

Adaptive Social Learning

Virginia Bordignon¹, *Graduate Student Member, IEEE*, Vincenzo Matta², *Member, IEEE*,
and Ali H. Sayed³, *Fellow, IEEE*

Abstract—This work proposes a novel strategy for social learning by introducing the critical feature of *adaptation*. In social learning, several distributed agents update continually their belief about a phenomenon of interest through: *i*) direct observation of streaming data that they gather *locally*; and *ii*) diffusion of their beliefs through *local cooperation* with their neighbors. Traditional social learning implementations are known to learn well the underlying hypothesis (which means that the belief of every individual agent peaks at the true hypothesis), achieving steady improvement in the learning accuracy under stationary conditions. However, these algorithms do not perform well under nonstationary conditions commonly encountered in online learning, exhibiting a significant inertia to track drifts in the streaming data. In order to address this gap, we propose an Adaptive Social Learning (ASL) strategy, which relies on a small step-size parameter to tune the adaptation degree. First, we provide a detailed characterization of the *learning performance* by means of a steady-state analysis. Focusing on the small step-size regime, we establish that the ASL strategy achieves consistent learning under standard global identifiability assumptions. We derive reliable Gaussian approximations for the probability of error (i.e., of choosing a wrong hypothesis) at each individual agent. We carry out a large deviations analysis revealing the universal behavior of adaptive social learning: the error probabilities decrease exponentially fast with the inverse of the step-size, and we characterize the resulting exponential learning rate. Second, we characterize the *adaptation performance* by means of a detailed transient analysis, which allows us to obtain useful analytical formulas relating the *adaptation time* to the step-size. The revealed dependence of the adaptation time and the error probabilities on the step-size highlights the fundamental trade-off between adaptation and learning emerging in adaptive social learning.

Index Terms—Social learning, adaptation, diffusion strategies, large deviations.

I. INTRODUCTION AND MOTIVATION

SOcial learning is a collective process whereby some agents form their opinions about a phenomenon of interest through the local exchange of information [2]–[12].

Manuscript received April 4, 2020; revised February 9, 2021; accepted June 10, 2021. Date of publication July 5, 2021; date of current version August 25, 2021. This work was supported in part by the Swiss National Science Foundation (SNSF) under Grant 205121-184999. This article was presented in part at the European Signal Processing Conference (EUSIPCO). (Corresponding author: Virginia Bordignon.)

Virginia Bordignon and Ali H. Sayed are with the School of Engineering, École Polytechnique Fédérale de Lausanne EPFL, 1015 Lausanne, Switzerland (e-mail: virginia.bordignon@epfl.ch; ali.sayed@epfl.ch).

Vincenzo Matta is with the Department of Information and Electrical Engineering and Applied Mathematics (DIEM), University of Salerno, 84084 Fisciano (SA), Italy, and also with the National Inter-University Consortium for Telecommunications (CNIT), Italy (e-mail: vmatta@unisa.it).

Communicated by R. Venkataramanan, Associate Editor for Machine Learning.

Color versions of one or more figures in this article are available at <https://doi.org/10.1109/TIT.2021.3094633>.

Digital Object Identifier 10.1109/TIT.2021.3094633

More formally, given a set of H hypotheses, $\Theta = \{1, 2, \dots, H\}$, there is one true state of nature $\theta_0 \in \Theta$. Each agent k , at time i , collects streaming data $\xi_{k,i}$ (bold notation is used for random objects) drawn from a distribution that depends on the underlying hypothesis θ_0 . By exchanging local information with its neighbors, each agent assigns a belief $\mu_{k,i}(\theta)$ to each hypothesis $\theta \in \Theta$, with the belief vector $\boldsymbol{\mu}_{k,i} = [\mu_{k,i}(1), \mu_{k,i}(2), \dots, \mu_{k,i}(H)]^\top$ being a probability vector. Proper social learning occurs when the highest credibility is assigned to the true hypothesis, i.e., when the belief $\mu_{k,i}(\theta)$ is maximized at $\theta = \theta_0$.

Several social learning strategies have been proposed in the literature. As a common feature, all of them exhibit the desirable property that, as time goes to infinity, the belief function converges to 1 at θ_0 . In other words, if the amount of streaming data is sufficiently large, maximum credibility is assigned to the correct hypothesis whereas minimum (i.e., zero) credibility is assigned to the wrong hypotheses [13]–[20]. Moreover, for most social learning implementations, convergence to the true hypothesis is exponentially fast.

However, such remarkably good convergence properties have a subtle consequence that has been overlooked so far in the literature. This is because the exponentially increasing accuracy in learning the true hypothesis makes all agents stubborn! We illustrate this phenomenon through a simple example.

Consider a weather forecast problem solved by an online social learning algorithm. Assume that the agents are collecting data that drive them to believe that “*tomorrow will be sunny*.” After some time, however, assume the streaming dataset available for the decision evolves in response to changes in weather conditions with the most recent evidences suggesting markedly that “*tomorrow will be rainy*.” The traditional (existing) social learning algorithms discourage agents from changing their “mind” and it will be virtually impossible for the agents to adapt to the new situation and revise their earlier conclusion. This effect is clearly visible in the example of Fig. 1. In this example we considered a network of 10 agents that collect data originating from one of three possible hypotheses, namely, “*sunny*,” “*cloudy*,” “*rainy*.” The data are initially consistent with the hypothesis “*sunny*.” We observe from the blue curve in the top plot of Fig. 1 that the belief of agent 1 for the hypothesis “*sunny*” approaches the value one and, therefore, this agent is able to arrive at the correct determination about the state of nature. However, in our simulation, the state of nature is made to change to “*rainy*” at instant $i = 200$ (not shown in the figure). It is observed that the beliefs of agent 1 start changing only around $i = 350$ and the agent first transitions to believing that it




Fig. 1. *Traditional social learning strategy. Top panel.* Belief evolution of agent 1, with a state of nature drifting at time $i = 200$, from “sunny” to “rainy.” *Bottom panel.* The instantaneous decision taken by agent 1 by choosing the hypothesis that maximizes the current belief. We see that traditional social learning is not able to adapt to the new state of nature.

is “cloudy” (the green curve) before switching to believing that it is “rainy” many iterations later around $i = 550$. This example shows that, under traditional social learning schemes, agents are not able to recover sufficiently fast to adapt their beliefs and track changes in the state of nature. The outcome of the social learning algorithm (we display in the figure the belief of agent 1, with similar behavior being observed for other agents) shows clearly that the agents learn well until instant $i = 200$, since they give almost full credibility to the hypothesis according to which the data are drawn, but react far slower afterward when the state of nature changes. As a matter of fact, the traditional social learning algorithm has a delayed reaction to the change, only perceiving that something has changed at instant $i \approx 350$, but *still not detecting the true state*, because the agent gives maximum credibility to the wrong intermediate hypothesis “cloudy.” After a prohibitive number of iterations, at $i \approx 550$, agents manage to overcome their stubbornness and opt for the correct hypothesis “rainy.”

This behavior can be problematic for an online algorithm continuously fed by streaming data since, in many practical scenarios, the system operating conditions (e.g., the underlying state of nature as in the introductory example, or the network topology, the quality of data, the statistical models,...) are reasonably expected to undergo some changes over time. For this reason, a good learning algorithm must be able to *adapt* to drifts in the streaming information collected by the agents. This work proposes an Adaptive Social Learning (ASL) strategy to fill this gap. One instance of such strategy is shown in Fig. 2 with reference to the same example from Fig. 1. We see that the ASL algorithm reacts much faster (almost instantly) and is able to track the target change at instant $i \approx 200$, exhibiting an *adaptation capacity that is remarkably higher than that of the classic social learning algorithm*.

There are at least two advantages in devising the ASL algorithm. The first one is related to a modeling perspective.

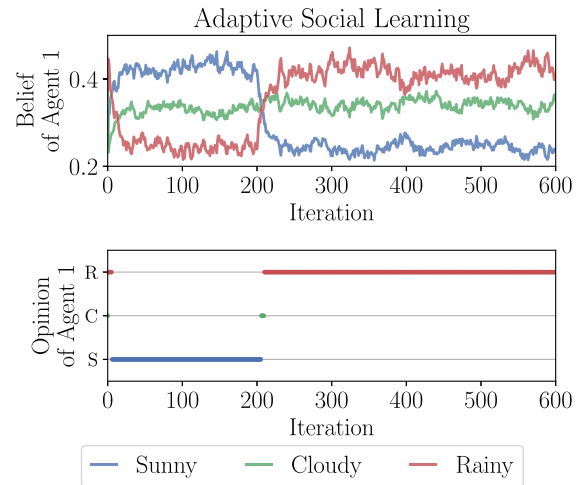


Fig. 2. *Adaptive social learning strategy proposed in this work. Top panel.* Belief evolution of agent 1, with a state of nature drifting at time $i = 200$, from “sunny” to “rainy.” *Bottom panel.* The instantaneous decision taken by agent 1 by choosing the hypothesis that maximizes the current belief. We see that the proposed strategy is able to adapt to the new state of nature.

As already indicated, the existing social learning strategies are not able to endow agents with adaptation abilities whereas the proposed ASL model will be able to do so. The second implication is related to a designing perspective. Social learning algorithms are useful not only in modeling opinion formation over social networks. They are also useful in designing man-made engineered systems (such as robotic swarms) tasked to solve decision problems collectively. Endowing such systems with adaptation abilities is critical for many applications.

The main contributions of this work are as follows. First, we introduce a novel social learning strategy that enables adaptation. Then, we provide an accurate analytical characterization of this strategy. In particular, by exploiting recent advances in the field of distributed detection over adaptive networks — see [21] for an overview — we furnish a detailed characterization of the social learning performance at each individual agent, in terms of i) convergence of the system at the steady-state (Theorem 1); ii) achievability of consistent learning (Theorem 2); iii) a Gaussian approximation for the learning performance (Theorem 3); iv) the error exponents for the learning error probabilities (Theorem 4); and v) the transient evolution for the instantaneous error probabilities. As the analysis will show, the ASL model allows the user to design the adaptation time, at the expense of losing learning accuracy, i.e., agents no longer achieve full confidence around the true hypothesis. Instead, agents maintain some skepticism regarding the true hypothesis, as illustrated in the belief curves of Fig. 2.

II. BACKGROUND AND PROBLEM FORMULATION

The agents of the network collect streaming observations (or data) about a phenomenon of interest. Agent $k = 1, 2, \dots, N$, at time epoch $i = 1, 2, \dots$, collects a “private” observation $\xi_{k,i}$ belonging to a certain space \mathcal{X}_k . The qualification “private” comes from the assumption that the raw observations cannot

be shared among agents in order to, for example, minimize communication costs or preserve secrecy. The dependence of the space \mathcal{X}_k upon k allows for a possible heterogeneity in the types of data at the different agents. The data will be assumed statistically independent over time, i.e., over the index i , whereas they can be dependent across agents.¹

In social learning, since the inter-agent dependence is usually not known to the agents, the focus is on *marginal* distributions, i.e., on the distribution pertaining to any *individual* agent. Specifically, it is assumed that the distribution of $\xi_{k,i}$ belongs to a set of H admissible models that are identified by a discrete parameter (or hypothesis) $\theta \in \Theta = \{1, 2, \dots, H\}$. The likelihood of agent k evaluated at θ is denoted by:

$$L_k(\xi|\theta), \quad \xi \in \mathcal{X}_k. \quad (1)$$

The presence of subscript k highlights that the likelihoods are allowed to vary across the agents. In our treatment, $L_k(\xi|\theta)$ (regarded as a function of ξ) can be either a probability density or mass function, depending on whether $\xi_{k,i}$ is continuous or discrete, respectively. Moreover, in order to avoid trivialities we assume the following regularity condition on Kullback-Leibler (KL) divergences [22].

Assumption 1 (Finiteness of KL Divergences): For each $k = 1, 2, \dots, N$ and each pair of distinct hypotheses θ and θ' , the Kullback-Leibler divergence between $L_k(\xi|\theta)$ and $L_k(\xi|\theta')$ is finite. \square

In social learning implementations, the two main objects of the learning process are: an intermediate belief $\psi_{k,i}(\theta)$, which each agent k shares at time i with its neighbors; and the belief $\mu_{k,i}(\theta)$, which agent k obtains at time i by combining the intermediate beliefs received from its neighbors. For the algorithm initialization, we assume the following standard condition.

Assumption 2 (Positive Initial Beliefs): All agents start with a strictly positive belief for all hypotheses, i.e., $\mu_{k,0}(\theta) > 0$ for each agent k and all $\theta \in \Theta$. \square

In order to capture the essence of our *adaptive* social learning strategy, it is useful to introduce first some background on *traditional* social learning. We refer in particular to the social learning strategy presented in [17]–[20], which is a two-step algorithm that iterates over time as follows.

In the first step, each agent k constructs an *intermediate* belief vector $\psi_{k,i}$ by incorporating the current observation $\xi_{k,i}$ into the belief of the preceding time epoch, $\mu_{k,i-1}$, through the following *Bayesian update*:

$$\psi_{k,i}(\theta) = \frac{\mu_{k,i-1}(\theta)L_k(\xi_{k,i}|\theta)}{\sum_{\theta' \in \Theta} \mu_{k,i-1}(\theta')L_k(\xi_{k,i}|\theta')}, \quad (2)$$

where the denominator is a normalization factor that makes $\psi_{k,i}$ a probability vector.

In the second step, each agent k aggregates into its own current belief $\mu_{k,i}$ the intermediate beliefs received from its neighbors by combining linearly the *logarithm* of the received intermediate beliefs, and then using exponentiation and normalization to get back an admissible probability vector.

¹Some of the forthcoming results (Theorems 3 and 4) will be proved under the additional assumption of independence across agents.

Specifically, each agent k at time i applies the following combination rule:

$$\mu_{k,i}(\theta) = \frac{\exp \left\{ \sum_{\ell \in \mathcal{N}_k} a_{\ell k} \log \psi_{\ell,i}(\theta) \right\}}{\sum_{\theta' \in \Theta} \exp \left\{ \sum_{\ell \in \mathcal{N}_k} a_{\ell k} \log \psi_{\ell,i}(\theta') \right\}}, \quad (3)$$

using a collection of convex combination weights:

$$0 < a_{\ell k} < 1, \quad \sum_{\ell=1}^N a_{\ell k} = 1, \quad a_{\ell k} = 0 \text{ for } \ell \notin \mathcal{N}_k, \quad (4)$$

where \mathcal{N}_k denotes the neighborhood of agent k , with k itself being included. The combination weights can be conveniently arranged into the *nonnegative and left-stochastic* combination matrix $A = [a_{\ell k}]$.

In the forthcoming treatment, we assume that the network is strongly connected (i.e., for any two nodes ℓ and k , there exists always a path with nonzero weights linking them in both directions, and at least one node in the network has a self-loop, i.e., $a_{kk} > 0$ for at least one agent k) [23]. Under these assumptions, the (nonnegative and left-stochastic) matrix A is *primitive*, implying, in view of the Perron-Frobenius theorem, that the Perron eigenvector π associated with the matrix A has all strictly positive entries [23][Lemma F.4, p. 775]:

$$A\pi = \pi, \quad \sum_{\ell=1}^N \pi_{\ell} = 1, \quad \pi_{\ell} > 0 \text{ for all } \ell = 1, 2, \dots, N, \quad (5)$$

and that the columns of the matrix powers A^m converge, as $m \rightarrow \infty$, to the Perron eigenvector at an exponential rate governed by the second largest-magnitude eigenvalue of A , as stated in the following property [24][Th. 8.5.1, p. 516].

Property 1 (Convergence of Matrix Powers): Let β_2 be the second largest-magnitude eigenvalue of A . Then, for any positive β such that $|\beta_2| < \beta < 1$, there exists a positive constant κ (depending only on A and β), such that, for all $\ell, k = 1, 2, \dots, N$, and for all $m = 1, 2, \dots$, we have that:

$$\left| [A^m]_{\ell k} - \pi_{\ell} \right| \leq \kappa \beta^m. \quad (6)$$

\square

In traditional social learning, a stationary setting is assumed where the data collected by the agents are generated from one particular model (the *true* hypothesis) and the goal of social learning is to let the agents learn this hypothesis from the data. It has been shown that, under the aforementioned assumptions, the algorithm described by (2)–(3) leads each agent to learn the true hypothesis almost surely as $i \rightarrow \infty$ [17]–[20].

In the adaptive context, the statistical conditions governing the data can change over time. While effective learning must still be guaranteed under stationary conditions, it is also critical to guarantee that the learning algorithms are able to react fast to drifting conditions. Under these changing environments, traditional social learning algorithms do not perform well. The fundamental goal is therefore to devise a social learning algorithm that allows agents to promptly react to these drifts and start learning the “new” model.

A. Adaptive Social Learning

An *adaptive* algorithm can be broadly defined as one that is tailored for environments in which the characteristics of the collected data might drift over time. The adaptive algorithm should be able to react promptly in view of these changes and deliver proper inference performance in a reasonable reaction time.

While certainly intuitive, the above description of adaptation is only qualitative. In this work, we will present a rigorous analysis of the adaptation/learning trade-off in social learning. To this aim, it is necessary to identify formally the concepts of adaptation and learning, and the technical framework that will be used to characterize these concepts.

- *Learning*. In the context of social learning, “learning” means “guessing the right hypothesis.” In order to quantify the learning performance, we specialize the standard prescriptions of adaptation theory to the social learning context. Given that the data are steadily generated according to a certain true likelihood model, what is the probability that an agent guesses the true state of nature? In the theory of adaptation, this analysis is commonly referred to as *steady-state* analysis [23].

- *Adaptation*. Assume that the system has been in operation for an arbitrary time. During this time, several phenomena can have occurred, i.e., variations of the true hypothesis, variations in the statistical conditions (i.e., malfunctioning of the system giving rise to distributions different from the nominal ones), missing observations, and so on. Due to the recursive nature of the social learning algorithms, at a given time i_0 all these variations are simply summarized in a certain initial belief vector μ_{i_0} . From $i_0 + 1$ onward, assume that the system becomes stable and the data are steadily generated according to a given likelihood model. Accordingly, the adaptation ability will be quantified by measuring how long it takes (*adaptation time*), given an arbitrary initial belief μ_{i_0} , for an agent to enter the steady-state regime and reach a prescribed probability of guessing the true hypothesis. In the theory of adaptation, this analysis is commonly referred to as *transient* analysis [23].

III. ASL STRATEGY

Examining (2), we see that the classic Bayesian update incorporates the new information into the past belief by giving *equal* weight to both $\mu_{k,i-1}$ and the likelihood of the new data $L_k(\xi_{k,i}|\theta)$. In order to promote adaptation, it is necessary to increase the relative credit given to the new data with respect to the belief accumulated over time by learning from past data. To this end, we turn the update step in (2) into the following *adaptive* form:

$$\psi_{k,i}(\theta) = \frac{\mu_{k,i-1}^{1-\delta}(\theta) L_k^\delta(\xi_{k,i}|\theta)}{\sum_{\theta' \in \Theta} \mu_{k,i-1}^{1-\delta}(\theta') L_k^\delta(\xi_{k,i}|\theta')}, \quad (7)$$

where $0 < \delta < 1$ is a design parameter employed by each agent to modulate the relative weights assigned to the past and new information. In particular, relatively large values for δ give more importance to the new data, whereas small values for δ give more importance to the past beliefs. In this way, as we

will show later in Sec. VII, the step-size parameter δ infuses the social learning algorithm with an adaptation mechanism.

A. ASL–Stochastic Gradient Interpretation

While the role of the parameter δ in favoring adaptation might be intuitive, the rationale behind the particular update rule in (7) might not be. In order to clarify this important issue, let us focus on the logarithmic belief ratio between two hypotheses θ and θ' :

$$\log \frac{\mu_{k,i}(\theta)}{\mu_{k,i}(\theta')}. \quad (8)$$

This quantity is the essential building block in the process of opinion formation, since an agent would opt for the hypothesis that maximizes the belief. In other words, agent k at time i would opt for opinion θ if its log-belief ratio $\log \frac{\mu_{k,i}(\theta)}{\mu_{k,i}(\theta')}$ is positive for all $\theta' \neq \theta$. Combining (7) with (3) and developing the associated recursion we can obtain the time-evolution of the log-belief ratio as:

$$\begin{aligned} \log \frac{\mu_{k,i}(\theta)}{\mu_{k,i}(\theta')} &= (1 - \delta) \sum_{\ell \in \mathcal{N}_k} a_{\ell k} \log \frac{\mu_{\ell,i-1}(\theta)}{\mu_{\ell,i-1}(\theta')} \\ &+ \delta \sum_{\ell \in \mathcal{N}_k} a_{\ell k} \log \frac{L_\ell(\xi_{\ell,i}|\theta)}{L_\ell(\xi_{\ell,i}|\theta')} \\ &= (1 - \delta)^i \sum_{\ell=1}^N [A^i]_{\ell k} \log \frac{\mu_{\ell,0}(\theta)}{\mu_{\ell,0}(\theta')} \\ &+ \delta \sum_{m=0}^{i-1} \sum_{\ell=1}^N (1 - \delta)^m [A^{m+1}]_{\ell k} \log \frac{L_\ell(\xi_{\ell,i-m}|\theta)}{L_\ell(\xi_{\ell,i-m}|\theta')}. \quad (9) \end{aligned}$$

It is straightforward to prove that the time-evolution of the log-belief ratio in (9) has the form of a *distributed stochastic gradient* algorithm with step-size δ and with quadratic cost function — see [21], [23]. In recent studies, it has been shown how this form of distributed adaptive inference can be useful for distributed binary hypothesis testing [21], [27], [30]. These studies motivate our choice for (7), whose goodness for social learning will be rigorously established in the forthcoming analysis.

