pair of hypotheses. For such agents, the assertion of the theorem then follows directly from Lemma 1, and update rules (3) and (14).

<u>Case 2</u>: We now focus only on regular agents i satisfying $|\mathcal{N}_i| \geq (2f+1)$. A key property of the LFRHE algorithm (Algorithm 2) that will be used throughout the proof is as follows. For any $i \in \mathcal{R}$, and any $\theta \in \Theta$, the filtering operation in line 7 of Algorithm 2 ensures that at each $t \in \mathbb{N}$, we have

$$\mu_{j,t}(\theta) \in \text{Conv}(\Psi_{i,t}^{\theta}), \quad \forall j \in \mathcal{M}_{i,t}^{\theta},$$
 (56)

where

$$\Psi_{i,t}^{\theta} \triangleq \{\mu_{l,t}(\theta) : l \in \mathcal{N}_i \cap \mathcal{R}\},\tag{57}$$

and $\operatorname{Conv}(\Psi_{i,t}^{\theta})$ is used to denote the convex hull formed by the points in the set $\Psi_{i,t}^{\theta}$ (recall that $\mathcal{M}_{i,t}^{\theta}$ was defined in line 8 of Algorithm 2 to be the set of agents in \mathcal{N}_i whose beliefs are retained by agent i after it removes the highest f and lowest f beliefs $\mu_{j,t}(\theta), j \in \mathcal{N}_i$). In words, any neighboring belief (on a particular hypothesis) that agent i uses in the update rule (13) lies in the convex hull of the actual beliefs of its regular neighbors (on that particular hypothesis). To see why (56) is true, partition the neighbor set \mathcal{N}_i of a regular agent into three sets $\mathcal{U}_{i,t}^{\theta}, \mathcal{M}_{i,t}^{\theta}$, and $\mathcal{J}_{i,t}^{\theta}$ as follows. Sets $\mathcal{U}_{i,t}^{\theta}$ and $\mathcal{J}_{i,t}^{\theta}$ are each of cardinality f, and contain neighbors of agent i that transmit the highest f and the

lowest f actual beliefs, respectively, on the hypothesis θ , to agent i at time-step t. The set $\mathcal{M}_{i,t}^{\theta}$ contains the remaining neighbors of agent i, and is non-empty at every time-step since $|\mathcal{N}_i| \geq (2f+1)$. If $\mathcal{M}_{i,t}^{\theta} \cap \mathcal{A} = \emptyset$, then (56) holds trivially. Thus, consider the case when there are adversaries in the set $\mathcal{M}_{i,t}^{\theta}$, i.e., $\mathcal{M}_{i,t}^{\theta} \cap \mathcal{A} \neq \emptyset$. Given the f-locality of the adversarial model, and the nature of the filtering operation in the LFRHE algorithm, we infer that for each $j \in \mathcal{M}_{i,t}^{\theta} \cap \mathcal{A}$, there exist regular agents $u, v \in \mathcal{N}_i \cap \mathcal{R}$, such that $u \in \mathcal{U}_{i,t}^{\theta}$, $v \in \mathcal{J}_{i,t}^{\theta}$, and $\mu_{v,t}(\theta) \leq \mu_{j,t}(\theta) \leq \mu_{u,t}(\theta)$. This establishes our claim regarding (56).

With the above property in hand, let $\bar{\Omega} \subseteq \Omega$ denote the set of sample paths for which assertions (i)-(iii) in Lemma 1 (see Appendix A) hold when restricted to the set of regular agents \mathcal{R} . Since the evolution of the local beliefs are unaffected by the presence of adversaries, Lemma 1 implies $\mathbb{P}^{\theta^*}(\bar{\Omega}) = 1$. Now as in Lemma 2, fix a sample path $\omega \in \bar{\Omega}$. Define $\gamma_1 \triangleq$ $\min_{i \in \mathcal{R}} \pi_{i,0}(\theta^*)$, pick a small number $\delta > 0$ satisfying $\delta < \gamma_1$, and observe that arguments similar to those in the proof of Lemma 2 imply the existence of a time-step $t'(\omega)$, such that for all $t \geq t'(\omega), \pi_{i,t}(\theta^*) \geq \gamma_1 - \delta > 0, \forall i \in \mathcal{R}$. Let $\gamma_2(\omega) \triangleq$ $\min_{i\in\mathcal{R}}\{\mu_{i,t'(\omega)}(\theta^*)\}$. As before, we claim $\gamma_2(\omega)>0$. To establish this claim, we need to answer the following question: Can an adversarial agent cause its out-neighbors to set their actual beliefs on θ^* to be 0 by setting its own actual belief on θ^* to be 0? We argue that this is impossible under the LFRHE algorithm. By way of contradiction, suppose there exists a time-step $\bar{t}(\omega)$ satisfying:

$$\bar{t}(\omega) = \min\{t \in \mathbb{N} : \exists i \in \mathcal{R} \text{ with } \mu_{i,t}(\theta^*) = 0\}.$$
 (58)

In words, $\bar{t}(\omega)$ represents the first time-step when some regular agent i sets its actual belief on the true hypothesis to be zero. Clearly, $\bar{t}(\omega) \neq 0$ based on line 1 of Algorithm 2. Suppose $\bar{t}(\omega)$ is some positive integer, and focus on how agent i updates $\mu_{i,\bar{t}(\omega)}(\theta^*)$ based on (13). Following similar arguments as in the proof of Lemma 2, we know that $\pi_{i,t}(\theta^*)>0, \forall t\in\mathbb{N}, \forall i\in\mathcal{R}.$ At the same time, every belief featuring in the set $\Psi_{i,\bar{t}(\omega)-1}^{\theta^*}$ (as defined in equation (57)) is strictly positive based on the way $\bar{t}(\omega)$ is defined. In light of the previous arguments, and based on (56) and (57), we infer:

$$\min\{\{\mu_{j,\bar{t}(\omega)-1}(\theta^{\star})\}_{j\in\mathcal{M}_{i,\bar{t}(\omega)-1}^{\theta^{\star}}}, \pi_{i,\bar{t}(\omega)}(\theta^{\star})\} > 0.$$
 (59)

Thus, based on (13), we must have $\mu_{i,\bar{t}(\omega)}(\theta^*) > 0$, yielding the desired contradiction. With $\eta(\omega) \triangleq \min\{\gamma_1 - \delta, \gamma_2(\omega)\} > 0$, one can easily verify the following by referring to (13):

$$\mu_{i,t}(\theta^*) > \eta(\omega), \quad \forall t > t'(\omega) \quad \forall i \in \mathcal{R}.$$
 (60)

In particular, (60) follows by (i) noting that for each $i \in \mathcal{R}$, $\pi_{i,t'(\omega)+1}(\theta^\star) \geq \eta(\omega)$, and each belief featuring in the set $\Psi_{i,t'(\omega)}^{\theta^\star}$ is lower bounded by $\eta(\omega)$, (ii) leveraging (56) and (57), and (iii) using a similar string of arguments as those used to arrive at (25). Thus, we have established an analogous result as in Lemma 2 for the regular agents.

To proceed, let us fix a false hypothesis $\theta \neq \theta^*$, and define $\tilde{K}(\theta^*,\theta) \triangleq \min_{v \in \mathcal{S}(\theta^*,\theta) \cap \mathcal{R}} K_v(\theta^*,\theta)$. Then, given any $\epsilon > 0$, Lemma 1 implies the existence of a time-step $\tilde{t}_1(\omega,\theta,\epsilon)$, such

that:

$$\pi_{i,t}(\theta) < e^{-(\tilde{K}(\theta^*,\theta) - \epsilon)t}, \quad \forall t \ge \tilde{t}_1(\omega,\theta,\epsilon), \quad \forall i \in \mathcal{S}(\theta^*,\theta) \cap \mathcal{R}.$$
(61)