B. ASL–Bayesian Update Interpretation

Examining (7), we see that the ASL strategy implements a convex combination of probability functions at the exponent, by discounting both the past belief and the new likelihood through the weights $1 - \delta$ and δ , respectively. However, the update (7) cannot be considered a *Bayesian* update because the likelihood exponentiated to δ does not integrate to one (w.r.t. ξ). Notably, the same problem does not occur for the discounted belief, since it can be normalized, both at the numerator and denominator of (7), by dividing by $\sum_{\theta' \in \Theta} \mu_{k,i-1}^{1-\delta}(\theta')$. Exploiting the latter property, we now show how it is possible to modify (7) to get an *adaptive Bayesian* update. First, agent k takes the past belief $\mu_{k,i-1}(\theta)$ and builds a new belief as:

$$\tilde{\mu}_{k,i-1}(\theta) = \frac{\mu_{k,i-1}^{1-\delta}(\theta)}{\sum_{\theta' \in \Theta} \mu_{k,i-1}^{1-\delta}(\theta')}. \quad (10)$$

Second, agent k implements the Bayesian update rule in (2) by replacing $\mu_{k,i-1}(\theta)$ with $\tilde{\mu}_{k,i-1}(\theta)$, which yields:

$$\psi_{k,i}(\theta) = \frac{\tilde{\mu}_{k,i-1}(\theta)L_k(\xi_{k,i}|\theta)}{\sum_{\theta' \in \Theta} \tilde{\mu}_{k,i-1}(\theta')L_k(\xi_{k,i}|\theta')}. \quad (11)$$

Notably, the exponentiation and normalization in (10) has the physical meaning of *flattening* the belief vector, i.e., of making it more uniform across θ . In this way, if an agent had a particularly peaked belief around a certain hypothesis, perhaps due to a bias accumulated over time, flattening the belief helps to give more credit to new data. Observe from (11) that the limiting choice $\delta = 0$ (i.e., no adaptation) gives back the classic Bayesian update in (2). In contrast, the update in (7) cannot be reduced to (2) for any selection of $\delta \in (0, 1)$. This notwithstanding, we now show that the ASL strategies (7) and (11) are in fact equivalent. To this aim, we can develop the recursion obtained by combining (3), (10), and (11) and get:

$$\begin{aligned} \log \frac{\mu_{k,i}(\theta)}{\mu_{k,i}(\theta')} &= (1 - \delta) \sum_{\ell \in \mathcal{N}_k} a_{\ell k} \log \frac{\mu_{\ell,i-1}(\theta)}{\mu_{\ell,i-1}(\theta')} \\ &+ \sum_{\ell \in \mathcal{N}_k} a_{\ell k} \log \frac{L_{\ell}(\xi_{\ell,i}|\theta)}{L_{\ell}(\xi_{\ell,i}|\theta')} \\ &= (1 - \delta)^i \sum_{\ell=1}^N [A^i]_{\ell k} \log \frac{\mu_{\ell,0}(\theta)}{\mu_{\ell,0}(\theta')} \\ &+ \sum_{m=0}^{i-1} \sum_{\ell=1}^N (1 - \delta)^m [A^{m+1}]_{\ell k} \log \frac{L_{\ell}(\xi_{\ell,i-m}|\theta)}{L_{\ell}(\xi_{\ell,i-m}|\theta')}. \quad (12) \end{aligned}$$

Let us now compare (9) against (12). We see that in both equations there is a term that dies out exponentially fast with time, and which is due to the initialization term $\log \frac{\mu_{\ell,0}(\theta)}{\mu_{\ell,0}(\theta')}$. We remark that this initialization term is determined by the beliefs set by the agents in the *absence of data*. Thus, it is zero if the agents set a uniform, non-informative prior, or it can be non-uniform if the agents have unbalanced prior convictions. The relevant term that determines the algorithms' evolution over time is given by the double summations appearing in (9) and (12). Comparing these summations, we see that they differ only by a scaling factor δ . As a result, we conclude that the time-evolution of the log-belief ratios for the two ASL strategies is equivalent. For example, the opinion that maximizes the belief function would be the same under both strategies, implying the same error probability. In fact, proportionality of the log-belief ratios implies that the belief function of one strategy is simply an exponentiated (and normalized) version of the belief function of the other strategy. This does not mean that the beliefs of the two strategies would take on the same values. In particular, our results will show that, as $\delta \rightarrow 0$, the steady-state log-belief ratios are stable under (7), which immediately implies that they diverge (i.e., achieving a belief close to 1 at the true hypothesis) under (11). While immaterial from a technical perspective, these differences might matter from a *behavioral* perspective [6], namely, to understand which update strategy reflects better the way of reasoning that an individual agent uses in social learning environments. For the sake of clarity, in the presentation of our technical results we opt for sticking to the update rule in (7), since it

automatically stabilizes the log-belief ratio without necessity of additional scaling factors.

The adaptation properties of the ASL strategy are enabled by a learning mechanism that is fundamentally different from that of classic social learning. To see why, let us assume that the true hypothesis remains stable for a sufficiently long time interval. Different from what happens in classic social learning — e.g., in (2) — in the ASL strategy the belief will *not* converge as time i increases. In contrast, the belief will *vary* indefinitely, preserving a random behavior also in the steady state. The learning performance will then be assessed by examining the statistical behavior of the beliefs in steady state. We will provide an accurate characterization of such statistical behavior in the regime of small step-sizes, i.e., by performing an asymptotic analysis as $\delta \rightarrow 0$. Under this regime, we will show that the probability of guessing the right hypothesis approaches 1 for sufficiently small step-sizes. We will furthermore characterize the transient performance by obtaining closed-form relationships that reveal how the adaptation time grows with smaller step-sizes. The overall analysis will highlight well the adaptation/learning trade-off: small (resp., large) values of δ mean less (resp., more) adaptation and higher (resp., lower) learning accuracy.

IV. STATISTICAL DESCRIPTORS OF THE LEARNING PERFORMANCE

Assume that the algorithm has been running until a certain time i_0 , with the evolution of the system up to i_0 being summarized in the “initial” belief vectors μ_{k,i_0} . Starting from i_0 , the ASL algorithm behavior will exhibit two important phases: a *transient* phase where, given the (possibly wrong) initial belief, each agent must suddenly adapt in order to depart from μ_{k,i_0} and start learning the correct hypothesis; and a *steady-state* phase where, given sufficient time to learn ($i \rightarrow \infty$), each agent must achieve high confidence in learning the correct hypothesis. According to the theory of adaptive inference, the performance of an adaptive learning strategy is characterized under the *steady-state* regime.

The following property is relevant for steady-state analysis. By examining the algorithm recursions (3)–(7), it is straightforward to see that, in light of Assumptions 1 and 2, the belief remains always nonzero at any θ during the algorithm evolution.² Now, assume that the algorithm has been running up to time i_0 , and that from $i_0 + 1$ onward the system remains stationary for sufficiently long time, with the data being generated according to hypothesis θ_0 . In order to perform a steady-state analysis from $i_0 + 1$ onward, we need to consider μ_{k,i_0} as initial state. Since we have observed that the beliefs are always nonzero, we can see that the initial belief vector μ_{k,i_0} fulfills Assumption 2.

In summary, for the purpose of the steady-state analysis and without loss of generality, we will assume that the steady-state

²This property follows by induction once we observe that, starting from a belief that is nonzero at any θ : *i*) the intermediate belief remains nonzero at any θ because the likelihoods in the update step (7) cannot be zero (but for an ensemble of zero probability) otherwise Assumption 1 would be violated; and *ii*) the final belief in (3) remains nonzero at any θ since the combination weights are convex.

analysis starts at time $i_0 = 0$ and consider an initial belief vector $\boldsymbol{\mu}_{k,0}$ that fulfills Assumption 2. The true hypothesis θ_0 is kept constant over time, yielding:

$$\boldsymbol{\xi}_{k,i} \sim L_k(\xi|\theta_0), \quad k = 1, 2, \dots, N, \quad i = 1, 2, \dots \quad (13)$$

Therefore, for the purpose of the steady-state analysis, we will always imply that expectations and probabilities are evaluated under the distributions $L_k(\xi|\theta_0)$. Note also that, under the steady-state regime, the data $\{\boldsymbol{\xi}_{k,i}\}$ are independent and identically distributed (i.i.d.) over time, i.e., over the index i . We will assume that they can have different distributions across the agents, i.e., across the index k . Statistical independence across the agents will be only used to prove some of the forthcoming results (Theorems 3 and 4 further ahead).

A. Log-Belief Ratios and Error Probabilities

In order to characterize the learning performance, it is convenient to introduce the logarithm of the ratio between the belief evaluated at θ_0 and the belief evaluated at a generic hypothesis $\theta \neq \theta_0$:

$$\boldsymbol{\lambda}_{k,i}^{(\delta)}(\theta) \triangleq \log \frac{\boldsymbol{\mu}_{k,i}(\theta_0)}{\boldsymbol{\mu}_{k,i}(\theta)}, \quad (14)$$

which is well-defined since, as already remarked, the belief remains nonzero at any θ during the algorithm evolution. Before continuing, it is important to make a notational remark. With the symbol $\boldsymbol{\lambda}_{k,i}^{(\delta)}(\theta)$ we denote a *random* (bold notation) function of: the agent index $k = 1, 2, \dots, N$, the time index $i = 0, 1, \dots$, the hypothesis $\theta \in \Theta \setminus \theta_0$, and the adaptation parameter δ . When we omit the argument θ and write $\boldsymbol{\lambda}_{k,i}^{(\delta)}$, we will be referring to the $(H-1) \times 1$ vector of log-belief ratios, namely,

$$\boldsymbol{\lambda}_{k,i}^{(\delta)} = \left[\boldsymbol{\lambda}_{k,i}^{(\delta)}(\theta_1), \boldsymbol{\lambda}_{k,i}^{(\delta)}(\theta_2), \dots, \boldsymbol{\lambda}_{k,i}^{(\delta)}(\theta_{H-1}) \right]^\top, \quad (15)$$

where the elements in the set of wrong-hypotheses have been indexed as:

$$\Theta \setminus \theta_0 = \{\theta_1, \theta_2, \dots, \theta_{H-1}\}. \quad (16)$$

One natural way for the agents to choose a hypothesis is to select the hypothesis that maximizes the belief. Therefore, the error probability at each time i can be expressed as

$$p_{k,i}^{(\delta)} = \mathbb{P} \left[\arg \max_{\theta \in \Theta} \boldsymbol{\mu}_{k,i}(\theta) \neq \theta_0 \right]. \quad (17)$$

It is useful to rewrite the error probability as a function of the log-belief ratios. To this end, observe that the event within brackets in (17) corresponds to saying that the belief is not maximized at θ_0 , which in turn corresponds to saying that the log-belief ratios in (14) are less than or equal to zero for at least one $\theta \neq \theta_0$. Therefore, the *instantaneous* error probability can be equivalently rewritten as:

$$p_{k,i}^{(\delta)} = \mathbb{P} \left[\exists \theta \neq \theta_0 : \boldsymbol{\lambda}_{k,i}^{(\delta)}(\theta) \leq 0 \right]. \quad (18)$$

Finally, we introduce the *steady-state* error probability:

$$p_k^{(\delta)} \triangleq \lim_{i \rightarrow \infty} p_{k,i}^{(\delta)}. \quad (19)$$




Fig. 3. Evolution of the error probability of two agents in a network running the ASL algorithm.

There are two fundamental questions related to the concept of steady-state error probability. The first question regards its *existence*, which is in principle not guaranteed. Theorem 1 will provide an affirmative answer to this question by characterizing the steady-state behavior of the log-belief ratios. The second question regards the *evaluation* of the steady-state error probability. An exact evaluation is generally a formidable task. Therefore, to tackle this critical problem, we will perform an asymptotic analysis in the regime of small δ , which will allow us to obtain reliable predictions of the steady-state performance.

In Fig. 3 we show an example of evolution for the error probability of two agents in a network implementing the ASL strategy.³ All the probabilities are estimated empirically by Monte Carlo simulation. We see how the instantaneous error probability $p_{k,i}^{(\delta)}$ converges to a steady-state *nonzero* value $p_k^{(\delta)}$ as i increases. It is useful to remark that this behavior is different from that of classic social learning, where, *under stationary conditions*, the error probability of each agent vanishes as time elapses. This is one instance of the adaptation/learning trade-off: non-adaptive strategies can increase their accuracy indefinitely under stationary conditions. However, astronomically low values of the error probabilities lead to a detrimental inertia in responding to possible changes.

B. Log-Likelihood Ratios

For $k = 1, 2, \dots, N$, $i = 0, 1, \dots$, and $\theta \neq \theta_0$, we introduce the log-likelihood ratio:

$$\mathbf{x}_{k,i}(\theta) \triangleq \log \frac{L_k(\boldsymbol{\xi}_{k,i}|\theta_0)}{L_k(\boldsymbol{\xi}_{k,i}|\theta)}, \quad (20)$$

and its expectation:

$$d_k(\theta) \triangleq \mathbb{E}[\mathbf{x}_{k,i}(\theta)] < \infty, \quad (21)$$

namely, the KL divergence between $L_k(\xi|\theta_0)$ and $L_k(\xi|\theta)$, which is finite in view of Assumption 1, implying that the log-likelihood ratios cannot diverge (but for an ensemble of realizations with zero probability). We recall that the expectation in (21) is computed assuming that the random variable $\boldsymbol{\xi}_{k,i}$ is distributed according to model $L_k(\xi|\theta_0)$. Since we focus on the steady state, this distribution is constant

³The details of the network topology as well as of the statistical learning problem are immaterial at this stage of the presentation.

over time, which explains why $d_k(\theta)$ does not depend on i . Furthermore, since the true hypothesis θ_0 is held fixed during the steady-state analysis, in order to avoid a heavier notation we are not emphasizing the dependence of the KL divergence $d_k(\theta)$ on θ_0 .

We continue by introducing an average variable that will play a role in the forthcoming results, namely, the *network average* of log-likelihood ratios, for all $\theta \neq \theta_0$:

$$\mathbf{x}_{\text{ave},i}(\theta) = \sum_{\ell=1}^N \pi_\ell \mathbf{x}_{\ell,i}(\theta). \quad (22)$$

The random variable $\mathbf{x}_{\text{ave},i}(\theta)$ appearing in (22) is obtained by combining linearly the local log-likelihood ratios $\mathbf{x}_{\ell,i}(\theta)$. The combination weight assigned to the log-likelihood ratio of the ℓ -th agent is given by the *limiting* combination weight, i.e., by the ℓ -th entry, π_ℓ , of the Perron eigenvector. We will see in the following that the asymptotic properties of the ASL strategy as $\delta \rightarrow 0$ are directly related to the statistical properties of the vector of average variables, $\mathbf{x}_{\text{ave},i}$.

V. STEADY-STATE ANALYSIS

As we have remarked in the introduction, different from the classic social learning setting, in the *adaptive* setting the belief will not converge as $i \rightarrow \infty$. In contrast, the belief of each agent will preserve a *random* behavior. This everlasting randomness is critical to ensure that the algorithm will adapt quickly to a change in the environment. On the other hand, it makes the steady-state analysis more difficult, since the beliefs preserve a random character even when $i \rightarrow \infty$. In order to carry out a meaningful steady-state analysis, the fundamental preliminary step becomes then to establish whether such random fluctuations lead to *stable* random variables as $i \rightarrow \infty$. Theorem 1 further ahead ascertains that this is the case.

Before stating the theorem, let us examine the evolution of the log-belief ratios.

Exploiting (3) and (7), we end up with the following recursion, for every $\theta \neq \theta_0$:

$$\lambda_{k,i}^{(\delta)}(\theta) = \sum_{\ell \in \mathcal{N}_k} a_{\ell k} \left\{ (1 - \delta) \lambda_{\ell,i-1}^{(\delta)}(\theta) + \delta \mathbf{x}_{\ell,i}(\theta) \right\}, \quad (23)$$

which can be rewritten as the following two-step recursion:

$$\nu_{\ell,i}^{(\delta)}(\theta) = (1 - \delta) \lambda_{\ell,i-1}^{(\delta)}(\theta) + \delta \mathbf{x}_{\ell,i}(\theta), \quad (24)$$

$$\lambda_{k,i}^{(\delta)}(\theta) = \sum_{\ell \in \mathcal{N}_k} a_{\ell k} \nu_{\ell,i}^{(\delta)}(\theta). \quad (25)$$

The time-evolution of the log-belief ratios in (24) and (25) is in the form of a *diffusion* algorithm with constant *step-size* δ — see, e.g. [23]. This is why we referred to δ as the step-size.

Developing the recursion in (23) and recalling that $A = [a_{\ell k}]$ is the combination matrix we can write, for all $\theta \neq \theta_0$:

$$\begin{aligned} \lambda_{k,i}^{(\delta)}(\theta) &= (1 - \delta)^i \underbrace{\sum_{\ell=1}^N [A^i]_{\ell k} \lambda_{\ell,0}(\theta)}_{\text{transient term}} \\ &+ \delta \sum_{m=0}^{i-1} \sum_{\ell=1}^N (1 - \delta)^m [A^{m+1}]_{\ell k} \mathbf{x}_{\ell,i-m}(\theta). \end{aligned} \quad (26)$$

Since the transient term dies out as $i \rightarrow \infty$, in order to evaluate the steady-state behavior of $\lambda_{k,i}(\theta)$, we can ignore it and focus on the second term:

$$\widehat{\lambda}_{k,i}^{(\delta)}(\theta) = \delta \sum_{\ell=1}^N \sum_{m=0}^{i-1} (1 - \delta)^m [A^{m+1}]_{\ell k} \mathbf{x}_{\ell,i-m}(\theta). \quad (27)$$

A. Steady-State Log-Belief Ratios

The goal of the steady-state analysis is to evaluate the performance (i.e., the error probability) for large i . For this evaluation to be meaningful, we must ascertain that the error probability in (18) converges as $i \rightarrow \infty$. To this end, we will now establish that there exists a certain limiting random vector, $\widetilde{\lambda}_k^{(\delta)}$, such that the *probability distribution* of the vector of log-belief ratios, $\widehat{\lambda}_{k,i}^{(\delta)}$, converges, as $i \rightarrow \infty$, to the probability distribution of $\widetilde{\lambda}_k^{(\delta)}$. This notion of convergence can be formally defined as follows.

We say that the sequence (over the index i) of random vectors $\widehat{\lambda}_{k,i}^{(\delta)}$ converges *in distribution* or *weakly* as $i \rightarrow \infty$ if we can define a random vector $\widetilde{\lambda}_k^{(\delta)}$ such that [26]:

$$\lim_{i \rightarrow \infty} \mathbb{P} \left[\widehat{\lambda}_{k,i}^{(\delta)} \in \mathcal{B} \right] = \mathbb{P} \left[\widetilde{\lambda}_k^{(\delta)} \in \mathcal{B} \right] \quad (28)$$

for all measurable sets \mathcal{B} whose boundary $\partial \mathcal{B}$ has zero probability under the limiting distribution, namely, for all measurable sets \mathcal{B} fulfilling the condition:

$$\mathbb{P} \left[\widetilde{\lambda}_k^{(\delta)} \in \partial \mathcal{B} \right] = 0. \quad (29)$$

In the following, weak convergence will be compactly denoted as:

$$\widehat{\lambda}_{k,i}^{(\delta)} \xrightarrow{i \rightarrow \infty} \widetilde{\lambda}_k^{(\delta)}, \quad (30)$$

and the vector $\widetilde{\lambda}_k^{(\delta)}$ will be referred to as the *steady-state* log-belief vector, since it provides the statistical characterization of the log-belief vector $\widehat{\lambda}_{k,i}^{(\delta)}$ as $i \rightarrow \infty$.

We are now ready to present the theorem that establishes the existence of steady-state log-belief ratios.

Theorem 1 (Steady-State Log-Belief Ratios): Let Assumptions 1 and 2 hold, and let

$$\widetilde{\lambda}_{k,i}^{(\delta)}(\theta) \triangleq \delta \sum_{\ell=1}^N \sum_{m=0}^{i-1} (1 - \delta)^m [A^{m+1}]_{\ell k} \mathbf{x}_{\ell,m+1}(\theta) \quad (31)$$

be the random sum obtained from (27) by taking the summands in reversed order.

First, we have that all the N inner sums in (31) are almost-surely absolutely convergent as $i \rightarrow \infty$, implying that $\widetilde{\lambda}_{k,i}^{(\delta)}(\theta)$ converges almost surely to the random series:

$$\widetilde{\lambda}_k^{(\delta)}(\theta) \triangleq \delta \sum_{\ell=1}^N \sum_{m=0}^{\infty} (1 - \delta)^m [A^{m+1}]_{\ell k} \mathbf{x}_{\ell,m+1}(\theta). \quad (32)$$

Second, we have that the vector of log-belief ratios $\widehat{\lambda}_{k,i}^{(\delta)}$ (with the original, i.e., non-reversed ordering of summation) converges in distribution to the vector $\widetilde{\lambda}_k^{(\delta)}$, namely,

$$\widehat{\lambda}_{k,i}^{(\delta)} \xrightarrow{i \rightarrow \infty} \widetilde{\lambda}_k^{(\delta)}. \quad (33)$$

Proof: See Appendix B. ■




Fig. 4. Comparison of the random sequences $\hat{\lambda}_{k,i}^{(\delta)}(\theta)$ and $\tilde{\lambda}_{k,i}^{(\delta)}(\theta)$ for $\delta = 0.1$, for the Gaussian setting described in Sec. VIII further ahead.

It is useful to make some comments on Theorem 1. First, finiteness of the expectation of $\mathbf{x}_{k,i}$ is sufficient (through Assumption 1) to guarantee the existence of a steady-state random variable. No assumption is made on higher-order moments.

Second, it is important to notice that (32) does not correspond to letting $i \rightarrow \infty$ in the summation in (27). In order to explain why, let us compare the random sums:

$$\hat{\lambda}_{k,i}^{(\delta)}(\theta) = \delta \sum_{m=0}^{i-1} \sum_{\ell=1}^N (1-\delta)^m [A^{m+1}]_{\ell k} \mathbf{x}_{\ell,i-m}(\theta), \quad (34)$$

and

$$\tilde{\lambda}_{k,i}^{(\delta)}(\theta) = \delta \sum_{m=0}^{i-1} \sum_{\ell=1}^N (1-\delta)^m [A^{m+1}]_{\ell k} \mathbf{x}_{\ell,m+1}(\theta). \quad (35)$$

In Fig. 4 we examine a sample path for these sums, and we can see that they exhibit different behavior. The random sum in (34), displayed with solid line in Fig. 4, exhibits steadily random fluctuations as time elapses. In contrast, the random sum in (35), displayed with dashed line, converges as time elapses, specifically to the random value $\tilde{\lambda}_k^{(\delta)}(\theta)$ defined in (32). Both behaviors are consistent with what we have already shown in Theorem 1. These profoundly different behaviors depend on the different ordering of the summands in (34) and (35). In particular, in (35) the most recent term, $\mathbf{x}_{\ell,i}(\theta)$, takes the smallest weight $(1-\delta)^{i-1}$, which lets the remainder of the series vanish (almost surely). In contrast, in (34) the most recent term, $\mathbf{x}_{\ell,i}(\theta)$ takes the highest weight $(1-\delta)^0 = 1$, thus keeping fluctuations (hence, adaptation) alive.

Even though the sums in (34) and (35) exhibit a markedly different behavior in terms of their time-evolution (i.e., on the sample paths), one notable conclusion from Theorem 1 is that

their *probability distributions* converge to the same distribution, that is the distribution of the limiting variable $\tilde{\lambda}_k^{(\delta)}$. This equivalence can be explained as follows. With reference to the top panel in Fig. 4, consider a sufficiently large i (say, $i = 300$) and take the corresponding values of the dashed curve and the solid curve, namely, $\tilde{\lambda}_{k,300}(2)$ and $\hat{\lambda}_{k,300}(2)$. These values are different. However, if we now repeat the experiment in Fig. 4 several times, the realizations of $\hat{\lambda}_{k,300}(2)$ across different experiments will be distributed in the same way as the realizations of $\tilde{\lambda}_{k,300}(2)$.