Let $\tilde{t}_2 = \max\{t'(\omega), \tilde{t}_1(\omega, \theta, \epsilon)\}$, where we have suppressed the dependence of \tilde{t}_2 on ω, θ and ϵ . For any agent $i \in \mathcal{S}(\theta^*, \theta) \cap \mathcal{R}$, observe that based on (56), (57), and (60):

$$\min\{\{\mu_{j,t}(\theta^*)\}_{i\in\mathcal{M}_{i,1}^{\theta^*}}, \pi_{i,t+1}(\theta^*)\} \ge \eta(\omega), \quad \forall t \ge \tilde{t}_2.$$
 (62)

Combining the aforementioned with a similar line of argument as used to arrive at (28), we obtain

$$\mu_{i,t}(\theta) < C_1(\omega)e^{-(\tilde{K}(\theta^*,\theta)-\epsilon)t}, \quad \forall t \ge \tilde{t}_2 + 1,$$

$$\forall i \in \mathcal{S}(\theta^*,\theta) \cap \mathcal{R}, \tag{63}$$

where $C_1(\omega) = \eta(\omega)^{-1}$. If $\mathcal{V} \setminus \mathcal{S}(\theta^*, \theta)$ is empty, then we are essentially done. Else, define

$$\mathcal{L}_1(\theta^*, \theta) \triangleq \{ i \in \mathcal{V} \setminus \mathcal{S}(\theta^*, \theta) : |\mathcal{N}_i \cap \mathcal{S}(\theta^*, \theta)| \ge (2f + 1) \}.$$
(64)

Whenever $\mathcal{V}\setminus\mathcal{S}(\theta^\star,\theta)$ is non-empty, we claim that $\mathcal{L}_1(\theta^\star,\theta)$ (as defined above) is also non-empty based on the hypothesis of the theorem. To see this, note that if $\mathcal{L}_1(\theta^\star,\theta)$ is empty, then $\mathcal{C}=\mathcal{V}\setminus\mathcal{S}(\theta^\star,\theta)$ is not (2f+1)-reachable, violating the fact that \mathcal{G} is strongly (2f+1)-robust w.r.t. $\mathcal{S}(\theta^\star,\theta)$. We claim that the following holds for each $i\in\mathcal{L}_1(\theta^\star,\theta)\cap\mathcal{R}$:

$$\min_{j \in \mathcal{M}_{i,t}^{\theta}} \mu_{j,t}(\theta) < C_1(\omega) e^{-(\tilde{K}(\theta^*,\theta) - \epsilon)t}, \quad \forall t \ge \tilde{t}_2 + 1. \quad (65)$$

To verify the above claim, pick any agent $i \in \mathcal{L}_1(\theta^*, \theta) \cap \mathcal{R}$, and suppose $t \geq \tilde{t}_2 + 1$. When $|\mathcal{M}_{i,t}^{\theta} \cap \{\mathcal{S}(\theta^{\star}, \theta) \cap \mathcal{R}\}| > 0$, the claim follows immediately based on (63). Consider the case when $|\mathcal{M}_{i,t}^{\theta} \cap \{\mathcal{S}(\theta^{\star},\theta) \cap \mathcal{R}\}| = 0$. Since $i \in \mathcal{L}_1(\theta^{\star},\theta)$, it has at least (2f + 1) neighbors in $S(\theta^*, \theta)$, out of which at least $f + \theta$ 1 are regular based on the f-locality of the adversarial model. Since the set $\mathcal{J}_{i,t}^{\theta}$ has cardinality f, it must then be that $|\mathcal{U}_{i,t}^{\theta}|$ $\{\mathcal{S}(\theta^\star,\theta)\cap\mathcal{R}\}|>0.\operatorname{Let}u\in\mathcal{U}_{i,t}^\theta\cap\{\mathcal{S}(\theta^\star,\theta)\cap\mathcal{R}\}.\operatorname{Based on }$ the way $\mathcal{M}_{i,t}^{\theta}$ is defined, it must be that $\mu_{j,t}(\theta) \leq \mu_{u,t}(\theta) < 0$ $C_1(\omega)e^{-(\tilde{K}(\dot{\theta}^\star,\theta)-\epsilon)t}, \forall j\in\mathcal{M}_{i,t}^{\theta}$, where the last inequality follows from (63). This establishes our claim regarding (65). Now consider the update of $\mu_{i,t+1}(\theta)$ based on (13), when $t \geq \tilde{t}_2 + 1$. In light of the above arguments, the numerator of the fraction on the RHS of (13) is upper-bounded by $C_1(\omega)e^{-(\tilde{K}(\theta^*,\theta)-\epsilon)t}$, whereas the denominator is lower-bounded by $\eta(\omega)$. We conclude that for all $i \in \mathcal{L}_1(\theta^*, \theta) \cap \mathcal{R}$:

$$\mu_{i,t}(\theta) < (C_1(\omega))^2 C_2(\theta, \epsilon) e^{-(\tilde{K}(\theta^*, \theta) - \epsilon)t}, \quad \forall t \ge \tilde{t}_2 + 2,$$
(66)

where $C_2(\theta, \epsilon) = e^{(\tilde{K}(\theta^*, \theta) - \epsilon)}$. With $\mathcal{L}_0(\theta^*, \theta) \triangleq \mathcal{S}(\theta^*, \theta)$, we recursively define the sets $\mathcal{L}_r(\theta^*, \theta), 1 \leq r \leq (n-1)$ as:

$$\mathcal{L}_{r}(\theta^{*},\theta) \triangleq \left\{ i \in \mathcal{V} \setminus \bigcup_{q=0}^{r-1} \mathcal{L}_{q}(\theta^{*},\theta) : | \mathcal{N}_{i} \right.$$

$$\left. \cap \left\{ \bigcup_{q=0}^{r-1} \mathcal{L}_{q}(\theta^{*},\theta) \right\} | \geq (2f+1) \right\}. \tag{67}$$

We claim that the following is true for all $i \in \mathcal{L}_r(\theta^*, \theta) \cap \mathcal{R}$:

$$\mu_{i,t}(\theta) < (C_1(\omega))^{r+1} (C_2(\theta, \epsilon))^r e^{-(\tilde{K}(\theta^*, \theta) - \epsilon)t},$$

$$\forall t \ge \tilde{t}_2 + (r+1). \tag{68}$$

To prove the claim, we proceed via induction on r. The base cases when $r \in \{0,1\}$ have already been established. Suppose (68) holds for all $r \in \{0, ..., m-1\}$, where $m \in \{2, ..., n-1\}$. The claim easily extends to the case when r = m by noting that (i) $\mathcal{L}_m(\theta^\star,\theta)$ is non-empty if $\mathcal{V}\setminus\{\bigcup_{q=0}^{(m-1)}\mathcal{L}_q(\theta^\star,\theta)\}$ is non-empty (based on the hypothesis of the theorem), (ii) any agent $i \in \mathcal{L}_m(\theta^*, \theta) \cap \mathcal{R}$ has at least (2f + 1) neighbors in the set $\bigcup_{q=0}^{(m-1)} \mathcal{L}_q(\theta^{\star}, \theta)$, of which at least f+1 are regular (based on the f-locality of the adversarial model), and (iii) using the induction hypothesis and arguments similar to those used to arrive at (66). We have thus verified the correctness of (68). Now taking the natural log on both sides of (68), dividing throughout by t, simplifying, and then taking the limit inferior on both sides of the resulting inequality immediately leads to (15). Finally, to complete the proof, it suffices to note that $\bigcup_{q=0}^{(n-1)} \mathcal{L}_q(\theta^*, \theta) = \mathcal{R}.$