The existence of a limiting distribution for the log-belief vector $\hat{\lambda}_{k,i}^{(\delta)}$ makes the definition of a *steady-state* error probability meaningful, since from Eqs. (18) and (19) we see that the steady-state error probability can be computed as⁴:

$$p_k^{(\delta)} = \mathbb{P} \left[\exists \theta \neq \theta_0 : \tilde{\lambda}_k^{(\delta)}(\theta) \leq 0 \right]. \quad (36)$$

However, it should be noticed that Theorem 1 constitutes only a first, albeit fundamental step towards the characterization of the ASL performance, since it establishes only the existence of a steady-state error probability without providing any explicit characterization thereof. Such characterization is in general not available. In the next sections we tackle this challenging problem by focusing on an asymptotic characterization of $\tilde{\lambda}_k^{(\delta)}$ in the regime of small δ .

VI. SMALL- δ ANALYSIS

We have ascertained that it makes sense to define *steady-state* random variables characterizing the log-belief ratios. Then, the steady-state learning performance can be determined by examining the probability that these random variables fulfill certain conditions. For example, the steady-state probability that an agent learns the truth is the probability that the steady-state log-belief ratio of that agent is positive only at the true value θ_0 . However, in general the exact characterization of these steady-state variables is a formidable task. For this reason we will resort to an asymptotic analysis in the regime of small δ . We will provide three types of asymptotic results.

- **Sec. VI-A: Weak law of small step-sizes (Theorem 2).** We will show that, for small δ , the steady-state vector $\tilde{\lambda}_k^{(\delta)}$ concentrates around the weighted average of the agents' KL divergences defined in (37). This concentration property guarantees that, with high probability as $\delta \rightarrow 0$, the true hypothesis is chosen by each agent. This result will require only finiteness of the first moments of the log-likelihood ratios, i.e., finiteness of the KL divergences.
- **Sec. VI-B: Asymptotic normality (Theorem 3).** We will obtain a Central Limit Theorem (CLT) that will provide a normal approximation, holding for small δ , for the error probabilities of each individual agent. This result will be proved assuming independence across agents and will

⁴According to the definition of convergence in distribution, the result in (36) holds provided that the limiting random variable $\tilde{\lambda}_k^{(\delta)}$ has no point mass at 0. However, we rule out such pathological case that is in practice the exception rather than the rule.

require finiteness of the variance of the log-likelihood ratios. We remark that previous results of asymptotic normality for adaptive distributed detection assumed finiteness of higher-order moments [27]. To the best of our knowledge, the result in Theorem 3 (which is based on part 5 of Lemma 1) is the first result that assumes the minimal requirement of finiteness of second moments.

- **Sec. VI-C: Large deviations analysis (Theorem 4).** We will characterize the exponential rate of decay of the error probabilities as $\delta \rightarrow 0$. This result will be proved assuming independence across agents and will require the existence of the moment generating function of the log-likelihood ratios.

Notably, the above three steps reflect perfectly a classic path in asymptotic statistics. However, in order to avoid misunderstandings, it is necessary to clarify one fundamental difference between the small- δ analysis and classic results. In order to illustrate this difference let us refer, for example, to the CLT result. In the traditional setting of asymptotic statistics, one examines the asymptotic behavior of sums of random variables when the number of terms of the sum goes to infinity. In contrast, the CLT proved in this work does *not* affirm that the sums involved in (27) converge to a Gaussian as $i \rightarrow \infty$. As a matter of fact, we have shown in Theorem 1 that the sums in (27) converge to certain random variables, but these variables are *not Gaussian*, in general. The CLT that we prove deals instead with the behavior, as δ goes to zero, of the steady-state random vector $\tilde{\lambda}_k^{(\delta)}$. The same distinction applies to the other two types of asymptotic results, namely, the weak law and the large deviations analysis. For this reason, as explained in [21], the correct way to deal with the asymptotic regime of small step-sizes in the adaptation context is made of two steps:

- First introduce a proper steady-state vector $\tilde{\lambda}_k^{(\delta)}$, which already embodies the effect of combining an infinite number of summands. This steady-state vector will be non-degenerate (i.e., no weak law as $i \rightarrow \infty$), will be non-Gaussian (i.e., no CLT as $i \rightarrow \infty$), and will be non-vanishing (i.e., no large deviations as $i \rightarrow \infty$).
- Then, characterize the asymptotic behavior of the steady-state random vector $\tilde{\lambda}_k^{(\delta)}$ as δ goes to zero.

It is worth noticing that, in the adaptation literature, the critical role of the first step is usually not emphasized. This is because the adaptation literature mostly focuses on *estimation* problems, where one usually quantifies the performance by evaluating convergence of the *moments* [23]. In contrast, when dealing with *decision* problems (as in our case), the performance is quantified through *probabilities*, namely, the probabilities of making a wrong (or correct) decision. In order to evaluate probabilities at the steady state, it is critical to obtain first a representation of the steady-state random variables [21].

A. Consistent Social Learning

We will establish that the ASL strategy achieves consistent social learning under the following standard assumption of global identifiability.

Assumption 3 (Global Identifiability): For each wrong hypothesis $\theta \neq \theta_0$, there is at least one agent that has strictly positive KL divergence. \square

Let us provide some intuition behind Assumption 3. Consider agent k and hypothesis $\theta \neq \theta_0$. Now, if the likelihoods $L_k(\xi|\theta)$ and $L_k(\xi|\theta_0)$ are equal, θ is not distinguishable from θ_0 at agent k , i.e., the classification problem is *locally non-identifiable*. Clearly, if there exists a hypothesis θ that is indistinguishable from θ_0 at *all* agents, there is no hope for the system to classify correctly, because the agents will be necessarily uncertain between θ and θ_0 . Therefore, a minimal requirement for global identifiability is that, for each $\theta \neq \theta_0$, there exists at least one agent for which model $L_k(\xi|\theta)$ is distinct from $L_k(\xi|\theta_0)$. This is exactly what Assumption 3 requires. It is also useful to highlight that Assumption 3 does not imply in any manner that agent k would be able to classify locally. In fact, saying that agent k is able to distinguish θ from θ_0 does not mean that it can distinguish θ_0 from the remaining hypotheses $\theta' \notin \{\theta, \theta_0\}$.

We are now ready to state the theorem that establishes achievability of consistent learning. To this end, it is useful to introduce the expectation of the average log-likelihood ratio in (22):

$$m_{\text{ave}}(\theta) \triangleq \mathbb{E}[\mathbf{x}_{\text{ave},i}(\theta)] = \sum_{\ell=1}^N \pi_{\ell} d_{\ell}(\theta), \quad (37)$$

which does not depend on i owing to the identical distribution over time implied by the steady-state analysis.

Theorem 2 (Consistency of ASL): Under Assumptions 1 and 2, we have the following convergence:

$$\tilde{\lambda}_k^{(\delta)} \xrightarrow{\delta \rightarrow 0} m_{\text{ave}} \quad \text{in probability.} \quad (38)$$

Since under Assumption 3 all entries of m_{ave} are strictly positive, Eq. (38) implies that each agent learns correctly the true hypothesis as $\delta \rightarrow 0$, namely, for all $\theta \neq \theta_0$ we have that the steady-state error probability of all agents $k = 1, 2, \dots, N$ converges to zero as δ approaches zero:

$$\lim_{\delta \rightarrow 0} p_k^{(\delta)} = 0. \quad (39)$$

Proof: See Appendix C. \blacksquare

The result of Theorem 2 relies on the weak law of small step-sizes proved in Lemma 1, part 3. Technically, this law requires finiteness of only the first moments $d_{\ell}(\theta)$, which is guaranteed by Assumption 1. Moreover, the result of Theorem 2 requires that $m_{\text{ave}}(\theta) > 0$ for all $\theta \neq \theta_0$. Since the entries of the Perron eigenvector are all strictly positive, we see that $m_{\text{ave}}(\theta)$ is strictly greater than zero for every θ if, for every θ , there exists at least one agent ℓ for which the KL divergence $d_{\ell}(\theta)$ is strictly positive. In other words, in order to achieve consistent learning, it is sufficient that at least one of the first moments (i.e., the KL divergence) is nonzero, which is guaranteed by Assumption 3.

Therefore, we see that Assumption 3 provides one important motivation for agents' cooperation in social learning. In fact, we assume that the learning problem can be non-identifiable (i.e., can be singular) *locally*, meaning that an individual agent can have one or more hypotheses $\theta \neq \theta_0$ that are

indistinguishable from the true one (zero KL divergence). If this happens, an individual agent is not able to learn properly. On the other hand, under a *global* identifiability condition, the network is able (as shown in Theorem 2) to identify the true hypothesis by fusing the information coming from distinct agents.

We have shown that the ASL strategy allows correct learning of the true hypothesis for sufficiently small step-sizes. In other words, we have established that the error probability vanishes as $\delta \rightarrow 0$. On the other hand, we have not established *how* it vanishes. There are at least two good reasons to examine the way this probability converges to zero. The first reason is to get manageable formulas for the evaluation of the social learning performance. The second reason is to characterize the fundamental scaling laws of the system. We will see that the ASL strategy is characterized by an exponential law, since the error probability of each individual agent decays exponentially fast as a function of the inverse step-size $1/\delta$.

B. Normal Approximation for Small δ

We will now prove a central limit theorem for the steady-state random vector $\tilde{\lambda}_k^{(\delta)}$. To this end, we will assume finiteness of second-order moments for the log-likelihoods. We furthermore assume statistical independence across the agents.

In order to state the CLT, it is convenient to define some useful quantities. First, we introduce the covariance between the log-likelihood ratios at θ and θ' , that is:

$$\rho_\ell(\theta, \theta') = \mathbb{E} \left[\left(\mathbf{x}_{\ell,i}(\theta) - d_\ell(\theta) \right) \left(\mathbf{x}_{\ell,i}(\theta') - d_\ell(\theta') \right) \right]. \quad (40)$$

Then we introduce the covariance between the average variables $\mathbf{x}_{\text{ave},i}(\theta)$ and $\mathbf{x}_{\text{ave},i}(\theta')$ which, exploiting independence across agents, can be evaluated as:

$$\mathbf{C}_{\text{ave}}(\theta, \theta') \triangleq \sum_{\ell=1}^N \pi_\ell^2 \rho_\ell(\theta, \theta'). \quad (41)$$

Next, it is necessary to examine the behavior of the first two moments of the log-belief ratios. In view of Lemma 1, part 2, it is possible to conclude that the expectation of the steady-state random vector $\tilde{\lambda}_k^{(\delta)}$ can be expressed as:

$$\mathbf{m}_k^{(\delta)}(\theta) \triangleq \mathbb{E} \left[\tilde{\lambda}_k^{(\delta)}(\theta) \right] = \mathbf{m}_{\text{ave}}(\theta) + O(\delta), \quad (42)$$

where $O(\delta)$ is a quantity such that the ratio $O(\delta)/\delta$ remains bounded as $\delta \rightarrow 0$. Likewise, using part 4 of Lemma 1, we conclude that the covariance of the steady-state random vector $\tilde{\lambda}_k^{(\delta)}$ is:

$$\begin{aligned} & c_k^{(\delta)}(\theta, \theta') \\ & \triangleq \mathbb{E} \left[\left(\tilde{\lambda}_k^{(\delta)}(\theta) - \mathbf{m}_k^{(\delta)}(\theta) \right) \left(\tilde{\lambda}_k^{(\delta)}(\theta') - \mathbf{m}_k^{(\delta)}(\theta') \right) \right] \\ & = \frac{\mathbf{C}_{\text{ave}}(\theta, \theta')}{2} \delta + O(\delta^2). \end{aligned} \quad (43)$$

Equations (42) and (43) can be rewritten in vector and matrix form, respectively as:

$$\mathbf{m}_k^{(\delta)} = \mathbf{m}_{\text{ave}} + O(\delta), \quad \mathbf{C}_k^{(\delta)} = \frac{\mathbf{C}_{\text{ave}}}{2} \delta + O(\delta^2), \quad (44)$$

where $\mathbf{C}_k^{(\delta)} = [c_k^{(\delta)}(\theta, \theta')]$ and $\mathbf{C}_{\text{ave}} = [c_{\text{ave}}(\theta, \theta')]$ are the matrices that collect the individual covariances. We see from (44) that, as $\delta \rightarrow 0$, there is a leading term that does not depend on the agent index k (whose impact is implicitly included in the higher order corrections, i.e., the $O(\cdot)$ terms).

The first relation in (44) reveals that the expectation vector of the steady-state log-belief ratios, $\mathbf{m}_k^{(\delta)}$, approximates, for small δ , the expectation vector of the *average* log-likelihood ratios, \mathbf{m}_{ave} . In comparison, the second relation in (44) reveals that the covariance matrix of the steady-state log-belief ratios, $\mathbf{C}_k^{(\delta)}$, goes to zero as $\mathbf{C}_{\text{ave}} \delta/2$, where \mathbf{C}_{ave} is the covariance matrix of the *average* log-likelihood ratios, namely,

$$\lim_{\delta \rightarrow 0} \frac{2\mathbf{C}_k^{(\delta)}}{\delta} = \mathbf{C}_{\text{ave}}. \quad (45)$$

We are now ready to state our central limit theorem.

Theorem 3 (Asymptotic Normality): Assume that the data $\{\xi_{k,i}\}$ are independent across the agents (recall that they are always assumed i.i.d. over time), and that the log-likelihood ratios have finite variance. Then, under Assumptions 1, 2 and 3, the following convergence holds:

$$\frac{\tilde{\lambda}_k^{(\delta)} - \mathbf{m}_{\text{ave}}}{\sqrt{\delta}} \overset{\delta \rightarrow 0}{\rightsquigarrow} \mathcal{G} \left(0, \frac{\mathbf{C}_{\text{ave}}}{2} \right), \quad (46)$$

where the symbol \rightsquigarrow denotes convergence in distribution, and $\mathcal{G}(0, C)$ is a zero-mean multivariate Gaussian with covariance matrix equal to C .

Proof: See Appendix D. ■

Theorem 3 entails the following approximation, holding for $\delta \approx 0$:

$$\tilde{\lambda}_k^{(\delta)} \approx \mathcal{G} \left(\mathbf{m}_{\text{ave}}, \frac{\mathbf{C}_{\text{ave}}}{2} \delta \right). \quad (47)$$

We see that such approximation does *not* depend on the agent index k . As shown in [21], in order to capture differences in performance across the agents, it is possible to replace the limiting expectation vector \mathbf{m}_{ave} and the limiting covariance matrix $\mathbf{C}_{\text{ave}} \delta/2$ with their exact counterparts, i.e., with the series appearing in (42) and (43), yielding the refined approximation:

$$\tilde{\lambda}_k^{(\delta)} \approx \mathcal{G} \left(\mathbf{m}_k^{(\delta)}, \mathbf{C}_k^{(\delta)} \right). \quad (48)$$

The approximations in (47) and (48) will be tested in the section devoted to numerical experiments.

C. Large Deviations for Small δ

In this section we focus on another relevant type of asymptotic analysis, namely, a *large deviations* analysis [28], [29]. The application of large deviations to adaptive networks was used in [21], [27], [30].

The basic aim of the LD analysis is to estimate the exponential decay rate of the probabilities associated to certain *rare* events. In our setting, the rare event is the probability that an agent opts for the wrong hypothesis. We will show that, at the steady state, this type of event becomes in fact rare as δ approaches zero.

More formally, the LD analysis will furnish the following type of representation for the steady-state error probability [28], [29]:

$$p_k^{(\delta)} \doteq e^{-\Phi/\delta}, \quad (49)$$

where the notation \doteq means equality to the leading exponential order (as $\delta \rightarrow 0$) or, more explicitly:

$$\lim_{\delta \rightarrow 0} \delta \log p_k^{(\delta)} = -\Phi, \quad (50)$$

for a certain value Φ that is called the *error exponent*. Notably, in the exponent Φ we did not put any dependence on the agent index k . This is because, as shown in Theorem 4 further ahead, *all agents will exhibit the same error exponent*.

On the other hand, it should be remarked that the equality at the leading exponential order in (49) does not imply in any way that we can approximate the probability of error as $e^{-\Phi/\delta}$, namely,

$$p_k^{(\delta)} \not\approx e^{-\Phi/\delta}. \quad (51)$$

This is because any LD analysis neglects sub-exponential corrections. For example, it is immediate to check that the probabilities $e^{-\Phi/\delta}$ and $100e^{-\Phi/\delta}$ have the same LD exponent (equal to Φ), but the second probability is two orders of magnitude larger. These sub-exponential corrections embody higher-order differences in the error probabilities (see, e.g., Fig. 3) that can arise across the agents due to different factors, for example, due to differences between very “central” agents with a high number of neighbors as opposed to “peripheral” agents with few neighbors. In order to compensate for sub-exponential corrections, a refined LD framework exists, usually referred to as “exact asymptotics,” which has been applied to binary adaptive detection in [21], [30].

In summary, the aim of a large deviations analysis is to evaluate the asymptotic decay rate of the error probabilities, which is a meaningful and significant index of the inferential performance. Since the error exponent is a compact statistical descriptor of the learning performance, it can be useful to compare different systems (e.g., ASL strategies with different network graphs) and/or to optimize some system parameters (e.g., the network graph) to achieve the fastest learning rate.

Before stating the main result about the LD analysis, it is necessary to introduce the Logarithmic Moment Generating Function (LMGF), a.k.a. cumulant generating function, of the log-likelihood ratios:

$$\Lambda_k(t; \theta) = \log \mathbb{E} \left[e^{t \mathbf{x}_{k,i}(\theta)} \right]. \quad (52)$$

We recall that, in the steady-state regime, the expectation is computed under the true model $L_k(\xi|\theta_0)$, which does not change over time, and this explains why $\Lambda_k(t; \theta)$ does not depend on i . It is also useful to introduce the LMGF of the average variable $\mathbf{x}_{\text{ave},i}(\theta)$ which, under the assumption that the data are independent across the agents, is:

$$\Lambda_{\text{ave}}(t; \theta) = \log \mathbb{E} \left[e^{t \mathbf{x}_{\text{ave},i}(\theta)} \right] = \sum_{\ell=1}^N \Lambda_{\ell}(\pi_{\ell} t; \theta). \quad (53)$$

Theorem 4 (Error Exponents): Assume that the data $\{\xi_{k,i}\}$ are independent across the agents (recall that they are always

assumed i.i.d. over time), and that the logarithmic moment generating function of $\mathbf{x}_{k,i}(\theta)$ exists everywhere, namely, for all $k = 1, 2, \dots, N$ and $\theta \neq \theta_0$:

$$\Lambda_k(t; \theta) < +\infty \quad \forall t \in \mathbb{R}. \quad (54)$$

Let

$$\phi(t; \theta) = \int_0^t \frac{\Lambda_{\text{ave}}(\tau; \theta)}{\tau} d\tau. \quad (55)$$

Then, under Assumptions 1, 2 and 3 we have the following two results holding for every agent $k = 1, 2, \dots, N$. First, we have that:

$$\mathbb{P} \left[\tilde{\lambda}_k^{(\delta)}(\theta) \leq 0 \right] \doteq e^{-\Phi(\theta)/\delta}, \quad \Phi(\theta) = -\inf_{t \in \mathbb{R}} \phi(t; \theta). \quad (56)$$

Second, the error probability is dominated by the worst-case (i.e., smaller) exponent:

$$p_k^{(\delta)} \doteq e^{-\Phi/\delta}, \quad \Phi = \min_{\theta \neq \theta_0} \Phi(\theta). \quad (57)$$

Proof: See Appendix E ■

The main message conveyed by Theorem 4 is that the steady-state error probability of each individual agent converges to zero as $\delta \rightarrow 0$, exponentially fast as a function of $1/\delta$. This exponential law provides a *universal* law for adaptive social learning, which reflects the universal scaling law of distributed adaptive detection — see [21]. The exponent Φ governing such an exponential decay is computed from the logarithmic moment generating function of the average log-likelihood, where the weights of this average are the limiting weights, i.e., the entries of the Perron eigenvector.

The need for cooperation has been already motivated in relation to social learning problems that are locally non-identifiable. Theorem 4 implies another potential benefit of cooperation, namely, that *cooperation improves the learning accuracy*. We will illustrate this aspect through one example. Assume the most favorable case where all agents could learn the true hypothesis individually. Consider then a doubly-stochastic combination matrix, yielding a Perron eigenvector with uniform entries $\pi_{\ell} = 1/N$ for all $\ell = 1, 2, \dots, N$. Exploiting (57), we can easily see that in this particular case the error exponent of the network is given by:

$$\Phi = N\Phi_{\text{ind}}, \quad (58)$$

where Φ_{ind} is the error exponent of an individual agent. According to (58), we see that the network error exponent is N times larger than the individual error exponent, which in turn implies an N -fold exponential improvement in the learning accuracy. Intuitively, a network of N agents observes N times as much data as a single agent at each time instant. The strong-connectivity of the network allows for the data to fully propagate across agents and yields the aforementioned learning performance improvement.

VII. TRANSIENT ANALYSIS

A. Qualitative Description of the Transient Phase

Preliminarily, we deem it is useful to provide a qualitative overview of the transient behavior of adaptive social learning in comparison to traditional social learning. To this

end, we consider initially a simple example consisting of a single-agent (indices k and ℓ dropped) binary ($\Theta = \{1, 2\}$) problem, with symmetric KL divergences:

$$\mathbb{E}_1 \left[\log \frac{L(\xi_i|1)}{L(\xi_i|2)} \right] = -\mathbb{E}_2 \left[\log \frac{L(\xi_i|1)}{L(\xi_i|2)} \right] \triangleq x > 0, \quad (59)$$

where \mathbb{E}_θ denotes expectation under the distribution $L(\xi|\theta)$. We assume that at time $i = 1$, the true underlying hypothesis is $\theta_0 = 1$, and the situation remains stationary until a certain time T_1 , after which data start being generated according to $\theta_0 = 2$, and that is why a transient analysis is necessary to see how the learning algorithm is able to track this drift.

In order to examine how the learning process progresses over time, it is sufficient to consider the time-evolution of the log-belief ratio:

$$r_i \triangleq \log \frac{\mu_i(1)}{\mu_i(2)}, \quad (60)$$

whose positive (resp., negative) values will let the agent opt for $\theta = 1$ (resp., $\theta = 2$). Specializing (2) and (3) to the single-agent binary setting, traditional social learning evolves according to the recursion (we add a superscript to distinguish traditional from adaptive social learning):

$$r_i^{\text{SL}} = r_{i-1}^{\text{SL}} + \log \frac{L(\xi_i|1)}{L(\xi_i|2)}, \quad r_0^{\text{SL}} = 0. \quad (61)$$

Likewise, replacing (2) with (7), the adaptive social learning strategy in this single-agent binary case evolves according to the recursion:

$$r_i^{\text{ASL}} = (1 - \delta)r_{i-1}^{\text{ASL}} + \delta \log \frac{L(\xi_i|1)}{L(\xi_i|2)}, \quad r_0^{\text{ASL}} = 0. \quad (62)$$

For the sake of concreteness, in both (61) and (62) we assume flat initial priors (i.e., $r_0^{\text{SL}} = r_0^{\text{ASL}} = 0$).