REFERENCES

- [1] A. Jadbabaie, P. Molavi, A. Sandroni, and A. Tahbaz-Salehi, "Non-Bayesian social learning," *Games Econ. Behav.*, vol. 76, no. 1, pp. 210–225, 2012.
- [2] A. Jadbabaie, P. Molavi, and A. Tahbaz-Salehi, "Information heterogeneity and the speed of learning in social networks," Columbia Bus. School Res. Paper no. 13-28, 2013.
- [3] P. Molavi, A. Tahbaz-Salehi, and A. Jadbabaie, "A theory of non-Bayesian social learning," *Econometrica*, vol. 86, no. 2, pp. 445–490, 2018.
- [4] V. V. Veeravalli, T. Basar, and H. V. Poor, "Decentralized sequential detection with a fusion center performing the sequential test," *IEEE Trans. Inf. Theory*, vol. 39, no. 2, pp. 433–442, Mar. 1993.
- [5] R. Viswanathan and P. K. Varshney, "Distributed detection with multiple sensors Part I. Fundamentals," *Proc. IEEE*, vol. 85, no. 1, pp. 54–63, Jan. 1997.
- [6] J. N. Tsitsiklis, "Decentralized detection by a large number of sensors," Math. Control, Signals, Syst., vol. 1, no. 2, pp. 167–182, 1988.
- [7] Q. Liu, A. Fang, L. Wang, and X. Wang, "Social learning with time-varying weights," J. Syst. Sci. Complexity, vol. 27, no. 3, pp. 581–593, 2014.
- [8] H. Salami, B. Ying, and A. H. Sayed, "Social learning over weakly connected graphs," *IEEE Trans. Signal Inf. Process. Netw.*, vol. 3, no. 2, pp. 222–238, Jun. 2017.
- [9] K. R. Rad and A. Tahbaz-Salehi, "Distributed parameter estimation in networks," in *Proc.* 49th IEEE Decis. Control Conf., 2010, pp. 5050–5055.
- [10] S. Shahrampour and A. Jadbabaie, "Exponentially fast parameter estimation in networks using distributed dual averaging," in *Proc. 52nd Decis. Control Conf.*, 2013, pp. 6196–6201.
- [11] S. Shahrampour, A. Rakhlin, and A. Jadbabaie, "Distributed detection: Finite-time analysis and impact of network topology," *IEEE Trans. Autom. Control*, vol. 61, no. 11, pp. 3256–3268, Nov. 2016.
- [12] A. Nedić, A. Olshevsky, and C. A. Uribe, "Fast convergence rates for distributed non-Bayesian learning," *IEEE Trans. Autom. Control*, vol. 62, no. 11, pp. 5538–5553, Nov. 2017.
- [13] A. Nedić, A. Olshevsky, and C. A. Uribe, "Nonasymptotic convergence rates for cooperative learning over time-varying directed graphs," in *Proc. Amer. Control Conf.*, 2015, pp. 5884–5889.
- [14] A. Lalitha, T. Javidi, and A. Sarwate, "Social learning and distributed hypothesis testing," *IEEE Trans. Inf. Theory*, vol. 64, no. 9, pp. 6161–6179, Sep. 2018.
- [15] A. Lalitha and T. Javidi, "Large deviation analysis for learning rate in distributed hypothesis testing," in *Proc. 49th Asilomar Conf. Signals, Syst.*, *Comput.*, 2015, pp. 1065–1069.
- [16] L. Su and N. H. Vaidya, "Defending Non-Bayesian learning against adversarial attacks," *Distrib. Comput.*, vol. 32, pp. 277–289, 2019.

- [17] R. Olfati-Saber, E. Franco, E. Frazzoli, and J. S. Shamma, "Belief consensus and distributed hypothesis testing in sensor networks," in *Proc. Netw. Embedded Sens. Control*, 2006, pp. 169–182.
- [18] V. Saligrama, M. Alanyali, and O. Savas, "Distributed detection in sensor networks with packet losses and finite capacity links," *IEEE Trans. Signal Process.*, vol. 54, no. 11, pp. 4118–4132, Nov. 2006.
- [19] G. L. Gilardoni and M. K. Clayton, "On reaching a consensus using DeGroot's iterative pooling," Ann. Statist., vol. 21, pp. 391–401, 1993.
- [20] N. A. Lynch, Distributed Algorithms. San Mateo, CA, USA: Morgan Kaufmann, 1996.
- [21] A. Mitra, J. A. Richards, and S. Sundaram, "A new approach for distributed hypothesis testing with extensions to Byzantine-resilience," in *Proc. Amer. Control Conf.*, 2019, pp. 261–266.
- [22] A. Nedić, A. Olshevsky, and C. A. Uribe, "Distributed learning for cooperative inference," 2017, arXiv:1704.02718.
- [23] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. Hoboken, NJ, USA: Wiley, 2012.
- [24] A. Jadbabaie, J. Lin, and A. S. Morse, "Coordination of groups of mobile autonomous agents using nearest neighbor rules," *IEEE Trans. Autom. Control*, vol. 48, no. 6, pp. 988–1001, Jun. 2003.
- [25] A. Nedić and A. Olshevsky, "Distributed optimization over time-varying directed graphs," *IEEE Trans. Autom. Control*, vol. 60, no. 3, pp. 601–615, Mar. 2015.
- [26] P. Molavi, A. Jadbabaie, K. R. Rad, and A. Tahbaz-Salehi, "Reaching consensus with increasing information," *IEEE J. Sel. Topics Signal Process.*, vol. 7, no. 2, pp. 358–369, Apr. 2013.
- [27] S. Park and N. C. Martins, "Design of distributed LTI observers for state omniscience," *IEEE Trans. Autom. Control*, vol. 62, no. 2, pp. 561–576, Feb. 2017.
- [28] A. Mitra and S. Sundaram, "Distributed observers for LTI systems," *IEEE Trans. Autom. Control*, vol. 63, no. 11, pp. 3689–3704, Nov. 2018.
- [29] D. Acemoglu, A. Ozdaglar, and A. ParandehGheibi, "Spread of (mis) information in social networks," *Games Econ. Behav.*, vol. 70, no. 2, pp. 194–227, 2010.
- [30] D. Dolev, N. A. Lynch, S. S. Pinter, E. W. Stark, and W. E. Weihl, "Reaching approximate agreement in the presence of faults," *J. ACM*, vol. 33, no. 3, pp. 499–516, 1986.
- [31] N. H. Vaidya, L. Tseng, and G. Liang, "Iterative approximate Byzantine consensus in arbitrary directed graphs," in *Proc. ACM Symp. Princ. Distrib. Comput.*, 2012, pp. 365–374.
- [32] H. J. LeBlanc, H. Zhang, X. Koutsoukos, and S. Sundaram, "Resilient asymptotic consensus in robust networks," *IEEE J. Sel. Areas Commun.*, vol. 31, no. 4, pp. 766–781, Apr. 2013.
- [33] S. M. Dibaji and H. Ishii, "Resilient consensus of second-order agent networks: Asynchronous update rules with delays," *Automatica*, vol. 81, pp. 123–132, 2017.
- [34] L. Su and N. H. Vaidya, "Fault-tolerant multi-agent optimization: Optimal iterative distributed algorithms," in *Proc. ACM Symp. Princ. Distrib. Comput.*, 2016, pp. 425–434.
- [35] S. Sundaram and B. Gharesifard, "Distributed optimization under adversarial nodes," *IEEE Trans. Autom. Control*, vol. 64, no. 3, pp. 1063–1076, Mar. 2019.
- [36] A. Mitra and S. Sundaram, "Byzantine-resilient distributed observers for LTI systems," *Automatica*, vol. 108, 2019, Art. no. 108487.
- [37] J. Usevitch and D. Panagou, "Resilient leader-follower consensus to arbitrary reference values," in *Proc. Annu. Amer. Control Conf.*, 2018, pp. 1292–1298.
- [38] C.-Y. Koo, "Broadcast in radio networks tolerating Byzantine adversarial behavior," in *Proc. ACM Symp. Princ. Distrib. Comput.*, 2004, pp. 275– 282.
- [39] H. Park and S. A. Hutchinson, "Fault-tolerant rendezvous of multirobot systems," *IEEE Trans. Robot.*, vol. 33, no. 3, pp. 565–582, Jun. 2017.
- [40] N. Vaidya, "Matrix representation of iterative approximate Byzantine consensus in directed graphs," 2012, arXiv:1203.1888.
- [41] W. Mulzer and D. Werner, "Approximating Tverberg points in linear time for any fixed dimension," *Discrete Comput. Geometry*, vol. 50, no. 2, pp. 520–535, 2013.
- [42] A. Mitra, J. A. Richards, and S. Sundaram, "A communication-efficient algorithm for exponentially fast non-Bayesian learning in networks," in *Proc. 58th Conf. Decis. Control.*, 2019, pp. 8347–8352.
- 43] H. Royden and P. Fitzpatrick, *Real Analysis*. Englewood Cliffs, NJ, USA: Prentice-Hall, 2010.
- [44] W. Hoeffding, "Probability inequalities for sums of bounded random variables," in *The Collected Works of Wassily Hoeffding*. Berlin, Germany: Springer-Verlag, 1994, pp. 409–426.