In order to get a flavor of the main trade-offs involved in the transient behavior, let us focus on the time-evolution of the *expected values*. Taking expectations in (61), at time T_1 we have:

$$\mathbb{E} [r_{T_1}^{\text{SL}}] = T_1 x, \quad (63)$$

where x is the symmetric KL divergence introduced in (59). Equation (63) shows that the expected value of the log-belief ratio grows linearly with the stationarity interval T_1 . This linear growth is a reflection of the increasing knowledge acquired by the agent as it aggregates new information represented by the log-likelihood ratio $\log \frac{L(\xi_i|1)}{L(\xi_i|2)}$. In a virtual asymptotic regime, this knowledge becomes a certainty, i.e., as $T_1 \rightarrow +\infty$, $r_{T_1}^{\text{SL}} \rightarrow +\infty$, which implies that if hypothesis 1 remains in force indefinitely, the belief of the agent regarding this hypothesis achieves full confidence. Unfortunately, this increasing confidence comes at a cost in terms of an elephant memory that makes the algorithm slow in adaptation. Indeed, since from time $T_1 + 1$ the true hypothesis is $\theta_0 = 2$, from (61) and (63) we have that:

$$\mathbb{E} [r_i^{\text{SL}}] = \mathbb{E} [r_{T_1}^{\text{SL}}] - ix = (T_1 - i)x. \quad (64)$$

Now, the adaptation time can be roughly identified by considering the time necessary to overcome the initial bias towards hypothesis 1 once the true hypothesis switches from




Fig. 5. Diagram of the time evolution of the log-belief ratio in expectation for the traditional social learning strategy (in blue) and for the ASL strategy (in violet) within the single-agent case.

1 to 2. In terms of our qualitative mean-value analysis, this is the time necessary for the expected log-belief ratio to change from positive to negative, which, in view of (64) implies that the adaptation time for the traditional social learning strategy is on the order of:

$$T_{\text{SL}} = T_1. \quad (65)$$

This behavior is clearly not admissible for an adaptive algorithm, since it implies that the time necessary to recover from a wrong opinion is proportional to the stationarity interval where this opinion was actually true! This behavior is illustrated in Fig. 5.

Let us switch to the adaptive strategy. Developing the recursion until time T_1 , from (62) we get, respectively:

$$\mathbb{E} [r_{T_1}^{\text{ASL}}] = \delta \sum_{m=0}^{T_1-1} (1 - \delta)^m x = (1 - (1 - \delta)^{T_1}) x \approx x, \quad (66)$$

where the approximation is motivated from assuming a sufficiently large T_1 . Considering then that from time $T_1 + 1$ onward the true hypothesis is $\theta_0 = 2$, Eqs. (62) and (66) yield, for any $i > T_1$:

$$\begin{aligned} \mathbb{E} [r_i^{\text{ASL}}] &= (1 - \delta)^i \mathbb{E} [r_{T_1}^{\text{ASL}}] - \delta \sum_{m=0}^{i-T_1-1} (1 - \delta)^m x \\ &\approx - (1 - 2(1 - \delta)^i) x. \end{aligned} \quad (67)$$

Now, equating (67) to zero to evaluate the adaptation time, we obtain:

$$T_{\text{ASL}} = \frac{\log 2}{\log(1 - \delta)^{-1}} \approx \frac{\log 2}{\delta}. \quad (68)$$

A visual comparison of the enhanced adaptation provided by the ASL strategy is exemplified in Fig. 5.

Comparing (68) against (65), we see that, in contrast to the undesirable behavior exhibited by traditional social learning, the adaptive formulation exhibits a controlled initial bias. This is because, after a relatively long stationarity interval T_1 , the expected log-belief is concentrated around a fixed value x , and the adaptation time will then increase roughly as $1/\delta$. In a nutshell, while the reaction capacity of traditional social learning is not controlled by design and is severely affected by the duration of previous stationarity intervals, in adaptive

social learning the adaptation time is not affected by previous stationarity intervals, and the effective memory is controlled through the step-size. This enhanced adaptivity comes at the price of learning accuracy. In fact, as we have established in the previous sections, the steady-state error probability does not converge to zero as time elapses, but converges to some stable value. However, this value vanishes exponentially fast as a function of $1/\delta$, highlighting the fundamental trade-off of adaptive social learning: the smaller the step-size δ , the smaller the error probability and the slower the adaptation.

In the theory of adaptation and learning, the transient analysis is typically performed by characterizing the evolution of suitable higher-order moments, such as second or fourth order moments of the pertinent statistics [23]. However, this analysis is more appropriate for estimation/regression problems where the focus of the transient analysis is to ascertain how long it takes for the pertinent system state to attain a prescribed neighborhood of the expected value. In our social learning setting, it is more appropriate to identify an adaptation time in terms of *error probabilities*. As established in Theorem 4, the behavior of these probabilities is governed by the logarithmic moment generating function of the observations which, as the name itself suggests, incorporates dependence upon *all moments*. Accordingly, a meaningful way to perform the transient analysis is to examine the time-evolution of logarithmic moment generating functions, rather than individual moments. This characterization constitutes the core of Theorem 5, which is introduced in the next section.

B. Quantitative Description of the Transient Phase

In this section, we provide a rigorous analysis to support the qualitative description of the transient behavior, seen in Sec. VII-A. We assume that the ASL strategy has been in operation for a certain arbitrary time i_0 . All the knowledge accumulated by the agents until this time is summarized in the belief vector μ_{i_0} . We remark that the evolution of the statistical models from $i = 0$ to $i = i_0$ is left completely arbitrary, that is, the system could have experienced several drifts in the statistical conditions, including change of the underlying hypotheses, data generated according to models that do not match the assumed likelihoods, and so on. From the ASL algorithm viewpoint, all these effects are summarized in the belief vector μ_{i_0} that acts as initial state at time i_0 . In order to perform the transient analysis, we assume that from $i_0 + 1$ onward, the true hypothesis is steadily equal to θ_0 , and will establish how much time is necessary to stay sufficiently close to the steady-state learning performance starting from a given (arbitrary) realization μ_{i_0} . As done before, to simplify the notation we set $i_0 = 0$ and the initial state becomes μ_0 .

As noticed at the end of the previous section, in a social learning problem the adaptation time should be properly related to the time-evolution of the error probability, and particularly to the time necessary for the *instantaneous* error probability to approach the *steady-state* error probability. Accordingly, in the next theorem we start by providing an upper bound on the *instantaneous* error probability introduced in (18).

Theorem 5 (Bounds on the Instantaneous Error Probability): The claim of the theorem holds under the same assumptions of Theorem 4. Let κ and β be the constants defined in Property 1, and let $t_\theta^* < 0$ be the unique solution to the equation:

$$\frac{\Lambda_{\text{ave}}(t_\theta^*, \theta)}{t_\theta^*} = 0. \quad (69)$$

Let

$$\lambda_{\text{ave},0}(\theta) = \sum_{\ell=1}^N \pi_\ell \lambda_{\ell,0}(\theta) \quad (70)$$

be the network average of the initial log-belief ratios $\lambda_{\ell,0}(\theta)$, and let, for all $\theta \neq \theta_0$:

$$\mathsf{K}_1(\theta) \triangleq |t_\theta^*| \left[m_{\text{ave}}(\theta) - \lambda_{\text{ave},0}(\theta) \right], \quad (71)$$

$$\mathsf{K}_2(\theta) \triangleq \kappa |t_\theta^*| \sum_{\ell=1}^N |\lambda_{\ell,0}(\theta)|. \quad (72)$$

Then, the instantaneous error probability $p_{k,i}^{(\delta)}$ is upper bounded as:

$$p_{k,i}^{(\delta)} \leq \sum_{\theta \neq \theta_0} e^{\frac{1}{\delta} [-\Phi(\theta) + \mathsf{K}_1(\theta)(1-\delta)^i + \mathsf{K}_2(\theta)(1-\delta)^i \beta^i + \mathcal{O}(\delta)]}, \quad (73)$$

where the notation $\mathcal{O}(\delta)$ signifies that the ratio $\mathcal{O}(\delta)/\delta$ stays bounded as $\delta \rightarrow 0$.

Proof: See Appendix F. ■

Theorem 5 reveals the main behavior of the transient error probability. Examining the error exponent of the upper bound in (73) we see, up to higher-order small- δ corrections embodied in the term $\mathcal{O}(\delta)$, the emergence of three terms: the *steady-state* error exponent $\Phi(\theta)$ already identified in Theorem 4, and two other terms that characterize the *transient* behavior. The first transient term decays as $(1-\delta)^i$, and is thus influenced solely by the step-size. The second transient term, $(1-\delta)^i \beta^i$, decays faster and is influenced also by the parameter β . This parameter, according to Property 1, is determined by the second largest-magnitude eigenvalue of A , and accordingly determines the mixing properties of A (i.e., the convergence rate of $[A^i]_{\ell k}$ to the Perron eigenvector entry π_ℓ). Therefore, the second transient term, with rate $(1-\delta)^i \beta^i$, determines a transient phenomenon that is related to the convergence of the matrix-powers to a “centralized” solution with combination weights π_ℓ . In comparison, the first term, with rate $(1-\delta)^i$, determines a transient phenomenon ruled by the step-size only.

In summary, Theorem 5 provides an upper bound on the instantaneous error probability that converges, as $i \rightarrow \infty$, to a sum of exponential terms with steady-state error exponent $\Phi = \min_{\theta \neq \theta_0} \Phi(\theta)$. Accordingly, we identify as a meaningful definition for the *adaptation time* the critical time instant after which the error probability decays with an error exponent $(1-\epsilon)\Phi$, for some small ϵ . This is made precise in the following corollary.

Corollary 1 (Adaptation Time): Under the same notation and assumptions of Theorem 5, let

$$\begin{aligned} \mathcal{K}_1 &\triangleq \max_{\theta \neq \theta_0} \mathcal{K}_1(\theta) = \max_{\theta \neq \theta_0} \left\{ |t_\theta^*| \left[m_{\text{ave}}(\theta) - \lambda_{\text{ave},0}(\theta) \right] \right\}, \\ \mathcal{K}_2 &\triangleq \max_{\theta \neq \theta_0} \mathcal{K}_2(\theta) = \kappa \max_{\theta \neq \theta_0} \left\{ |t_\theta^*| \sum_{\ell=1}^N |\lambda_{\ell,0}(\theta)| \right\}. \end{aligned} \quad (74)$$

Then, the upper bound:

$$p_{k,i}^{(\delta)} \leq e^{-\frac{1}{\delta}[(1-\epsilon)\Phi + \Theta(\delta)]} \quad (75)$$

holds for all $i > T_{\text{ASL}}$, where T_{ASL} is given by the following rules:

i) (Favorable case, all initial states are good).

If $\lambda_{\text{ave},0}(\theta) \geq m_{\text{ave}}(\theta)$ for all $\theta \neq \theta_0$:

$$T_{\text{ASL}} = \frac{1}{\log \beta^{-1}} \log \frac{\mathcal{K}_2}{\epsilon \Phi}, \quad \epsilon < \frac{\mathcal{K}_2}{\Phi}. \quad (76)$$

ii) (Unfavorable case, at least one initial state is bad).

If $\lambda_{\text{ave},0}(\theta) < m_{\text{ave}}(\theta)$ for at least one $\theta \neq \theta_0$:

$$T_{\text{ASL}} = \frac{1}{\log(1-\delta)^{-1}} \log \frac{\mathcal{K}_1}{\epsilon \Phi}, \quad \epsilon < \frac{\mathcal{K}_1}{\Phi}. \quad (77)$$

Proof: See Appendix G. ■

Let us now examine the main parameters and phenomena affecting the adaptation time T_{ASL} .

– *Memory.* The memory coming from the past algorithm evolution is summarized in the starting belief vector μ_0 , which in turn determines the average log-belief $\lambda_{\text{ave},0}(\theta)$.

First of all, we notice that an average initial state $\lambda_{\text{ave},0}(\theta)$ greater than $m_{\text{ave}}(\theta)$ creates already a (favorable) bias toward the true hypothesis. Accordingly, when $\lambda_{\text{ave},0}(\theta) \geq m_{\text{ave}}(\theta)$ the transient term $\mathcal{K}_1(\theta)(1-\delta)^i$ reduces the error probability since $\mathcal{K}_1(\theta) < 0$. In this case, the dominant transient term is $(1-\delta)^i \beta^i$, and the corresponding adaptation time in (76) is essentially determined by the mixing parameter β , i.e., by how fast the combination weights converge to the Perron eigenvector. Under this regime, the adaptation time *does not depend critically on the step-size*.

In comparison, the case where $\lambda_{\text{ave},0}(\theta) < m_{\text{ave}}(\theta)$ is the unfavorable case where we are, as $\lambda_{\text{ave},0}(\theta)$ decreases, progressively far from the steady-state. Under this regime, for small δ the dominant transient term is $\mathcal{K}_1(\theta)(1-\delta)^i$, and the adaptation time *scales with the step-size as* $1/\log(1-\delta)^{-1} \approx 1/\delta$.

One particularly interesting case is when the average initial state is negative. This happens, for example, when the initial state comes from a previous learning cycle where the agent converged to a certain hypothesis that has then changed at the beginning of the subsequent learning cycle. In line with intuition, the adaptation time (77) increases with increasing size of the wrong starting conditions. Moreover, this dependence upon the past states is only logarithmic, which reveals that the past algorithm evolution has not a dramatic impact on the adaptation time.

– *KL Divergences and Error Exponent.* By ignoring the initial state, Eq. (77) becomes:

$$T_{\text{ASL}} = \frac{1}{\log(1-\delta)^{-1}} \log \frac{\max_{\theta \neq \theta_0} [|t_\theta^*| m_{\text{ave}}(\theta)]}{\epsilon \Phi}. \quad (78)$$

From Property P2) in Lemma 2 (see Appendix F), we know that:

$$\Phi(\theta) \leq |t_\theta^*| m_{\text{ave}}(\theta), \quad (79)$$

which shows that the ratio $\max_{\theta \neq \theta_0} [|t_\theta^*| m_{\text{ave}}(\theta)] / \Phi$ appearing in (78) is greater than 1. Even if declaring a general behavior for this ratio for all statistical models is not obvious, we see that the numerator and the denominator are not independent. For example, having an “easier” detection problem where the KL divergences (numerator) increase typically corresponds to an increase of the error exponent (denominator) as well. However, in all cases the dependence on these parameters is not critical, since it is logarithmic.

– *Parameter t_θ^* .* First of all, to evaluate and interpret the bound on the adaptation time it is useful to remark that the term $|t_\theta^*|$ is comprised between $1/\pi_{\text{max}}$ and $1/\pi_{\text{min}}$ — see property P3) in Lemma 2. Apparently, these bounds introduce a dependence on the network parameters (i.e., on the Perron eigenvector). However, we should be careful here, and recall that the network error exponent Φ depends on the whole network as well. In order to get insights on this dependence, let us ignore the initial state and consider the case where all likelihoods are equal across agents and the combination matrix is doubly stochastic (yielding a uniform Perron eigenvector). Under these assumptions, from property P3) in Lemma 2 we get $t_\theta^* = -N$, and using (58) we obtain:


$$T_{\text{ASL}} = \frac{1}{\log(1-\delta)^{-1}} \log \frac{\max_{\theta \neq \theta_0} [m_{\text{ave}}(\theta) - \lambda_{\text{ave},0}(\theta)]}{\epsilon \Phi_{\text{ind}}}, \quad (80)$$

which shows how the network size appearing in the parameter $t_\theta^* = -N$ is perfectly compensated by the network size embodied in the network exponent $\Phi = N\Phi_{\text{ind}}$. Accordingly, we expect that the network parameters have a reduced impact on the transient time in (77), while, as observed before, the effect of the network is embodied in the parameter β controlling the higher-order transient term $(1-\delta)^i \beta^i$ in (73), which is neglected in the small- δ regime.

– *Parameter ϵ .* The smaller ϵ is, the closer the error exponent to the steady-state exponent Φ will be. Remarkably, the dependence is logarithmic in $1/\epsilon$, which means that this parameter is not critical.

– *Step-Size.* Finally, in the (more interesting) case where the initial state is not good, see (77), the adaptation time scales as $1/\delta$. We remark that this behavior matches well the qualitative analysis of Sec. VII-A.

The bottom line of Corollary 1 is that the adaptive capabilities of the ASL strategy are enhanced by a larger value of δ , by yielding a reduced adaptation time. A larger δ however is not always desirable, since it can reduce the accuracy in the decision-making process (as seen in Theorem 4, the steady-state probability of error is increased for larger δ). Both phenomena represent the trade-off adaptation vs. learning present in the ASL strategy and should be taken into account when designing δ . Such trade-off can be better summarized by combining Theorem 4 and Corollary 1, which shows that the error probability decays exponentially fast with

Fig. 6. Strongly-connected network topology with $N = 10$ agents.

the adaptation time, roughly as:

$$p_k^{(\delta)} \approx \exp \left\{ -\frac{\Phi}{\log[K_1 \times (\epsilon \Phi)^{-1}]} T_{\text{ASL}} \right\}. \quad (81)$$

– *Stability over Successive Learning Cycles.* The characterization of the transient stage provided by Theorem 5 and the related corollary is valid under an arbitrary choice of the starting state $\lambda_{k,0}$. However, as we have commented in the previous section, if we start from a wrong state the level of this state affects adversely the adaptation time. Therefore, some fundamental questions arise. Assume that the time axis is divided into successive intervals (learning cycles) wherein the system evolves under stationary conditions. Then, the belief accumulated at the end of a learning cycle can be wrong in relation to the subsequent learning cycle. How “wrong” are the initial beliefs at the beginning of a learning cycle as the algorithm progresses? Do these initial states compromise the learning capability of the algorithm over successive cycles? These fundamental questions can be answered by combined steady-state and transient analyses. In fact, from the steady-state analysis carried out in the previous sections, we learned that the steady-state log-belief ratios fluctuate in a small neighborhood (of size $\sim \sqrt{\delta}$) of the expected values of the pertinent KL divergences. This means that at the end of each cycle the ASL strategy converges to some *stable* state, i.e., a state that does not diverge as the step-size δ becomes small. As a result, the initial states of each learning cycle would evolve in a stable manner and, hence, do not compromise the learning performance of the algorithm, provided that the adaptation time is smaller than the duration of the learning cycles. These aspects will be more quantitatively illustrated in Sec. IX, with reference to specific illustrative examples.

VIII. ILLUSTRATIVE EXAMPLES

We consider the strongly-connected network of $N = 10$ agents displayed in Fig. 6. We assume that all agents have a self-loop (not displayed in the figure). Besides, the combination matrix is designed using an averaging rule, resulting in a left-stochastic matrix [23].

The network is faced with the following statistical learning problem. We consider a family of Laplace likelihood functions with scale parameter 1, seen in Fig. 7. Formally, we are given three Laplace densities:

$$f_n(\xi) = \frac{1}{2} \exp \{-|\xi - 0.1 n|\}, \quad (82)$$




Fig. 7. Family of Laplace likelihood functions.

TABLE I
IDENTIFIABILITY SETUP FOR THE NETWORK IN FIG. 6

Agent k	Likelihood Function: $L_k(\xi \theta)$		
	$\theta = 1$	$\theta = 2$	$\theta = 3$
1 – 3	$f_1(\xi)$	$f_1(\xi)$	$f_3(\xi)$
4 – 6	$f_1(\xi)$	$f_3(\xi)$	$f_3(\xi)$
7 – 10	$f_1(\xi)$	$f_2(\xi)$	$f_1(\xi)$

for $n \in \{1, 2, 3\}$. The likelihoods of the data collected by the agents are chosen from among these Laplace densities.

To make things more interesting, we assume that the inference problem is *not locally identifiable*. The setup for each agent’s family of likelihood functions can be seen in Table I.

In summary, the data $\{\xi_{k,i}\}$ are i.i.d. (across time and agents) Laplace random variables, with expectations that depend both on the agent k and the hypothesis θ . Accordingly, we will use the notation $e_k(\theta)$ to denote the expectation of $\xi_{k,i}$, computed under likelihood $L_k(\xi|\theta)$. For example, using Table I, we see that:

$$e_1(1) = 0.1, \quad e_4(3) = 0.3, \quad e_7(2) = 0.2. \quad (83)$$

We are now ready to delve into a detailed illustration of the numerical experiments. In particular, in this section we will test how the empirical performance matches the steady-state performance as characterized in Theorems 1–4. In order to examine the steady-state behavior *empirically*, we need that the ASL algorithm run for a sufficiently long period of time. In line with the prescriptions from Sec. VII, the duration of this this period is chosen as at least one order of magnitude larger than the inverse of the step-size, $1/\delta$.

A. Consistency

We consider that all agents are running the ASL algorithm for a fixed $\theta_0 = 1$ over 8000 time samples (after which we consider that they achieved the steady state). From Theorem 2, we saw that as δ approaches zero, all agents k are able to consistently learn — see (38). In order to show this effect, for each value of δ (50 sample points in the interval $\delta \in [0.001, 1]$ are taken), we consider a different realization of the observations. In Fig. 8, for agent 1 and $\theta = 2, 3$, we show how the log-belief ratios $\lambda_1^{(\delta)}(\theta)$ behave for decreasing values of δ . We see the weak-law of small step-sizes arising, since the limiting log-belief ratios tend to concentrate around m_{ave} .




Fig. 8. Consistency of the ASL strategy (Theorem 2). According to the weak-law of small step-sizes, the steady-state log-belief ratios for agent 1 concentrate around the predicted expectation values in m_{ave} as δ approaches zero.

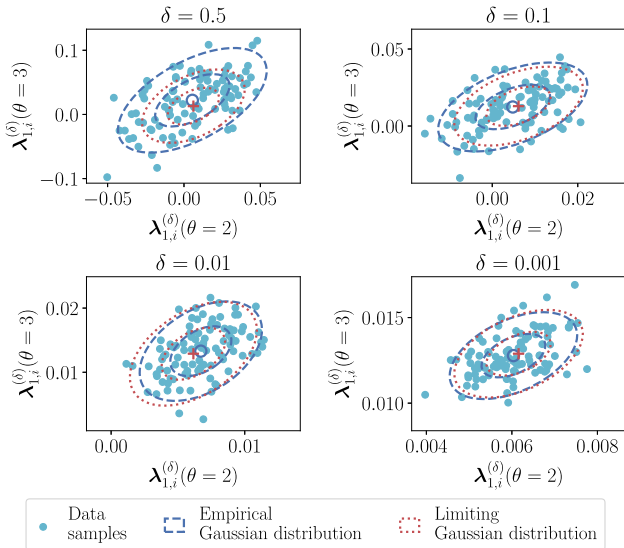


Fig. 9. Distribution of data samples at steady state compared with the limiting and empirical Gaussian distributions.

B. Asymptotic Normality

We consider 10000 time samples, where again all agents are collecting data under a true hypothesis $\theta_0 = 1$. From Theorem 3, we saw that in steady state we can approximate the log-belief ratios distribution by a multivariate Gaussian pdf, see Eqs. (47) and (48). In Fig. 9, we assume that the ASL algorithm has reached the steady state at $i = 10000$, and display the log-likelihood ratios corresponding to instant $i = 10000$. The experiment is repeated over 100 Monte Carlo runs, such that we obtain 100 realizations of the steady-state variable $\lambda_k^{(\delta)}$. Moreover, we consider 4 values of δ .

In dashed blue lines we see the ellipses representing the confidence intervals relative to one and two standard deviations computed for the empirical Gaussian approximation seen in (48): the smaller ellipse encompasses approximately 68% of the samples whereas the larger ellipse encompasses 95%. In red dotted lines, we see the corresponding ellipses for the limiting theoretical Gaussian approximation seen in (47), with the red cross indicating the limiting theoretical expectation m_{ave} . Note how as δ decreases, the ellipses tend to be smaller, which is in accordance with the scaling of the covariance matrices by δ in (47) and (48), and the distributions tend to

overlap, which is in accordance with the behavior predicted by Theorem 3.

C. Error Exponents

We start by evaluating the theoretical exponents for the Laplace example at hand. To this aim, we need to compute first the logarithmic moment generating function of the log-likelihood ratios $\mathbf{x}_{k,i}(\theta)$ in (20). Since the data follow a Laplace distribution, the log-likelihood ratio is:

$$\mathbf{x}_{k,i}(\theta) = |\xi_{k,i} - e_k(\theta)| - |\xi_{k,i} - e_k(\theta_0)|. \quad (84)$$

Before we proceed to characterize the random variable $\mathbf{x}_{k,i}(\theta)$, let us define the auxiliary quantity:

$$\Delta_{k,\theta} \triangleq e_k(\theta) - e_k(\theta_0). \quad (85)$$

We also introduce the centered variable $\tilde{\xi}_{k,i} = \xi_{k,i} - e_k(\theta_0)$, and therefore we can write:

$$\mathbf{x}_{k,i}(\theta) = |\tilde{\xi}_{k,i} - \Delta_{k,\theta}| - |\tilde{\xi}_{k,i}|. \quad (86)$$

For the case in which $\Delta_{k,\theta} > 0$, the random variable $\mathbf{x}_{k,i}(\theta)$ depends on the random variable $\tilde{\xi}_{k,i}$ in the following manner:

$$\mathbf{x}_{k,i}(\theta) = \begin{cases} -\Delta_{k,\theta}, & \text{if } \tilde{\xi}_{k,i} > \Delta_{k,\theta}, \\ \Delta_{k,\theta} - 2\tilde{\xi}_{k,i}, & \text{if } \tilde{\xi}_{k,i} \in [0, \Delta_{k,\theta}], \\ \Delta_{k,\theta}, & \text{if } \tilde{\xi}_{k,i} < 0. \end{cases} \quad (87)$$

We can then express the cumulative distribution function of $\mathbf{x}_{k,i}(\theta)$ as

$$\begin{aligned} \mathbb{P}[\mathbf{x}_{k,i}(\theta) \leq x] &= \begin{cases} 0, & \text{if } x < -\Delta_{k,\theta}, \\ \mathbb{P}\left[\tilde{\xi}_{k,i} \geq \frac{\Delta_{k,\theta} - x}{2}\right], & \text{if } x \in [-\Delta_{k,\theta}, \Delta_{k,\theta}], \\ 1, & \text{if } x > \Delta_{k,\theta}, \end{cases} \end{aligned} \quad (88)$$

where $\mathbb{P}[A]$ is the probability of event A , computed from the distribution of $\tilde{\xi}_{k,i}$. Note that its probability density function is given by $L_k(\xi + e_k(\theta_0)|\theta_0)$, which is a Laplace distribution with zero mean and scale parameter 1.

From the cumulative distribution function in (88), we can derive the density function of $\mathbf{x}_{k,i}(\theta)$ as:

$$\begin{aligned} p(x) &= \mathbb{P}\left[\tilde{\xi}_{k,i} > \Delta_{k,\theta}\right] \delta(x + \Delta_{k,\theta}) \\ &+ \mathbb{P}\left[\tilde{\xi}_{k,i} < 0\right] \delta(x - \Delta_{k,\theta}) \\ &+ \frac{1}{2} L_k\left(\frac{\Delta_{k,\theta} - x}{2} + e_k(\theta_0) \middle| \theta_0\right) \text{rect}\left(\frac{x}{2\Delta_{k,\theta}}\right), \\ &= \frac{1}{2} \exp[-\Delta_{k,\theta}] \delta(x + \Delta_{k,\theta}) + \frac{1}{2} \delta(x - \Delta_{k,\theta}) \\ &+ \frac{1}{4} \exp\left[-\frac{(\Delta_{k,\theta} - x)}{2}\right] \text{rect}\left(\frac{x}{2\Delta_{k,\theta}}\right), \end{aligned} \quad (89)$$

where $\text{rect}(\cdot)$ is the rectangle function, i.e., it is equal to 1 in the interval $]-\frac{1}{2}, \frac{1}{2}[$ and 0 elsewhere. Also we should distinguish the notation $\delta(x)$, which represents the Dirac delta-function, from the notation δ , which refers to the step-size parameter.




Fig. 10. Steady-state error probability. Markers refer to the empirical probability curves estimated via Monte Carlo simulation. The dotted line refers to the theoretical error probability in (18) computed using the Gaussian approximation in (47). The slope of the probability curves is compared against the slope Φ (i.e., the error exponent) predicted by Theorem 4, and shown with dashed line.

The LMGF of variable $x_{k,i}(\theta)$, whose expression was seen in (52), can be explicitly computed using (89):

$$\begin{aligned} \Lambda_k(t; \theta) &= \log \left(\int_{\mathbb{R}} e^{tx} p(x) dx \right) \\ &= \log \left[\frac{1}{2} \exp(-\Delta_{k,\theta}(t+1)) + \frac{1}{2} \exp(\Delta_{k,\theta}t) \right. \\ &\quad \left. + \frac{1}{2} \exp\left(-\frac{\Delta_{k,\theta}}{2}\right) \frac{\sinh(\Delta_{k,\theta}(t+1/2))}{t+1/2} \right]. \end{aligned} \quad (90)$$

If similar steps are followed for the case $\Delta_{k,\theta} < 0$, we would find the following expression for the LMGF:

$$\begin{aligned} \Lambda_k(t; \theta) &= \log \left[\frac{1}{2} \exp(\Delta_{k,\theta}(t+1)) + \frac{1}{2} \exp(-\Delta_{k,\theta}t) \right. \\ &\quad \left. - \frac{1}{2} \exp\left(\frac{\Delta_{k,\theta}}{2}\right) \frac{\sinh(\Delta_{k,\theta}(t+1/2))}{t+1/2} \right]. \end{aligned} \quad (91)$$

Assuming that the true state is $\theta_0 = 1$, we can then evaluate numerically $\Phi(\theta)$ by employing the expressions in Theorem 4, for $\theta = 2$ and $\theta = 3$, from which we obtain $\Phi(2) = 0.03348$ and $\Phi(3) = 0.05051$. Finally, the error probability dominant exponent is given by:

$$\Phi = \min_{\theta \in \{2,3\}} \Phi(\theta) = 0.03348 \quad (92)$$

Now we illustrate the details of the numerical experiments. We consider that the true state of nature is set as $\theta_0 = 1$, and we let all agents execute the ASL algorithm for 3000 iterations and for 20 values of δ in the interval $[1/150, 1/10]$. We run 20000 Monte Carlo experiments and we compute the steady-state empirical probability of error for each agent and each value of δ . In Fig. 10, the empirical probability curves of agents 1, 3, 6, 7, 9 are compared against the theoretical error probability in (18) computed using the Gaussian approximation in (47). The slope of these curves is compared against the slope Φ (i.e., the error exponent) predicted by Theorem 4.

IX. EVOLUTION OVER SUCCESSIVE LEARNING CYCLES

In this section we would like to focus on a specific nonstationary setting to illustrate in more detail the role of adaptation. We consider the time axis can be divided into successive *random* intervals (*learning cycles*) wherein the system conditions remain stationary. We do not focus here on situations where the system parameters can vary smoothly at each time instant following some “trajectory,” as happens, e.g., in tracking applications. While from the analysis of similar algorithms we can expect that the ASL strategy possesses some inherent tracking ability, the study of this scenario is left for future work.

We examine an environment where there are three different sources of nonstationarity, which will be modeled as (mutually independent) homogeneous Markov chains, as now specified:

- The true hypothesis can change over time. For $i = 1, 2, \dots$, the true state of nature at time i , denoted by $\theta_0(i)$, follows a Markov process with possible states in $\Theta = \{1, 2, 3\}$ and with transition probabilities described by the finite-state diagram in Fig. 11 (where only transition probabilities are displayed, with the complementary probabilities of remaining in a state being omitted).
- The combination policy can change over time. We assume that the agents employ two possible combination matrices, one doubly-stochastic, the other left-stochastic. For $i = 1, 2, \dots$, the combination matrix in force at time i , denoted by $\mathbf{A}(i)$, follows a Markov process with transition matrix represented by the corresponding finite-state diagram in Fig. 11.
- The system can be in one of three possible functioning states, namely, nominal, perturbed, and bad. For $i = 1, 2, \dots$, the operating state at time i is denoted by $\mathbf{f}(i)$. Under state $\mathbf{f}(i) = \text{nominal}$, the data are generated according to the true likelihood corresponding to hypothesis $\theta_0(i)$. Under state $\mathbf{f}(i) = \text{perturbed}$, some noise is added to perturb the true data model (while the agents still rely on the nominal likelihood to run their ASL strategy). State $\mathbf{f}(i) = \text{bad}$ corresponds to a failure of the system, where a large amount of noise is added to the data so as to impair the learning process. The transition matrix of the functioning process is encoded in the pertinent finite-state diagram in Fig. 11.

Let us evaluate the average duration of a learning cycle. In order to be conservative, we focus on the worst case, i.e., on the shorter average duration, which is obtained when the system is in the most unstable case (i.e., the state where transitions are more frequent). Examining Fig. 11, the most unstable state is obtained when: *i*) the hypothesis in force is $\theta_0(i) = 2$, since from such intermediate state the Markov chain can move leftward or rightward, while from the other states it cannot; *ii*) the combination policy is either left stochastic or doubly stochastic; and *iii*) the system works under a perturbed state of functioning, for the same reasons as in point *i*). Now, given that the overall system is in the joint state $\{\theta_0(i) = 2, \mathbf{A}(i) = \text{left stochastic}, \mathbf{f}(i) = \text{perturbed}\}$, the probability that the system remains stable for

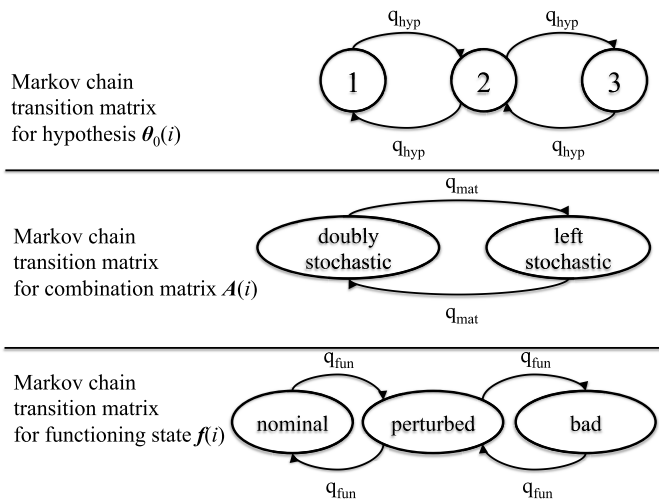


Fig. 11. Transition matrices of the Markov chains corresponding to the sources of nonstationarity illustrated in Sec. IX.

a single step is equal to:

$$q^* = (1 - 2q_{\text{hyp}})(1 - q_{\text{mat}})(1 - 2q_{\text{fun}}). \quad (93)$$

Likewise, the probability that the system remains stable for a certain number of steps is ruled by a geometric distribution of parameter q^* , yielding the following average duration for the worst-case learning cycle:

$$T_{\text{LC}} = \frac{q^*}{1 - q^*}. \quad (94)$$

In order to model a nonstationary environment where the system parameters remain stable during the learning cycles, we take inspiration from the Gilbert-Elliott model typically employed to model random bursts of errors over communication channels [31], [32]. According to the Gilbert-Elliott model, the transition probabilities between states of the chain are kept small so as to ensure that the chain remains in the same state for some contiguous time samples (i.e., we have “bursts” where the same state is repeatedly observed).

For what concerns the nominal likelihood models, we use the following family of Laplace likelihood functions, for $n \in \{1, 2, 3\}$:

$$f_n(\xi) = \frac{1}{2} \exp\{-|\xi - n|\}, \quad (95)$$

under the same identifiability setup as in Table I. The network topology is the same as in Fig. 6, on top of which we build two possible combination matrices: a left-stochastic matrix obtained through a uniform-averaging combination policy, and a doubly-stochastic matrix obtained through a Laplacian combination policy [23]. Under this setting, we evaluate the adaptation time exploiting (77). Regarding the initial states appearing in (77), we assume that in a given learning cycle the system comes from a previous learning cycle where the agents converged to a hypothesis different from that in force during the current learning cycle. Then we consider the worst-case initial state, and further the worst-case over all possible θ

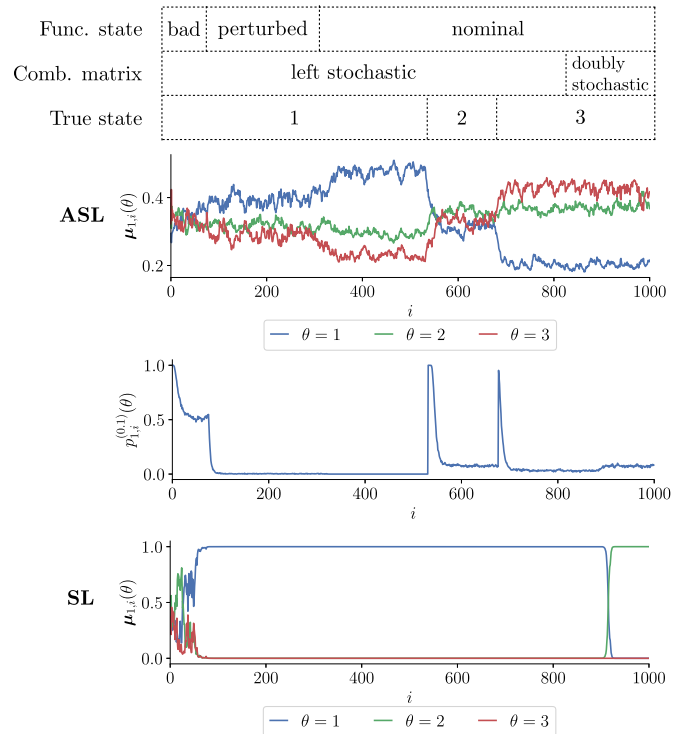


Fig. 12. Evolution of the learning strategies over successive learning cycles, with step-size $\delta = 0.1$ and average learning-cycle duration $T_{\text{LC}} \approx 100$. *First (top) row.* Observed transitions for the three sources of nonstationarity illustrated in the main text, namely, state of functioning, combination matrix, and hypothesis. *Second row.* Time-evolution of the belief at agent 1 for the *adaptive social learning strategy*. *Third row.* Time-evolution of the error probability at agent 1 for the *adaptive social learning strategy*. *Fourth row.* Time-evolution of the belief at agent 1 for the *traditional social learning strategy*.

and θ_0 . With these conservative choices, the time necessary to stay at 3 dB from the exponent Φ is equal to:

$$T_{\text{ASL}} \approx \frac{2.7286}{\delta}. \quad (96)$$

We now examine two settings that correspond to (relatively) short and long learning cycles, respectively.

– “Short” Learning Cycles. First of all, we consider that malfunctioning events and variations of the combination matrix are rare as compared to changes in the hypothesis. In particular, we set:

$$q_{\text{hyp}} = 5 \times 10^{-3}, \quad q_{\text{mat}} = 10^{-3}, \quad q_{\text{fun}} = 10^{-3}. \quad (97)$$

Exploiting (94), the average duration of a learning cycle can be approximated as $T_{\text{LC}} \approx 76$. If we equate the value found for T_{LC} to the adaptation time in (96), we get $\delta = 0.035$. For proper learning, we need that the adaptation time is smaller than the average duration of a learning cycle to ensure convergence to the correct hypothesis. In the experiments shown in Fig. 12 we made the choice:

$$\delta = 0.1, \quad (98)$$

which corresponds to an adaptation time not larger than one third of the average worst-case learning cycle. In Fig. 12, we display: in the second row, the time-evolution of the beliefs

at agent 1 corresponding to one realization of the process; in the third row, the corresponding error probability; and in the fourth row, the time-evolution of the beliefs at agent 1 for a traditional social learning strategy (same realization considered for the ASL strategy). During the considered time interval, several variations occurred, according to the nonstationary model described before.

First, we see that, except for the learning cycle corresponding to a bad state of functioning, the ASL strategy is able to learn well in all learning cycles, after a relatively short transient at the beginning of each learning cycle. The ability of learning well is showed by the time-evolution of the beliefs (second row), which shows how the maximum belief corresponds to the true hypothesis, after relatively short adaptation intervals necessary to react in face of nonstationarities. More quantitatively, the ability of learning is showed by the time-evolution of the error probabilities (third row), where we see some peaks (error probability close to 1) that clearly correspond to the changes, and that have a short duration dictated by the adaptation times. In sharp contrast, the traditional social learning strategy loses irremediably its learning ability yet after the first learning cycle.

Zooming in on Fig. 12, we see that nonstationarities in the hypotheses induce a perceivable change in the learning performance, whereas nonstationarities in the combination policy or in the state of functioning deserve a separate analysis.

For what concerns the combination policies, we see that the learning ability is preserved in face of a change, i.e., the system does *not* undergo an interval of poor performance. This behavior makes perfect sense, since from the theoretical analysis we know that the ASL strategy must consistently learn both with a left-stochastic or a doubly-stochastic matrix. What can be different is the learning performance, which depends on the Perron eigenvector. In this particular example, we have verified that, as opposed to the uniform Perron eigenvector corresponding to the doubly-stochastic matrix, the eigenvector of the left-stochastic matrix features higher weights corresponding to more informative agents (i.e., agents with higher KL divergences), which provides an explanation of the slightly increased performance observed in Fig. 12.

Regarding the state of functioning, we see that during the “bad” functioning state the data do not provide useful information, and the system undergoes an interval of failure (error probability ≈ 0.5). Remarkably, the adaptivity of the ASL strategy allows the agents to recover from this failure state in the successive learning cycles. In particular, the agents are able to recover and learn well already during the “perturbed” state of functioning. Actually, this regime of operation where the data do not follow any of the nominal likelihoods is not covered by our steady-state analysis. Our results could be in principle extended by allowing arbitrary distributions for the true data. In this case, as an inherent property of the ASL recursion, what is certainly preserved is the convergence to a *stable limiting state that does not diverge as the step-size δ becomes small*. What is lost is the relation between the limiting statistics and the KL divergences. Accordingly, the possibility of consistent learning would depend on the particular behavior of the limiting random variables under the non-nominal

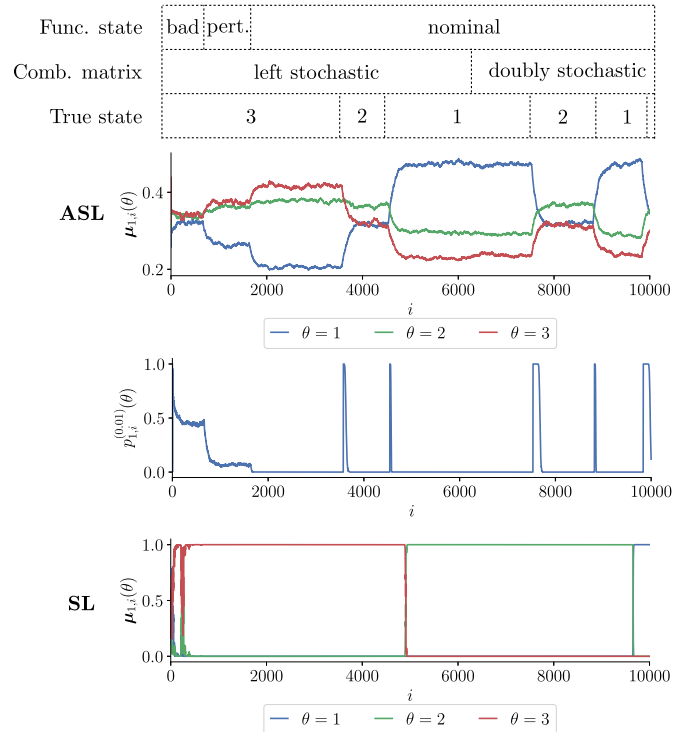


Fig. 13. Evolution of the learning strategies over successive learning cycles, with step-size $\delta = 0.01$ and average learning-cycle duration $T_{LC} \approx 1000$. *First (top) row.* Observed transitions for the three sources of nonstationarity illustrated in the main text, namely, state of functioning, combination matrix, and hypothesis. *Second row.* Time-evolution of the belief at agent 1 for the *adaptive* social learning strategy. *Third row.* Time-evolution of the error probability at agent 1 for the *adaptive* social learning strategy. *Fourth row.* Time-evolution of the belief at agent 1 for the *traditional* social learning strategy.

distributions. In light of these observations, it is not unexpected that, for reasonable amounts of perturbation, the agents are still able to learn, as happens in the considered example. Intuitively, we expect that, passing from a perturbed to a nominal state, the performance improves. Visually, this effect would be more clearly appreciated in the subsequent example shown in Fig. 13.

In summary, we see that the starting values at the beginning of each learning cycle are stable, since they arise as steady-state limiting values from at the end of the previous learning cycle. As such, these starting values do not diverge as time elapses, guaranteeing proper learning over successive learning cycles. This is a critical property, since it reveals that the number of variations of the underlying statistical conditions occurring during the entire algorithm evolution does not impair successful learning of the ASL strategy. What really matters is that the duration of the learning cycle is sufficiently large to allow a (small) value of δ to enable accurate learning.

– “Long” Learning Cycles. In Fig. 13, we consider the more favorable situation where the average duration of the learning cycle is increased by one order of magnitude, using the following transition probabilities for the pertinent Markov chains:

$$q_{\text{hyp}} = 5 \times 10^{-4}, \quad q_{\text{mat}} = 10^{-4}, \quad q_{\text{fun}} = 10^{-4}. \quad (99)$$

Accordingly, we expect that the adaptation properties of the system will be preserved if we reduce the step-size by one order of magnitude, yielding:

$$\delta = 0.01. \quad (100)$$

Comparing Fig. 13 against Fig. 12, we see that the general behavior is perfectly confirmed, and two notable effects emerge. First, the adaptation properties are preserved, i.e., the system is able to adapt to the changes sufficiently fast to guarantee a stable evolution over successive learning cycles. Second, the fluctuations around the limiting steady-state are reduced w.r.t. Fig. 12, yielding a smaller error probability, as it must be according to the theoretical analysis carried out in the previous sections since we are now using a smaller step-size $\delta = 0.01$.

X. CONCLUDING REMARKS

Social learning is a relevant inferential paradigm lying at the core of many multi-agent systems. Under this paradigm, the agents are able to learn progressively an underlying state of nature by continually updating their beliefs based on the incoming streaming data and the beliefs exchanged with their neighbors.