Aritra Mitra received the B.E. degree in electrical engineering from Jadavpur University, Kolkata, India, in 2013, and the M.Tech. degree in electrical engineering from the Indian Institute of Technology Kanpur, Kanpur, India, in 2015. He is currently working toward the Ph.D. degree in electrical engineering with the School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN, USA.

His current research interests include the design of distributed algorithms for estimation, in-

ference and learning, networked control systems, and secure control.

Mr. Mitra was a recipient of the University Gold Medal at Jadavpur University and the Academic Excellence Award at IIT Kanpur.



John A. Richards received the S.B. and M.Eng. degrees in electrical engineering, and the Ph.D. degree in electrical engineering and computer science from the Massachusetts Institute of Technology, Cambridge, MA, USA, in 1996 and 2001, respectively.

He is currently a Distinguished Member of Technical Staff with the Autonomy for Hypersonics Department, Sandia National Laboratories, Albuquerque, NM, USA. He has led numerous projects involving sensor exploitation,

autonomy, and signal and image processing applications. He is the author of numerous conference and journal papers, including the article on synthetic aperture radar (SAR) in the *Encyclopedia of Optical Engineering* (Taylor & Francis, 2015).

Dr. Richards is a member of the Executive Committee of the Automatic Target Recognition Working Group, a consortium of developers, sponsors, and users of sensor exploitation systems. He is widely recognized as a leading figure in the field of automatic target recognition for SAR and high-range-resolution radar.



Shreyas Sundaram (Senior Member, IEEE) received the M.S. and Ph.D. degrees in electrical engineering from the University of Illinois at Urbana-Champaign, Champaign, IL, USA, in 2005 and 2009, respectively.

He is currently an Associate Professor with the School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN, USA. He was a Postdoctoral Researcher with the University of Pennsylvania, Philadelphia, PA, USA, from 2009 to 2010, and an Assistant

Professor with the Department of Electrical and Computer Engineering, University of Waterloo, Waterloo, ON, USA, from 2010 to 2014. His research interests include network science, analysis of large-scale dynamical systems, fault-tolerant and secure control, linear system and estimation theory, game theory, and the application of algebraic graph theory to system analysis.

Dr. Sundaram is a recipient of the NSF CAREER Award, and an Air Force Research Lab Summer Faculty Fellowship. At Purdue, he received the Hesselberth Award for Teaching Excellence and the Ruth and Joel Spira Outstanding Teacher Award. At Waterloo, he received the Department of Electrical and Computer Engineering Research Award and the Faculty of Engineering Distinguished Performance Award. He was also a recipient of the M. E. Van Valkenburg Graduate Research Award and the Robert T. Chien Memorial Award from the University of Illinois, and was a finalist for the Best Student Paper Award at the 2007 and 2008 American Control Conferences.