Several social learning implementations are currently available. However, these implementations are not open to deal with *nonstationary* data. For example, even if the agents learned well the true state, when this state changes at a certain instant, in the traditional social learning implementations the agents tend to be stubborn and keep on believing in the old state. In this work we proposed an Adaptive Social Learning (ASL) strategy that overcomes this issue and examined its performance and convergence guarantees in some great detail. The key insight is the introduction of an *adaptive update* depending on a step-size parameter δ that allows to tune the degree of adaptation. The introduction of the step-size δ allows the user to explore the trade-off between accuracy in decision making and adaptation speed.

A careful analysis of the system performance has been provided. In the steady-state phase, with focus on the small step-size regime, we have ascertained that the ASL strategy is able to learn consistently, and we have provided reliable performance characterization of the learning performance at each individual agent. In the transient phase, we have shown how the learning performance evolves over time and how the choice of the step-size affects the adaptation time.

Several useful extensions and generalizations are possible. One extension refers to removing the assumption that the agents' updates and data sharing are synchronous. An asynchronous version of the ASL strategy could be devised and examined exploiting the tools applied in the context of adaptive diffusion algorithms [33].

Another interesting generalization concerns the analysis of the ASL performance in *tracking* applications where some system parameters vary "smoothly," as opposed to the setting considered in this work where the system evolves under stationary conditions during individual learning cycles.

Finally, the technical machinery used to prove our results can be exploited to characterize adaptive social learning under

broader non-ideal settings, for example: when the true distribution is different from the nominal likelihoods assumed by the agents [17], [20]; when the learning task amounts to solving an optimization problem with multiple optimal solutions (i.e., there does not exist a unique truth from the agents' perspective) [17]; or when the network is weakly connected, giving rise to a dichotomy between influential and influenced agents [20].

APPENDIX A

In the following, the symbols \mathcal{S}° and $\bar{\mathcal{S}}$ denote the interior and the closure of set \mathcal{S} , respectively.

Lemma 1 (Asymptotic Properties of Random Series Useful for Adaptation): For $m = 0, 1, \dots$, let $\{z_m\}$ be a sequence of i.i.d. integrable random variables with:

$$m_z \triangleq \mathbb{E}[z_m], \quad m_z^{\text{abs}} \triangleq \mathbb{E}[|z_m|] < \infty. \quad (101)$$

Let also $0 < \delta < 1$, and consider the following partial sums:

$$s_i(\delta) = \delta \sum_{m=0}^i (1-\delta)^m \alpha_m z_m, \quad (102)$$

where $0 < \alpha_m \leq 1$, with α_m converging to some value $\alpha > 0$ and obeying the following upper bound for all m :

$$|\alpha_m - \alpha| \leq \kappa \beta^m, \quad (103)$$

for some constant $\kappa > 0$ and for some $0 < \beta < 1$. Then, we have the following asymptotic properties.

1. *Steady-State Stability:* The partial sums in (102) are almost-surely absolutely convergent, namely, we can define the (almost-surely) convergent series:

$$s^{\text{abs}}(\delta) \triangleq \delta \sum_{m=0}^{\infty} (1-\delta)^m \alpha_m |z_m|, \quad (104)$$

$$s(\delta) \triangleq \delta \sum_{m=0}^{\infty} (1-\delta)^m \alpha_m z_m. \quad (105)$$

2. *First Moment:* The expectation of $s(\delta)$ is:

$$\mathbb{E}[s(\delta)] = m_z \delta \sum_{m=0}^{\infty} (1-\delta)^m \alpha_m = \alpha m_z + O(\delta), \quad (106)$$

where $O(\delta)$ is a quantity such that the ratio $O(\delta)/\delta$ remains bounded as $\delta \rightarrow 0$.

3. *Weak Law of Small Step-Sizes:* The series $s(\delta)$ converges to αm_z in probability as $\delta \rightarrow 0$, namely, for all $\epsilon > 0$ we have that:

$$\lim_{\delta \rightarrow 0} \mathbb{P}[|s(\delta) - \alpha m_z| > \epsilon] = 0. \quad (107)$$

4. *Second Moment:* If:

$$\sigma_z^2 \triangleq \text{VAR}[z_m] < \infty, \quad (108)$$

then:

$$\begin{aligned} \text{VAR}[s(\delta)] &= \sigma_z^2 \delta^2 \sum_{m=0}^{\infty} (1-\delta)^{2m} \alpha_m^2 \\ &= \frac{\alpha^2 \sigma_z^2}{2} \delta + O(\delta^2). \end{aligned} \quad (109)$$




Fig. 14. Typical shape of the rate function.

5. *Asymptotic Normality*: If z_m has finite variance σ_z^2 , then the following convergence in distribution holds:

$$\frac{s(\delta) - m_z}{\sqrt{\delta}} \xrightarrow{\delta \rightarrow 0} \mathcal{G}\left(0, \alpha^2 \sigma_z^2 / 2\right), \quad (110)$$

and, hence, $s(\delta)$ is asymptotically normal as $\delta \rightarrow 0$.

6. *Large Deviations*: Assume that z_m is non-deterministic and has LMGF finite everywhere:

$$\Lambda_z(t) = \log \mathbb{E} \left[e^{z_m t} \right] < +\infty, \quad \forall t \in \mathbb{R}. \quad (111)$$

Let $\Lambda_{\alpha z}(t) = \Lambda_z(\alpha t)$ be the LMGF of the scaled variable αz_m , where α is defined in (103). Denoting by $\Lambda_\delta(t)$ the LMGF of $s(\delta)$, we have that:

$$\lim_{\delta \rightarrow 0} \delta \Lambda_\delta(t/\delta) = \phi(t) = \int_0^t \frac{\Lambda_{\alpha z}(\tau)}{\tau} d\tau. \quad (112)$$

Then the following Large Deviations Principle (LDP) holds for any measurable set \mathcal{S} (the infimum over an empty set is taken as $+\infty$):

$$\liminf_{\delta \rightarrow 0} \delta \log \mathbb{P}[s(\delta) \in \mathcal{S}] \geq - \inf_{\gamma \in \mathcal{S}^o} \phi^*(\gamma), \quad (113)$$

$$\limsup_{\delta \rightarrow 0} \delta \log \mathbb{P}[s(\delta) \in \mathcal{S}] \leq - \inf_{\gamma \in \overline{\mathcal{S}}} \phi^*(\gamma), \quad (114)$$

where

$$\phi^*(\gamma) = \sup_{t \in \mathbb{R}} [\gamma t - \phi(t)] \quad (115)$$

is the Fenchel-Legendre transform of $\phi(t)$ [28], [29]. The function $\phi^*(\gamma)$ (which is allowed to be an extended real number) is usually called rate function [28], [29] and has the following properties.

- Let z_- and z_+ be the boundaries of the support of z_m , and let $\mathcal{D} = \{\gamma \in \mathbb{R} : \phi^*(\gamma) < +\infty\}$. Then \mathcal{D} is given by the following open interval:

$$\mathcal{D} = (\alpha z_-, \alpha z_+). \quad (116)$$

- The function $\phi^*(\gamma)$ is smooth and strictly convex on \mathcal{D} , and diverges to $+\infty$ at the boundaries of \mathcal{D} . In particular, if a boundary is finite, the rate function is equal to $+\infty$ at that boundary.
- $\phi^*(\gamma) \geq 0$, with equality if, and only if, $\gamma = \alpha m_z$.

A typical shape of the rate function is illustrated in Fig. 14. Exploiting the aforementioned regularity properties of $\phi^*(\gamma)$,

from (113)–(114) we have in particular that:

$$\lim_{\delta \rightarrow 0} \delta \log \mathbb{P}[s(\delta) \geq \gamma] = -\phi^*(\gamma), \quad \forall \gamma \geq \alpha m_z, \quad (117)$$

$$\lim_{\delta \rightarrow 0} \delta \log \mathbb{P}[s(\delta) \leq \gamma] = -\phi^*(\gamma), \quad \forall \gamma \leq \alpha m_z. \quad (118)$$

Proof: We prove sequentially the six parts of the lemma.

Part 1: In view of (101), the following series of (absolute) expectations is convergent:

$$\begin{aligned} \delta \sum_{m=0}^{\infty} (1-\delta)^m \alpha_m \mathbb{E} \left[|z_m| \right] &= m_z^{\text{abs}} \delta \sum_{m=0}^{\infty} (1-\delta)^m \alpha_m \\ &\leq m_z^{\text{abs}} \delta \sum_{m=0}^{\infty} (1-\delta)^m \\ &= m_z^{\text{abs}} < +\infty. \end{aligned} \quad (119)$$

In view of [34][Lemma 3.6'], convergence of the series of absolute first moments implies that the random series $s^{\text{abs}}(\delta)$ is almost-surely finite, which in turn implies that so is $s(\delta)$, and part 1 is proved.

Part 2: Since the series of (absolute) expectations is convergent, so is the series of expectations:

$$\sum_{m=0}^{\infty} (1-\delta)^m \alpha_m \mathbb{E}[z_m] = m_z \sum_{m=0}^{\infty} (1-\delta)^m \alpha_m. \quad (120)$$

On the other hand, by triangle inequality we have the following upper bound:

$$|s_i(\delta)| \leq \delta \sum_{m=0}^i (1-\delta)^m \alpha_m |z_m| \leq s^{\text{abs}}(\delta). \quad (121)$$

Now we observe that $s^{\text{abs}}(\delta)$ is a proper random variable in view of part 1. Furthermore, it is an integrable random variable from Beppo Levi's monotone convergence theorem [35][Th. 1.5.7, p. 27], thanks to the convergence of absolute expectations in (120).

We conclude that the random sequence $s_i(\delta)$ is upper bounded by an integrable random variable. Therefore, the dominated convergence theorem [35][Th. 1.5.8, p. 27] implies that the expectation of the almost-sure limit $s(\delta)$ is equal to the convergent series of expectations, and the first equality in (106) follows. Moreover, we can write:

$$\begin{aligned} \delta \sum_{m=0}^{\infty} (1-\delta)^m \alpha_m &= \delta \sum_{m=0}^{\infty} (1-\delta)^m (\alpha_m - \alpha) \\ &\quad + \underbrace{\alpha \delta \sum_{m=0}^{\infty} (1-\delta)^m}_{=1}. \end{aligned} \quad (122)$$

In view of (103), the absolute value of the first summation on the RHS in (122) is dominated by:

$$\kappa \delta \sum_{m=0}^{\infty} (\beta(1-\delta))^m = \frac{\kappa \delta}{1-\beta(1-\delta)} = O(\delta). \quad (123)$$

We conclude from (120), (122) and (123) that the second equality in (106) holds.

Part 3: Let

$$\zeta_m \triangleq \delta(1-\delta)^m \alpha_m, \quad (124)$$

and consider the following centered variables:

$$\tilde{\mathbf{s}}(\delta) = \mathbf{s}(\delta) - \mathbb{E}[\mathbf{s}(\delta)], \quad \tilde{\mathbf{z}}_m = \mathbf{z}_m - \mathbb{E}[\mathbf{z}_m]. \quad (125)$$

In view of parts 1 and 2, the centered partial sums:

$$\mathbf{s}_i(\delta) - \mathbb{E}[\mathbf{s}_i(\delta)] = \sum_{m=0}^i \zeta_m \tilde{\mathbf{z}}_m \quad (126)$$

converge in distribution to $\tilde{\mathbf{s}}(\delta)$ as $i \rightarrow \infty$. By Lévy's continuity theorem, the corresponding characteristic functions must converge [36][Th. 2, p. 431]. Since the \mathbf{z}_m 's are i.i.d. we can write:

$$\varphi_{\tilde{\mathbf{s}}}(t) \triangleq \mathbb{E} \left[e^{j\tilde{\mathbf{s}}(\delta)t} \right] = \prod_{m=0}^{\infty} \varphi_{\tilde{\mathbf{z}}_m}(\zeta_m t), \quad (127)$$

where $j = \sqrt{-1}$. We want to show that $\tilde{\mathbf{s}}(\delta)$ converges in probability to 0 as $\delta \rightarrow 0$. In view of Lévy's continuity Theorem this is tantamount to showing that $\varphi_{\tilde{\mathbf{s}}}(t)$ converges to 1 as $\delta \rightarrow 0$. Using (127) we can write⁵:

$$|\varphi_{\tilde{\mathbf{s}}}(t) - 1| \leq \sum_{m=0}^{\infty} |\varphi_{\tilde{\mathbf{z}}_m}(\zeta_m t) - 1|. \quad (129)$$

Consider, without loss of generality, a positive t . Since the random variables $\tilde{\mathbf{z}}_m$ have finite expectation, the first derivative of the characteristic function, $\varphi'_{\tilde{\mathbf{z}}_m}(t)$, is a continuous function, and by the mean-value theorem we can write (since in particular $\mathbb{E}[\tilde{\mathbf{z}}_m] = 0$):

$$\varphi_{\tilde{\mathbf{z}}_m}(\zeta_m t) = 1 + \zeta_m t \varphi'_{\tilde{\mathbf{z}}_m}(t_m), \text{ for some } t_m \in (0, \zeta_m t). \quad (130)$$

Accordingly we can write:

$$|\varphi_{\tilde{\mathbf{z}}_m}(\zeta_m t) - 1| \leq \zeta_m |t| \max_{\tau \in [0, \delta t]} |\varphi'_{\tilde{\mathbf{z}}_m}(\tau)|, \quad (131)$$

where the latter inequality follows from the fact that $\zeta_m \leq \delta$, see (124). Applying (131) to (129) we get:

$$|\varphi_{\tilde{\mathbf{s}}}(t) - 1| \leq |t| \max_{\tau \in [0, \delta t]} |\varphi'_{\tilde{\mathbf{z}}_m}(\tau)| \underbrace{\sum_{m=0}^{\infty} \zeta_m}_{\leq 1}. \quad (132)$$

On the other hand, since $\varphi'_{\tilde{\mathbf{z}}_m}(0) = \mathbb{E}[\tilde{\mathbf{z}}_m] = 0$, from the continuity of $\varphi'_{\tilde{\mathbf{z}}_m}(t)$ it follows that:

$$\lim_{\delta \rightarrow 0} \max_{\tau \in [0, \delta t]} |\varphi'_{\tilde{\mathbf{z}}_m}(\tau)| = 0, \quad (133)$$

which proves that $\mathbf{s}(\delta)$ converges to $\mathbb{E}[\mathbf{s}(\delta)]$ in probability as $\delta \rightarrow 0$. The claim in (107) then follows from (106).

Part 4: Since the variables \mathbf{z}_m have common finite variance σ_z^2 and are independent, it is immediate to see that:

$$\lim_{i \rightarrow \infty} \text{VAR}[\mathbf{s}_i(\delta)] = \sigma_z^2 \delta^2 \sum_{m=0}^{\infty} (1 - \delta)^{2m} \alpha_m^2 < \infty. \quad (134)$$

⁵The following inequality is known for complex numbers x_m, y_m , with $|x_m| \leq 1$ and $|y_m| \leq 1$ [36]:

$$\left| \prod_{m=0}^i x_m - \prod_{m=0}^i y_m \right| \leq \sum_{m=0}^i |x_m - y_m|, \quad (128)$$

Consider now the squared and centered variables:

$$\begin{aligned} & \left(\mathbf{s}_i(\delta) - \mathbb{E}[\mathbf{s}_i(\delta)] \right)^2 \\ &= \delta^2 \left(\sum_{m=0}^i (1 - \delta)^m \alpha_m (\mathbf{z}_m - \mathbf{m}_z) \right)^2. \end{aligned} \quad (135)$$

In view of parts 1 and 2 the quantity on the LHS converges almost surely, as $i \rightarrow \infty$, to:

$$\left(\mathbf{s}(\delta) - \mathbb{E}[\mathbf{s}(\delta)] \right)^2. \quad (136)$$

Given the convergence of the variance of the partial sums in (134), by Fatou's lemma we conclude that [35][Th. 1.5.5, p. 26]:

$$\text{VAR}[\mathbf{s}(\delta)] \leq \lim_{i \rightarrow \infty} \text{VAR}[\mathbf{s}_i(\delta)], \quad (137)$$

i.e., the limiting variable $\mathbf{s}(\delta)$ has finite variance. But since the limiting variable $\mathbf{s}(\delta)$ can be written as:

$$\mathbf{s}(\delta) = \mathbf{s}_i(\delta) + \delta \sum_{m=i+1}^{\infty} (1 - \delta)^m \alpha_m \mathbf{z}_m, \quad (138)$$

with the two quantities on the RHS being statistically independent, the variance of $\mathbf{s}(\delta)$ cannot be smaller than the variance of $\mathbf{s}_i(\delta)$ for all i , implying that:

$$\text{VAR}[\mathbf{s}(\delta)] \geq \lim_{i \rightarrow \infty} \text{VAR}[\mathbf{s}_i(\delta)]. \quad (139)$$

Combining (137) with (139) we see that the variance of the almost-sure limit $\mathbf{s}(\delta)$ is equal to the convergent series of variances, which is the first equality in (109).

In order to prove the second equality in (109) we write:

$$\begin{aligned} & \text{VAR} \left[\delta \sum_{m=0}^{\infty} (1 - \delta)^m \alpha_m \mathbf{z}_m \right] \\ &= \sigma_z^2 \delta^2 \sum_{m=0}^{\infty} (1 - \delta)^{2m} \alpha_m^2 \\ &= \sigma_z^2 \delta^2 \sum_{m=0}^{\infty} (1 - \delta)^{2m} (\alpha_m^2 - \alpha^2) \\ &+ \alpha^2 \sigma_z^2 \delta^2 \sum_{m=0}^{\infty} (1 - \delta)^{2m}. \end{aligned} \quad (140)$$

Reasoning as done to prove part 2, we can easily show that the first summation on the RHS in (140) is $O(\delta^2)$. The second summation is instead equal to:

$$\frac{\alpha^2 \sigma_z^2 \delta^2}{1 - (1 - \delta)^2} = \frac{\alpha^2 \sigma_z^2 \delta}{2 - \delta}, \quad (141)$$

and the second equality in (109) follows.

Part 5: Let

$$\sigma_{\text{lim}}^2 \triangleq \frac{\alpha^2 \sigma_z^2}{2}. \quad (142)$$

The claim in (110) is equivalent to prove that the random variable $\frac{\mathbf{s}(\delta) - \mathbf{m}_z}{\sqrt{\delta} \sigma_{\text{lim}}}$ converges in distribution to a standard Gaussian. On the other hand, we have that:

$$\frac{\mathbf{s}(\delta) - \mathbf{m}_z}{\sqrt{\delta} \sigma_{\text{lim}}} = \frac{\mathbf{s}(\delta) - \mathbb{E}[\mathbf{s}(\delta)]}{\sqrt{\delta} \sigma_{\text{lim}}} + \frac{\mathbb{E}[\mathbf{s}(\delta)] - \mathbf{m}_z}{\sqrt{\delta} \sigma_{\text{lim}}}. \quad (143)$$

Since the second term in (143) converges to zero in view of (106), from Slutsky's theorem [26][Th. 1.11, p. 60] it suffices to show that the random variable $\frac{s(\delta) - \mathbb{E}[s(\delta)]}{\sqrt{\delta}\sigma_{\text{lim}}}$ converges in distribution to a standard Gaussian. To this end, we start by introducing, with slight abuse of notation w.r.t. (124) and (125), the quantities:

$$\zeta_m \triangleq \frac{\sqrt{2\delta}(1-\delta)^m \alpha_m}{\alpha}, \quad (144)$$

and:

$$\tilde{s}(\delta) = \frac{s(\delta) - \mathbb{E}[s(\delta)]}{\sqrt{\delta}\sigma_{\text{lim}}}, \quad \tilde{z}_m = \frac{z_m - \mathbb{E}[z_m]}{\sigma_z}. \quad (145)$$

We notice that \tilde{z}_m has zero mean and unit variance.

We will now show that $\tilde{s}(\delta)$ converges in distribution to a standard Gaussian. In view of Lévy's continuity theorem, this claim is equivalent to the convergence, as $\delta \rightarrow 0$, of the characteristic function of $\tilde{s}(\delta)$ to the characteristic function $e^{-\frac{t^2}{2}}$. From (105), (142), (144) and (145) we see that:

$$\tilde{s}(\delta) = \sum_{m=0}^{\infty} \zeta_m \tilde{z}_m. \quad (146)$$

Reasoning as done to compute (127), the characteristic function of $\tilde{s}(\delta)$ in (145) can be written as:

$$\varphi_{\tilde{s}}(t) = \prod_{m=0}^{\infty} \varphi_{\tilde{z}}(\zeta_m t). \quad (147)$$

Using the triangle inequality for complex numbers we can write:

$$\left| \varphi_{\tilde{s}}(t) - e^{-\frac{t^2}{2}} \right| \leq \left| \varphi_{\tilde{s}}(t) - e^{-\frac{\sum_{m=0}^{\infty} \zeta_m^2 t^2}{2}} \right| + \left| e^{-\frac{\sum_{m=0}^{\infty} \zeta_m^2 t^2}{2}} - e^{-\frac{t^2}{2}} \right|. \quad (148)$$

Now, that the second term on the RHS of (148) converges to zero follows from part 4), since from (109) and the definition of ζ_m in (144) we conclude that:

$$\lim_{\delta \rightarrow 0} \sum_{m=0}^{\infty} \zeta_m^2 = 1. \quad (149)$$

Let us now focus on the first term on the RHS of (148). Since the characteristic functions have magnitude not greater than 1, in view of (128) and (147) we can write:

$$\begin{aligned} & \left| \varphi_{\tilde{s}}(t) - e^{-\frac{\sum_{m=0}^{\infty} \zeta_m^2 t^2}{2}} \right| \\ & \leq \sum_{m=0}^{\infty} \left| \varphi_{\tilde{z}}(\zeta_m t) - e^{-\frac{\zeta_m^2 t^2}{2}} \right| \\ & \leq \sum_{m=0}^{\infty} \left| \varphi_{\tilde{z}}(\zeta_m t) - 1 + \frac{\zeta_m^2 t^2}{2} \right| \\ & + \sum_{m=0}^{\infty} \left| e^{-\frac{\zeta_m^2 t^2}{2}} - 1 + \frac{\zeta_m^2 t^2}{2} \right|, \end{aligned} \quad (150)$$

where in the latter step we applied the triangle inequality. Now, the last term in (150) converges to zero since for any positive

s we have $|e^{-s} - 1 + s| \leq s^2/2$, and since it is immediate to show that (see the proof in [27]):

$$\lim_{\delta \rightarrow 0} \sum_{m=0}^{\infty} \zeta_m^4 = 0. \quad (151)$$

On the other hand, using [35][Lemma 3.3.19, p. 134] we can write, for an arbitrarily small $\epsilon > 0$:

$$\begin{aligned} & \left| e^{j\tilde{z}_m \zeta_m t} - 1 - j\tilde{z}_m \zeta_m t + \frac{1}{2}\tilde{z}_m^2 \zeta_m^2 t^2 \right| \\ & \leq \mathbb{I}\left\{|\tilde{z}_m| \zeta_m \leq \epsilon\right\} \frac{|\tilde{z}_m \zeta_m t|^3}{6} \\ & + \mathbb{I}\left\{|\tilde{z}_m| \zeta_m > \epsilon\right\} (\tilde{z}_m \zeta_m t)^2 \\ & \leq \epsilon \tilde{z}_m^2 \zeta_m^2 \frac{|t|^3}{6} + \tilde{z}_m^2 \mathbb{I}\left\{|\tilde{z}_m| \zeta_m > \epsilon\right\} \zeta_m^2 t^2 \\ & \leq \epsilon \tilde{z}_m^2 \zeta_m^2 \frac{|t|^3}{6} + \tilde{z}_m^2 \mathbb{I}\left\{|\tilde{z}_m| > \epsilon\alpha/\sqrt{2\delta}\right\} \zeta_m^2 t^2, \end{aligned} \quad (152)$$

where $\mathbb{I}\{\mathcal{E}\}$ is the indicator of event \mathcal{E} , and the last inequality follows because $\zeta_m \leq \sqrt{2\delta}/\alpha$ — see (144). Let now:

$$g(\delta) = \mathbb{E}\left[\tilde{z}_m^2 \mathbb{I}\left\{\tilde{z}_m^2 > \epsilon\alpha/\sqrt{2\delta}\right\}\right]. \quad (153)$$

Owing to identical distribution of \tilde{z}_m across index m , the function $g(\delta)$ does not depend on m . Since \tilde{z}_m has finite variance, we have that $g(\delta) \rightarrow 0$ as $\delta \rightarrow 0$. In view of (152), recalling that the magnitude of the expectation is upper bounded by the expectation of the magnitude, and that \tilde{z}_m has zero mean and unit variance, we have that:

$$\left| \varphi_{\tilde{z}}(\zeta_m t) - 1 + \frac{\zeta_m^2 t^2}{2} \right| \leq \sum_{m=0}^{\infty} \zeta_m^2 \left(\epsilon \frac{|t|^3}{6} + t^2 g(\delta) \right), \quad (154)$$

and, hence,

$$\limsup_{\delta \rightarrow 0} \left| \varphi_{\tilde{z}}(\zeta_m t) - 1 + \frac{\zeta_m^2 t^2}{2} \right| \leq \epsilon \frac{|t|^3}{6}, \quad (155)$$

finally implying, due to the arbitrariness of ϵ , that $\varphi_{\tilde{s}}(t)$ converges to $e^{-t^2/2}$ as $\delta \rightarrow 0$. We have therefore shown that $\tilde{s}(\delta)$ in (145) converges to a standard Gaussian as $\delta \rightarrow 0$, and this completes the proof of part 5.

Part 6: The convergence in (112) can be proved as done in [27, Appendix C]. Then the convergence in (112) implies the LDP in (113)–(114) in view of the Gärtner-Ellis theorem [28][Th. 2.3.6, p. 44], [29][Th. V.6, p. 54].

Next we focus on the regularity properties of the Fenchel-Legendre transform $\phi^*(\gamma)$. Following the development used in [27, Appendix C], we can prove that \mathcal{D}° is an interval, that $\phi^*(\gamma)$ is smooth and strictly convex for $\gamma \in \mathcal{D}^\circ$, and that $\phi^*(\gamma) \geq 0$ with equality if, and only if, $\gamma = \alpha m_z$.

Thus, it remains to characterize the boundaries of \mathcal{D}° and the behavior of the rate function at these boundaries. To this end, it is sufficient to prove the claim with $\alpha = 1$ and for the right boundary, since the proof for other values of α and for the left boundary is simply obtained using the scaling and reflection properties of the LMGF [28], [29].

Now, since it has been shown in [27, Appendix C] that the right boundary of \mathcal{D}° is equal to $\lim_{t \rightarrow \infty} \Lambda_z(t)/t$, we must now prove that this limit equals z_+ (recall that we are working

with $\alpha = 1$). We start by noticing that, letting $z_- < \bar{z} < z_+$, the LMGF $\Lambda_z(t)$ can be written as:

$$\Lambda_z(t) = \log \left(\mathbb{E} [\mathbb{I}\{z_m \leq \bar{z}\} e^{z_m t}] + \mathbb{E} [\mathbb{I}\{z_m > \bar{z}\} e^{z_m t}] \right). \quad (156)$$

From (156) we get, for all $t > 0$:

$$\frac{\Lambda_z(t)}{t} \geq \frac{\log \left(e^{\bar{z}t} \mathbb{E} [\mathbb{I}\{z_m > \bar{z}\}] \right)}{t} = \bar{z} + \frac{\log q}{t}, \quad (157)$$

where we set $q = \mathbb{P}[z_m > \bar{z}]$. We remark that $0 < q < 1$ since \bar{z} is internal to the support of z_m . From (157) we get:

$$\liminf_{t \rightarrow \infty} \frac{\Lambda_z(t)}{t} \geq \bar{z}. \quad (158)$$

If $z_+ = +\infty$ the result is proved due to arbitrariness of \bar{z} . If $z_+ < +\infty$, we can choose $\bar{z} = z_+ - \epsilon$, and conclude that the limit inferior in (158) is equal to z_+ . The fact that the corresponding limit superior is equal to z_+ follows by observing that, in view of (156), for all $t > 0$ the quantity $\Lambda_z(t)/t$ is upper bounded by z_+ .

Finally, we characterize the behavior of the rate function at the boundaries of \mathcal{D}° . We focus again on the right boundary z_+ . When $z_+ = +\infty$, it suffices to notice that the rate function $\phi^*(\gamma)$ is strictly convex in \mathcal{D}° and is strictly increasing for $\gamma > m_z$ (see Fig. 14) to conclude that the rate function diverges to $+\infty$ as $\gamma \rightarrow z_+$.

We move on to examine the case $z_+ < +\infty$. Exploiting (156) we can write, for all $t > 0$:

$$\begin{aligned} \Lambda_z(t) &\leq \log \left((1-q)e^{\bar{z}t} + qe^{z_+t} \right) \\ &= z_+t + \log \left((1-q)e^{-(z_+-\bar{z})t} + q \right). \end{aligned} \quad (159)$$

Since $z_+ > \bar{z}$, for any $\epsilon > 0$ there exists $t_\epsilon > 0$ such that:

$$(1-q)e^{-(z_+-\bar{z})t} \leq \epsilon q, \quad \text{for all } t \geq t_\epsilon, \quad (160)$$

implying, in view of (159):

$$\Lambda_z(t) \leq z_+t + \log((1+\epsilon)q), \quad \text{for all } t \geq t_\epsilon. \quad (161)$$

Using (161) in (112) we can thus write:

$$\begin{aligned} \phi(t) &= \int_0^t \frac{\Lambda_z(\tau)}{\tau} d\tau = \int_0^{t_\epsilon} \frac{\Lambda_z(\tau)}{\tau} d\tau + \int_{t_\epsilon}^t \frac{\Lambda_z(\tau)}{\tau} d\tau \\ &\leq \phi(t_\epsilon) + z_+(t-t_\epsilon) + \int_{t_\epsilon}^t \frac{\log((1+\epsilon)q)}{\tau} d\tau \\ &= \phi(t_\epsilon) + z_+(t-t_\epsilon) + \log((1+\epsilon)q) \log \frac{t}{t_\epsilon}. \end{aligned} \quad (162)$$

Plugging the latter inequality in (115) we get:

$$\begin{aligned} \phi^*(z_+) &\geq \sup_{t \geq t_\epsilon} [z_+t - \phi(t)] \geq -\phi(t_\epsilon) + z_+t_\epsilon \\ &\quad + \log \frac{1}{(1+\epsilon)q} \sup_{t \geq t_\epsilon} \log \frac{t}{t_\epsilon} = +\infty, \end{aligned} \quad (163)$$

where we have chosen ϵ so small to ensure that $(1+\epsilon)q < 1$. Finally, in view of (115) we can write, for a generic $t \in \mathbb{R}$:

$$\lim_{\gamma \rightarrow z_+} \phi^*(\gamma) \geq \lim_{\gamma \rightarrow z_+} [\gamma t - \phi(t)] = [z_+t - \phi(t)], \quad (164)$$

and from (163) we conclude that $\phi^*(\gamma) \rightarrow +\infty$ as $\gamma \rightarrow z_+$. ■

APPENDIX B

Proof of Theorem 1: We are interested in characterizing, for each agent k , the *joint* behavior of the random variables $\widehat{\lambda}_{k,i}^{(\delta)}(\theta)$ for all values of $\theta \neq \theta_0$. To this end, it is useful to consider the $(H-1) \times 1$ vector $\widehat{\lambda}_{k,i}^{(\delta)}$ similarly defined as the vector in (15). We also introduce, for a fixed time epoch i , the $N \times (H-1)$ data matrix \mathbf{X}_i , whose entries, for $\ell = 1, 2, \dots, N$ and $\theta \neq \theta_0$, are:

$$[\mathbf{X}_i]_{\ell\theta} = \mathbf{x}_{\ell,i}(\theta). \quad (165)$$

In light of (27) we can write:

$$\widehat{\lambda}_{k,i}^{(\delta)} = f_{k,i}^{(\delta)}(\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_i), \quad (166)$$

to highlight that the random vector $\widehat{\lambda}_{k,i}^{(\delta)}$ is a certain function $f_{k,i}^{(\delta)}$ of the data matrices $\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_i$. Since the data are i.i.d. over time, reversing the order of the data matrices in (166) does not change the distribution of the resulting random vector, i.e.:

$$\widetilde{\lambda}_{k,i}^{(\delta)} = f_{k,i}^{(\delta)}(\mathbf{X}_i, \mathbf{X}_{i-1}, \dots, \mathbf{X}_1) \stackrel{d}{=} \widehat{\lambda}_{k,i}^{(\delta)}, \quad (167)$$

where $\stackrel{d}{=}$ denotes equality in distribution. Considering this reversed order of the data matrices in (27) and exchanging the order of summation we obtain:

$$\widetilde{\lambda}_{k,i}^{(\delta)}(\theta) = \sum_{\ell=1}^N \delta \sum_{m=0}^{i-1} (1-\delta)^m [A^{m+1}]_{\ell k} \mathbf{x}_{\ell, m+1}(\theta). \quad (168)$$

From part 1) of Lemma 1 in the Appendix, each of the N inner partial sums (scaled by δ) converges *almost surely*. In fact, the random variables $\mathbf{x}_{\ell, m+1}(\theta)$ have finite first moment in view of Assumption 1, and the weights $[A^{m+1}]_{\ell k}$ fulfill condition (103) in view of Property 1. It makes thus sense to define a proper random variable as the (almost-surely convergent) value of the random series in (168), which corresponds to (32). This in turn implies the following almost-sure convergence, as $i \rightarrow \infty$, of the *vector* with reversed ordering, $\widetilde{\lambda}_{k,i}^{(\delta)}$, to the limiting random vector $\widetilde{\lambda}_k^{(\delta)}$. In view of (167), this almost-sure convergence implies the convergence *in distribution* of the original (i.e., with correct ordering of the data matrices \mathbf{X}_i) vector $\widehat{\lambda}_{k,i}$, finally yielding the claim of the theorem. ■

APPENDIX C

Proof of Theorem 2: We start by proving (38). Examining (32) we see that each one of the N inner series matches the conditions in Lemma 1, part 3, implying that the ℓ -th inner series converges in probability, as $\delta \rightarrow 0$, to the expected value $\pi_\ell \mathbb{E}[\mathbf{x}_{\ell, m+1}(\theta)] = \pi_\ell d_\ell(\theta)$. As a result, $\widetilde{\lambda}_k^{(\delta)}(\theta)$ converges in probability to $m_{\text{ave}}(\theta)$, which implies, for any $\epsilon > 0$:

$$\lim_{\delta \rightarrow 0} \mathbb{P} \left[\widetilde{\lambda}_k^{(\delta)}(\theta) < m_{\text{ave}}(\theta) - \epsilon \right] = 0. \quad (169)$$

Since under Assumption 3 the quantity $m_{\text{ave}}(\theta)$ is strictly positive, we conclude that:

$$\lim_{\delta \rightarrow 0} \mathbb{P} \left[\widetilde{\lambda}_k^{(\delta)}(\theta) \leq 0 \right] = 0, \quad (170)$$

which, by application of the union bound, in light of (18) gives:

$$\begin{aligned} p_k^{(\delta)} &= \mathbb{P} \left[\exists \theta \neq \theta_0 : \tilde{\lambda}_k^{(\delta)}(\theta) \leq 0 \right] \\ &\leq \sum_{\theta \neq \theta_0} \mathbb{P} \left[\tilde{\lambda}_k^{(\delta)}(\theta) \leq 0 \right] \xrightarrow{\delta \rightarrow 0} 0, \end{aligned} \quad (171)$$

and the claim of the theorem is proved. \blacksquare

APPENDIX D

Proof of Theorem 3: In the following we will refer to the elements $\theta_1, \theta_2, \dots, \theta_{H-1}$ in the set $\Theta \setminus \theta_0$ — see (16). Consider a zero-mean Gaussian random vector:

$$\mathbf{g} = [\mathbf{g}(\theta_1), \mathbf{g}(\theta_2), \dots, \mathbf{g}(\theta_{H-1})]^\top, \quad (172)$$

with covariance matrix equal to $\mathbf{C}_{\text{ave}}/2$. We recall that the (θ, θ') -th entry of \mathbf{C}_{ave} is the covariance $c_{\text{ave}}(\theta, \theta')$ defined in (41). What we want to show is that the random vector:

$$\frac{\tilde{\lambda}_k^{(\delta)} - m_{\text{ave}}}{\sqrt{\delta}} \quad (173)$$

converges in distribution to \mathbf{g} .

When dealing with convergence in distribution of random vectors, the standard path is to reduce the vector problem to a scalar problem through the following argument. In view of Lévy's continuity theorem for random vectors, convergence in distribution takes place if, and only if, convergence of the pertinent (multivariate) characteristic functions takes place [26]. This implies that⁶ our claim will be proved if we show that, for any sequence of real numbers $t(\theta_1), t(\theta_2), \dots, t(\theta_{H-1})$:

$$\sum_{\theta \neq \theta_0} t(\theta) \frac{\tilde{\lambda}_k^{(\delta)}(\theta) - m_{\text{ave}}(\theta)}{\sqrt{\delta}} \xrightarrow{\delta \rightarrow 0} \sum_{\theta \neq \theta_0} t(\theta) \mathbf{g}(\theta). \quad (174)$$

Obviously, the linear combination on the RHS in (174) is a Gaussian random variable with zero mean and with variance:

$$\text{VAR} \left[\sum_{\theta \neq \theta_0} t(\theta) \mathbf{g}(\theta) \right] = \sum_{\theta \neq \theta_0} \sum_{\theta' \neq \theta_0} t(\theta) t(\theta') \frac{c_{\text{ave}}(\theta, \theta')}{2}. \quad (175)$$

Let us now examine the LHS in (174). Using (168) we get:

$$\begin{aligned} &\sum_{\theta \neq \theta_0} t(\theta) \tilde{\lambda}_k^{(\delta)}(\theta) \\ &= \sum_{\ell=1}^N \delta \sum_{m=0}^{\infty} (1-\delta)^m [A^{m+1}]_{\ell k} \sum_{\theta \neq \theta_0} t(\theta) \mathbf{x}_{\ell, m+1}(\theta), \end{aligned} \quad (176)$$

whereas using (21) we have:

$$\sum_{\theta \neq \theta_0} t(\theta) m_{\text{ave}}(\theta) = \sum_{\ell=1}^N \pi_\ell \sum_{\theta \neq \theta_0} t(\theta) d_\ell(\theta). \quad (177)$$

⁶This corollary of Lévy's continuity theorem is also known as Cramér-Wold device or theorem [26][Th. 1.9, p. 56].

Let us now set:

$$\mathbf{z}_m^{(\ell)} \triangleq \sum_{\theta \neq \theta_0} t(\theta) \mathbf{x}_{\ell, m+1}(\theta), \quad (178)$$

$$\alpha_m^{(\ell)} \triangleq [A^{m+1}]_{\ell k}, \quad (179)$$

$$\mathbf{s}^{(\ell)}(\delta) \triangleq \delta \sum_{m=0}^{\infty} (1-\delta)^m \alpha_m^{(\ell)} \mathbf{z}_m^{(\ell)}. \quad (180)$$

We observe that:

$$\mathbb{E} \left[\mathbf{z}_m^{(\ell)} \right] = \sum_{\theta \neq \theta_0} t(\theta) d_\ell(\theta), \quad (181)$$

$$\text{VAR} \left[\mathbf{z}_m^{(\ell)} \right] = \sum_{\theta \neq \theta_0} \sum_{\theta' \neq \theta_0} t(\theta) t(\theta') \rho_\ell(\theta, \theta'). \quad (182)$$

Exploiting Eqs. (178)–(181), the LHS in (174) can be cast in the form:

$$\sum_{\ell=1}^N \frac{\mathbf{s}^{(\ell)}(\delta) - \mathbb{E} \left[\mathbf{z}_m^{(\ell)} \right]}{\sqrt{\delta}}. \quad (183)$$

We see from Eqs. (178)–(180) that the random variables $\mathbf{s}^{(\ell)}(\delta)$ match the structure of the random series used in Lemma 1. We now verify that $\mathbf{s}^{(\ell)}(\delta)$ fulfills the conditions of part 5 in Lemma 1, for every $\ell = 1, 2, \dots, N$. First we note that $\mathbf{z}_m^{(\ell)}$ has finite variance since it is a linear combination of random variables that have finite variance. Second we see that condition (103) is verified in view of Property 1. We conclude then from part 5 of Lemma 1 that the following convergence in distribution holds:

$$\frac{\mathbf{s}^{(\ell)}(\delta) - \mathbb{E} \left[\mathbf{z}_m^{(\ell)} \right]}{\sqrt{\delta}} \xrightarrow{\delta \rightarrow 0} \mathcal{G} \left(0, \frac{\pi_\ell^2}{2} \text{VAR} \left[\mathbf{z}_m^{(\ell)} \right] \right). \quad (184)$$

Since the data are independent across agents, we have that the random variables $\mathbf{s}^{(\ell)}(\delta)$ are independent across index ℓ . For this reason, and in view of (184), we conclude that the LHS in (174) is asymptotically normal, with zero mean and with variance given by:

$$\begin{aligned} &\frac{\pi_\ell^2}{2} \sum_{\ell=1}^N \text{VAR} \left[\mathbf{z}_m^{(\ell)} \right] \\ &= \sum_{\theta \neq \theta_0} \sum_{\theta' \neq \theta_0} t(\theta) t(\theta') \sum_{\ell=1}^N \frac{\pi_\ell^2}{2} \rho_\ell(\theta, \theta') \\ &= \sum_{\theta \neq \theta_0} \sum_{\theta' \neq \theta_0} t(\theta) t(\theta') \frac{c_{\text{ave}}(\theta, \theta')}{2}, \end{aligned} \quad (185)$$

where we have used (182). Since the RHS in (185) coincides with the variance in (175), the proof is complete. \blacksquare

APPENDIX E

Proof of Theorem 4: In light of (18), the error probability of *not* choosing θ_0 can be bounded as follows (with the lower bound holding for every $\theta \neq \theta_0$):

$$\mathbb{P} \left[\tilde{\lambda}_{k,i}^{(\delta)}(\theta) \leq 0 \right] \leq p_{k,i}^{(\delta)} \leq \sum_{\theta \neq \theta_0} \mathbb{P} \left[\tilde{\lambda}_{k,i}^{(\delta)}(\theta) \leq 0 \right], \quad (186)$$

where the upper bound is the union bound. At the steady state, Eq. (186) implies:

$$\mathbb{P}\left[\tilde{\lambda}_k^{(\delta)}(\theta) \leq 0\right] \leq p_k^{(\delta)} \leq \sum_{\theta \neq \theta_0} \mathbb{P}\left[\tilde{\lambda}_k^{(\delta)}(\theta) \leq 0\right]. \quad (187)$$

One key point to prove the claim of the theorem is the exponential characterization of the probability $\mathbb{P}\left[\tilde{\lambda}_k^{(\delta)}(\theta) \leq 0\right]$. Preliminarily, let us set:

$$\mathbf{z}_m^{(\ell)} \triangleq \mathbf{x}_{\ell, m+1}(\theta), \quad (188)$$

$$\alpha_m^{(\ell)} \triangleq [A^{m+1}]_{\ell k}, \quad (189)$$

$$\mathbf{s}^{(\ell)}(\delta) \triangleq \delta \sum_{m=0}^{\infty} (1-\delta)^m \alpha_m^{(\ell)} \mathbf{z}_m^{(\ell)}, \quad (190)$$

which yields:

$$\tilde{\lambda}_k^{(\delta)}(\theta) = \sum_{\ell=1}^N \mathbf{s}^{(\ell)}(\delta). \quad (191)$$

Recall that the log-likelihood ratios $\mathbf{x}_{\ell, m}(\theta)$ are assumed to be independent across agents (i.e., across ℓ). Thus $\mathbf{s}^{(\ell)}(\delta)$ are also independent random variables. Now, part 6 of Lemma 1 would provide the required exponential characterization for the individual variable $\mathbf{s}^{(\ell)}(\delta)$. We need instead the characterization for $\tilde{\lambda}_k^{(\delta)}(\theta)$, which is the sum of the (independent) variables $\mathbf{s}^{(\ell)}(\delta)$. Let us elaborate on this aspect. The starting point to prove part 6 in Lemma 1 is the convergence in (112). Exploiting additivity of the LMGF for independent variables, we conclude that the LMGF of $\tilde{\lambda}_k^{(\delta)}$, scaled by δ and evaluated at t/δ , converges to the sum:

$$\sum_{\ell=1}^N \int_0^t \frac{\Lambda_{\ell}(\pi_{\ell}\tau; \theta)}{\tau} d\tau = \int_0^t \frac{\Lambda_{\text{ave}}(\tau; \theta)}{\tau} d\tau \triangleq \phi(t; \theta), \quad (192)$$

where: *i*) we used the fact that the LMGF of $\mathbf{z}_m^{(\ell)}$ is $\Lambda_{\ell}(t; \theta)$; *ii*) the intermediate equality comes from (53) (having exchanged the integral with the sum); and *iii*) the last equality comes from (55). Moreover, the properties of the rate function in part 6 of Lemma 1 depend only on the fact that $\Lambda_{\alpha z}(t)$ is a logarithmic moment generating function that is finite for all $t \in \mathbb{R}$. Since $\Lambda_{\text{ave}}(\tau; \theta)$ is the LMGF of the average variable $\mathbf{x}_{\text{ave}, i}(\theta)$ (and is finite for all $t \in \mathbb{R}$ by assumption), all the remaining results in part 6 of Lemma 1 hold true, provided that the properties pertaining to $\alpha \mathbf{z}_m$ are now referred to $\mathbf{x}_{\text{ave}, i}(\theta)$.

We conclude that it is legitimate to use the exponential characterization provided in Lemma 1. In particular, since we have $\gamma = 0 < m_{\text{ave}}(\theta)$, the pertinent relation is given by (118) with the choice $\gamma = 0$, yielding:

$$\lim_{\delta \rightarrow 0} \delta \log \mathbb{P}\left[\tilde{\lambda}_k^{(\delta)}(\theta) \leq 0\right] = -\Phi(\theta), \quad (193)$$

where the exponent $\Phi(\theta)$ is accordingly computed as the value of the rate function at $\gamma = 0$, namely,

$$\Phi(\theta) = \sup_{t \in \mathbb{R}} [-\phi(t; \theta)] = -\inf_{t \in \mathbb{R}} \phi(t; \theta). \quad (194)$$

Using the lower bound in (187), we can readily conclude from (193) and from the definitions appearing in (57) and (194) that:

$$\liminf_{\delta \rightarrow 0} \delta \log p_k^{(\delta)} \geq \max_{\theta \neq \theta_0} \left(-\Phi(\theta)\right) = -\min_{\theta \neq \theta_0} \Phi(\theta) = -\Phi. \quad (195)$$

Let us now focus on the upper bound in (187). By definition, for all $\theta \neq \theta_0$ we have that $\Phi \leq \Phi(\theta)$. Accordingly, the convergence in (118) implies that, given an arbitrary $\epsilon > 0$, for sufficiently small δ we can write:

$$\mathbb{P}\left[\tilde{\lambda}_k^{(\delta)}(\theta) \leq 0\right] \leq e^{-(1/\delta)(\Phi - \epsilon)}. \quad (196)$$

Exploiting (196), the upper bound in (187) yields:

$$\delta \log p_k^{(\delta)} \leq \delta \log(H-1) - \Phi + \epsilon, \quad (197)$$

where we recall that H is the number of hypotheses or admissible models. Due to the arbitrariness of ϵ , we have:

$$\limsup_{\delta \rightarrow 0} \delta \log p_k^{(\delta)} \leq -\Phi. \quad (198)$$

Bridging (195) and (198) implies the desired claim. ■

APPENDIX F

We start by proving an auxiliary lemma.

Lemma 2 (Useful properties of the LMGF $\Lambda_{\ell}(t; \theta)$): The lemma is proved under the same assumptions used in Theorem 4. Let

$$\Lambda_{\ell}(t; \theta) = \log \mathbb{E}\left[e^{t \mathbf{x}_{\ell, i}(\theta)}\right] = \log \mathbb{E}\left[e^{t \log \frac{L_{\ell}(\mathbf{x}_{\ell, i}|\theta_0)}{L_{\ell}(\mathbf{x}_{\ell, i}|\theta)}}\right] \quad (199)$$

be the LMGF of the log-likelihood at the ℓ -th agent, let

$$\Lambda_{\text{ave}}(t; \theta) = \log \mathbb{E}\left[e^{t \mathbf{x}_{\text{ave}, i}(\theta)}\right] = \sum_{\ell=1}^N \Lambda_{\ell}(\pi_{\ell} t; \theta) \quad (200)$$

be the LMGF of the network average of log-likelihoods, $\mathbf{x}_{\text{ave}, i}(\theta) = \sum_{\ell=1}^N \pi_{\ell} \mathbf{x}_{\ell, i}(\theta)$, and let:

$$\phi(t; \theta) = \int_0^t \frac{\Lambda_{\text{ave}}(\tau; \theta)}{\tau} d\tau. \quad (201)$$

Then, we have the following properties:

P1) The error exponent $\Phi(\theta)$ is given by:

$$\Phi(\theta) = -\inf_{t \in \mathbb{R}} \phi(t; \theta) = -\phi(t_{\theta}^*; \theta), \quad (202)$$

where $t_{\theta}^* < 0$ is the unique solution to:

$$\frac{\Lambda_{\text{ave}}(t_{\theta}^*; \theta)}{t_{\theta}^*} = 0. \quad (203)$$

P2) For all $t \in \mathbb{R}$ we have:

$$\Lambda_{\ell}(t; \theta) \geq d_{\ell}(\theta)t, \quad (204)$$

implying in particular that:

$$\Phi(\theta) \leq |t_{\theta}^*| m_{\text{ave}}(\theta). \quad (205)$$

P3) Let π_{\min} and π_{\max} be the minimum and maximum entry of the Perron eigenvector, respectively. Then we have:

$$\frac{1}{\pi_{\max}} \leq |t_{\theta}^*| \leq \frac{1}{\pi_{\min}}. \quad (206)$$

Proof: From the convexity properties of $\phi(t; \theta)$ (see [27], [30] for a detailed summary) we know that the infimum of $\phi(t; \theta)$ in (202) is in fact a unique minimum located at the solution t_{θ}^* to the stationary equation:

$$\phi'(t_{\theta}^*; \theta) = 0, \quad (207)$$

where $'$ denotes derivative w.r.t. t . Therefore, Eq. (203) follows from (201). On the other hand, in view of the convexity properties of $\phi(t; \theta)$, the function $\phi'(t; \theta)$ is strictly increasing in t , and since $\phi'(0; \theta) = \Lambda'_{\text{ave}}(0; \theta) = m_{\text{ave}}(\theta) > 0$ (we use the fact that the first derivative of the LMGF evaluated in 0 is equal to the mean of the relative random variable), from (207) we conclude that the value t_{θ}^* that minimizes the function $\phi(t)$ in (202) cannot but be negative, and the proof of property P1) is complete [27], [30].

Regarding property P2), from the convexity of the local LMGF $\Lambda_{\ell}(t; \theta)$ we can write, for all $t \in \mathbb{R}$:

$$\Lambda_{\ell}(t; \theta) \geq t\Lambda'_{\ell}(0; \theta) = td_{\ell}(\theta). \quad (208)$$

Exploiting (200), (201), (202), and (208), we obtain:

$$\begin{aligned} \Phi(\theta) &= -\phi(t_{\theta}^*; \theta) = -\int_0^{t_{\theta}^*} \frac{\Lambda_{\text{ave}}(\tau; \theta)}{\tau} d\tau \\ &= \sum_{\ell=1}^N \int_{t_{\theta}^*}^0 \frac{\Lambda_{\ell}(\pi_{\ell}\tau; \theta)}{\tau} d\tau \\ &\leq |t_{\theta}^*| \sum_{\ell=1}^N \pi_{\ell} d_{\ell}(\theta) = |t_{\theta}^*| m_{\text{ave}}(\theta), \end{aligned} \quad (209)$$

and property P2) is proved.

Finally we prove property P3). Making explicit the definition of $\Lambda_{\text{ave}}(t; \theta)$, Eq. (203) can be written as:

$$\frac{\sum_{\ell=1}^N \Lambda_{\ell}(\pi_{\ell} t_{\theta}^*; \theta)}{t_{\theta}^*} = 0. \quad (210)$$

In view of (199), with expectation computed under the model $L_{\ell}(\xi|\theta_0)$, we have that $\Lambda_{\ell}(-1; \theta) = 0$. Accordingly, when $\pi_{\ell} = 1/N$ for all ℓ , Eq. (206) is obvious. Let us focus on the case where the Perron eigenvector is not uniform. From the strict convexity of $\Lambda_{\ell}(t; \theta)$, we know that:

$$\begin{aligned} \Lambda_{\ell}(\pi_{\ell} t; \theta) &> 0 \quad \text{for } t < -\frac{1}{\pi_{\ell}}, \\ \Lambda_{\ell}(\pi_{\ell} t; \theta) &< 0 \quad \text{for } -\frac{1}{\pi_{\ell}} < t < 0, \end{aligned} \quad (211)$$

see Fig. 15. Since the equality in (210) requires that $\Lambda_{\ell}(\pi_{\ell} t_{\theta}^*; \theta)$ takes on at least one positive and one negative value, Eq. (211) implies property P3). ■

Proof of Theorem 5: From (167) we know that $\tilde{\lambda}_{k,i}^{(\delta)}(\theta)$ and $\tilde{\lambda}_{k,i}^{(\delta)}(\theta)$ share the same distribution. Thus, from (26)




Fig. 15. Typical shape of the LMGF of the ℓ -th likelihood.

and (27) we have that (we recall that $\stackrel{d}{=}$ denotes equality in distribution):

$$\begin{aligned} \lambda_{k,i}^{(\delta)}(\theta) &\stackrel{d}{=} \tilde{\lambda}_{k,i}^{(\delta)}(\theta) + (1-\delta)^i \sum_{\ell=1}^N [A^i]_{\ell k} \lambda_{\ell,0}(\theta) \\ &\geq \tilde{\lambda}_{k,i}^{(\delta)}(\theta) + (1-\delta)^i \sum_{\ell=1}^N \pi_{\ell} \lambda_{\ell,0}(\theta) \\ &\quad - \kappa (1-\delta)^i \beta^i \sum_{\ell=1}^N |\lambda_{\ell,0}(\theta)| \\ &= \tilde{\lambda}_{k,i}^{(\delta)}(\theta) + (1-\delta)^i \sum_{\ell=1}^N \pi_{\ell} \lambda_{\ell,0}(\theta) - \frac{K_2(\theta)}{|t_{\theta}^*|} (1-\delta)^i \beta^i, \end{aligned} \quad (212)$$

where the inequality follows from Property 1, and in the last equality we used (72). In view of (212), and since $t_{\theta}^* < 0$, we can write:

$$\begin{aligned} \mathbb{P}[\lambda_{k,i}^{(\delta)}(\theta) \leq 0] &\leq \mathbb{P}\left[\tilde{\lambda}_{k,i}^{(\delta)}(\theta) \leq -(1-\delta)^i \lambda_{\text{ave},0}(\theta) + \frac{K_2(\theta)}{|t_{\theta}^*|} (1-\delta)^i \beta^i\right] \\ &\stackrel{(a)}{=} \mathbb{P}\left[\frac{t_{\theta}^*}{\delta} \tilde{\lambda}_{k,i}^{(\delta)}(\theta) \geq \frac{|t_{\theta}^*|}{\delta} (1-\delta)^i \lambda_{\text{ave},0}(\theta) - \frac{K_2(\theta)}{\delta} (1-\delta)^i \beta^i\right] \\ &\stackrel{(b)}{\leq} \frac{\mathbb{E}\left[\exp\left\{\frac{t_{\theta}^*}{\delta} \tilde{\lambda}_{k,i}^{(\delta)}(\theta)\right\}\right]}{\exp\left\{\frac{|t_{\theta}^*|}{\delta} (1-\delta)^i \lambda_{\text{ave},0}(\theta) - \frac{K_2(\theta)}{\delta} (1-\delta)^i \beta^i\right\}} \\ &\stackrel{(c)}{=} e^{\frac{1}{\delta} \left[\delta \Lambda_{k,i}^{(\delta)}\left(\frac{t_{\theta}^*}{\delta}; \theta\right) - (1-\delta)^i |t_{\theta}^*| \lambda_{\text{ave},0} + K_2(\theta) (1-\delta)^i \beta^i\right]}, \end{aligned} \quad (213)$$

where (a) follows from multiplying by t_{θ}^*/δ both sides of the inequality in the probability brackets and taking into account the fact that $t_{\theta}^* < 0$; (b) follows from applying Chernoff's bound; and in (c) we applied Property P3) and introduced the LMGF of $\tilde{\lambda}_{k,i}^{(\delta)}(\theta)$, which can be explicitly defined as:

$$\begin{aligned} \Lambda_{k,i}^{(\delta)}(t; \theta) &= \log \mathbb{E}\left[e^{t \tilde{\lambda}_{k,i}^{(\delta)}(\theta)}\right] \\ &= \sum_{\ell=1}^N \sum_{m=0}^{i-1} \Lambda_{\ell}(\delta(1-\delta)^m [A^{m+1}]_{\ell k} t; \theta), \end{aligned} \quad (214)$$

with $\Lambda_\ell(t; \theta)$ being the LMGF of the log-likelihood ratio $\mathbf{x}_{\ell, m+1}(\theta)$. Now, letting

$$c_i \triangleq (1 - \delta)^{i-1}, \quad (215)$$

and applying Eqs. (85) and (86) from [30] to the inner summation in (214), we have the following representation:

$$\begin{aligned} & \Lambda_{k,i}^{(\delta)} \left(\frac{t_\theta^*}{\delta}; \theta \right) \\ &= \frac{1}{\delta} \left[\sum_{\ell=1}^N \int_{c_i \pi_\ell t_\theta^*}^{\pi_\ell t_\theta^*} \frac{\Lambda_\ell(\tau; \theta)}{\tau} d\tau + \mathcal{O}(\delta) \right] \\ &\stackrel{(a)}{=} \frac{1}{\delta} \left[\phi(t_\theta^*; \theta) - \sum_{\ell=1}^N \int_0^{c_i \pi_\ell t_\theta^*} \frac{\Lambda_\ell(\tau; \theta)}{\tau} d\tau + \mathcal{O}(\delta) \right] \\ &\stackrel{(b)}{=} \frac{1}{\delta} \left[-\Phi(\theta) + \int_{-c_i \pi_\ell |t_\theta^*|}^0 \frac{\Lambda_\ell(\tau; \theta)}{\tau} d\tau + \mathcal{O}(\delta) \right] \\ &\stackrel{(c)}{\leq} \frac{1}{\delta} \left[-\Phi(\theta) + c_i |t_\theta^*| \sum_{\ell=1}^N \pi_\ell d_\ell(\theta) + \mathcal{O}(\delta) \right], \quad (216) \end{aligned}$$

where (a) follows from (201), while (b) and (c) from properties P1) and P2) in Lemma 2, respectively. Using now (216) in (213) and using the definition of $K_1(\theta)$ in (71) we get the upper bound in (73). ■

APPENDIX G

Proof of Corollary 1: We now determine the adaptation time as the critical instant after which we stay close to the exponent Φ , in the precise sense specified by (75). Let us consider first the case where $\lambda_{\text{ave}}(\theta) \geq m_{\text{ave}}(\theta)$ for all $\theta \neq \theta_0$. In this case, we have $K_1(\theta) \leq 0$ for all θ and, hence, in view of (73), condition (75) will be met if we ensure that:

$$i > \frac{1}{\log \beta^{-1}} \log \frac{K_2}{\epsilon \Phi} \Rightarrow K_2 \beta^i < \epsilon \Phi, \quad (217)$$

which shows that the choice for T_{ASL} in (76) guarantees (75) for all $i > T_{\text{ASL}}$.

We continue by examining the unfavorable case where $\lambda_{\text{ave}}(\theta) < m_{\text{ave}}(\theta)$ for at least one value $\theta \neq \theta_0$. In this case we have $K_1 = \max_{\theta \neq \theta_0} K_1(\theta) > 0$, and we can write:

$$i > \frac{1}{\log(1 - \delta)^{-1}} \log \frac{K_1}{\epsilon \Phi} \Rightarrow (1 - \delta)^i K_1 < \epsilon \Phi. \quad (218)$$

Then, if we set the adaptation time T_{ASL} according to the law in (218), the quantity β^i appearing in (73) would decay to zero as $\approx \beta^{1/\delta}$, and, hence, would be incorporated into the higher-order term $\mathcal{O}(\delta)$, and the claim of the corollary is proved. ■

ACKNOWLEDGMENT

The authors wish to thank the anonymous reviewers for the careful reading of the manuscript and for the valuable suggestions that helped improve the quality of the work.

REFERENCES

- [1] V. Bordignon, V. Matta, and A. H. Sayed, "Adaptation in online social learning," in *Proc. 28th Eur. Signal Process. Conf. (EUSIPCO)*, Jan. 2021, pp. 2170–2174.
- [2] C. Chamley, *Rational Herds: Economic Models of Social Learning*. Cambridge, U.K.: Cambridge Univ. Press, 2004.
- [3] 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.
- [4] P. Molavi, A. Jadbabaie, K. Rahnama 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.
- [5] 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.
- [6] P. Molavi, A. Tahbaz-Salehi, and A. Jadbabaie, "A theory of non-Bayesian social learning," *Econometrica*, vol. 86, no. 2, pp. 445–490, Mar. 2018.
- [7] V. Krishnamurthy and H. V. Poor, "Social learning and Bayesian games in multiagent signal processing: How do local and global decision makers interact?" *IEEE Signal Process. Mag.*, vol. 30, no. 3, pp. 43–57, May 2013.
- [8] V. Krishnamurthy, "Interactive sensing and decision making in social networks," *Found. Trends Signal Process.*, vol. 7, nos. 1–2, pp. 1–196, 2014.
- [9] C. Chamley, A. Scaglione, and L. Li, "Models for the diffusion of beliefs in social networks: An overview," *IEEE Signal Process. Mag.*, vol. 30, no. 3, pp. 16–29, May 2013.
- [10] E. Yildiz, A. Ozdaglar, D. Acemoglu, A. Saberi, and A. Scaglione, "Binary opinion dynamics with stubborn agents," *ACM Trans. Econ. Comput.*, vol. 1, no. 4, p. 19, Dec. 2013.
- [11] D. Acemoglu, A. Ozdaglar, and A. ParandehGheibi, "Spread of (mis)information in social networks," *Games Econ. Behav.*, vol. 70, no. 2, pp. 194–227, Nov. 2010.
- [12] B. Golub and M. O. Jackson, "Naïve learning in social networks and the wisdom of crowds," *Amer. Econ. J., Microecon.*, vol. 2, no. 1, pp. 112–149, Feb. 2010.
- [13] D. Acemoglu and A. Ozdaglar, "Opinion dynamics and learning in social networks," *Dyn. Games Appl.*, vol. 1, no. 1, pp. 3–49, Mar. 2011.
- [14] A. Jadbabaie, P. Molavi, and A. Tahbaz-Salehi, "Information heterogeneity and the speed of learning in social networks," *Columbia Bus. School Res. Paper*, nos. 13–28, pp. 1–38, May 2013.
- [15] X. Zhao and A. H. Sayed, "Learning over social networks via diffusion adaptation," in *Proc. Asilomar Conf. Signals, Syst. Comput.*, Nov. 2012, pp. 709–713.
- [16] H. Salami, B. Ying, and A. H. Sayed, "Social learning over weakly connected graphs," *IEEE Trans. Signal Inf. Process. Over Netw.*, vol. 3, no. 2, pp. 222–238, Jun. 2017.
- [17] A. Nedic, 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.
- [18] A. Lalitha, T. Javidi, and A. D. Sarwate, "Social learning and distributed hypothesis testing," *IEEE Trans. Inf. Theory*, vol. 64, no. 9, pp. 6161–6179, Sep. 2018.
- [19] V. Matta, A. Santos, and A. H. Sayed, "Exponential collapse of social beliefs over weakly-connected heterogeneous networks," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, May 2019, pp. 5267–5271.
- [20] V. Matta, V. Bordignon, A. Santos, and A. H. Sayed, "Interplay between topology and social learning over weak graphs," *IEEE Open J. Signal Process.*, vol. 1, pp. 99–119, Jul. 2020.
- [21] V. Matta and A. H. Sayed, "Estimation and detection over adaptive networks," in *Cooperative and Graph Signal Processing*, P. Djuric and C. Richard, Eds. Amsterdam, The Netherlands: Elsevier, 2018, pp. 69–106.
- [22] T. Cover and J. Thomas, *Elements of Information Theory*. Hoboken, NJ, USA: Wiley, 1991.
- [23] A. Sayed, "Adaptation, learning, and optimization over networks," *Found. Trends Mach. Learn.*, vol. 7, nos. 4–5, pp. 311–801, 2014.
- [24] R. Horn and C. Johnson, *Matrix Analysis*. Cambridge, U.K.: Cambridge Univ. Press, 2013.
- [25] H. Chernoff, "A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations," *Ann. Math. Statist.*, vol. 23, no. 4, pp. 493–507, Dec. 1952.
- [26] J. Shao, *Mathematical Statistics*. New York, NY, USA: Springer, 2003.

- [27] V. Matta, P. Braca, S. Marano, and A. H. Sayed, "Diffusion-based adaptive distributed detection: Steady-state performance in the slow adaptation regime," *IEEE Trans. Inf. Theory*, vol. 62, no. 8, pp. 4710–4732, Aug. 2016.
- [28] A. Dembo and O. Zeitouni, *Large Deviations Techniques and Applications*. New York, NY, USA: Springer, 1998.
- [29] F. den Hollander, *Large Deviations*. Providence, RI, USA: American Mathematical Society, 2008.
- [30] V. Matta, P. Braca, S. Marano, and A. H. Sayed, "Distributed detection over adaptive networks: Refined asymptotics and the role of connectivity," *IEEE Trans. Signal Inf. Process. Over Netw.*, vol. 2, no. 4, pp. 442–460, Dec. 2016.
- [31] E. N. Gilbert, "Capacity of a burst-noise channel," *Bell Syst. Tech. J.*, vol. 39, no. 5, pp. 1253–1265, Sep. 1960.
- [32] E. O. Elliott, "Estimates of error rates for codes on burst-noise channels," *Bell Syst. Tech. J.*, vol. 42, no. 5, pp. 1977–1997, Sep. 1963.
- [33] A. H. Sayed and X. Zhao, "Asynchronous adaptive networks," in *Cooperative and Graph Signal Processing*, P. Djuric and C. Richard, Eds. Amsterdam, The Netherlands: Elsevier, 2018, pp. 3–68.
- [34] M. Loève, "On almost sure convergence," in *Proc. 2nd Berkeley Symp. Math. Statist. Prob. (Univ. Calif. Press)*, 1951, pp. 279–303.
- [35] R. Durrett, *Probability Theory and Examples*. Cambridge, U.K.: Cambridge Univ. Press, 2019.
- [36] W. Feller, *An Introduction to Probability Theory and its Applications*, vol. 2. Hoboken, NJ, USA: Wiley, 1971.

Virginia Bordignon (Graduate Student Member, IEEE) received the degrees in engineering from École Centrale de Lyon, France, and the Federal University of Rio Grande do Sul (UFRGS), Brazil, and the M.S. degree in electrical engineering from UFRGS. She is currently pursuing the Ph.D. degree with the Adaptive Systems Laboratory, École Polytechnique Fédérale de Lausanne (EPFL), Switzerland. Her research interests include statistical inference, distributed learning, and information processing over networks.

Vincenzo Matta (Member, IEEE) is currently a Full Professor of telecommunications with the Department of Information and Electrical Engineering and Applied Mathematics (DIEM), University of Salerno. His research interests include adaptation and learning over networks, information processing in multi-agent systems, and network security. He serves as an Associate Editor for IEEE OPEN JOURNAL OF SIGNAL PROCESSING, and served as an Associate Editor for IEEE TRANSACTIONS ON AEROSPACE AND ELECTRONIC SYSTEMS, IEEE SIGNAL PROCESSING LETTERS, and IEEE TRANSACTIONS ON SIGNAL AND INFORMATION PROCESSING OVER NETWORKS, and as a Senior Area Editor for IEEE SIGNAL PROCESSING LETTERS.

Ali H. Sayed (Fellow, IEEE) is Dean of engineering at EPFL, Switzerland, where he also leads the Adaptive Systems Laboratory (<http://asl.ee.ucla.edu/>). His research involves several areas including adaptation and learning theories, data and network sciences, statistical inference, information processing theories, and multi-agent systems. He is recognized as a highly cited researcher, a member of the U.S. National Academy of Engineering, and served as the President of the IEEE Signal Processing Society from 2018 to 2019. His work has been recognized with several major awards from IEEE and other professional societies.