Adaptive Social Learning for Tracking Rare Transition Markov Chains

Malek Khammassi
EPFL
Lausanne, Switzerland
malek.khammassi@epfl.ch

Virginia Bordignon
EPFL
Lausanne, Switzerland
virginia.bordignon@epfl.ch

Vincenzo Matta
University of Salerno
Fisciano, Italy
vmatta@unisa.it

Ali H. Sayed
EPFL
Lausanne, Switzerland
ali.sayed@epfl.ch

***Abstract*—Adaptive Social Learning (ASL) enables consistent truth learning in nonstationary environments. In this framework, agents linked by a graph, exchange their local beliefs with neighbors to track some underlying state of interest. This state can drift over time. Previous works have examined the adaptation and learning properties of ASL without relating them to the speed of the drifts. This study assesses the performance of ASL by modeling the true state as a Markov chain. We determine an asymptotic characterization of the ASL tracking performance, revealing the fundamental scaling laws that rule the rare transition regime. We demonstrate that ASL achieves a vanishing probability of error when the average drift time of the Markov chain is smaller than the adaptation time of the ASL algorithm. Simulations illustrate our theoretical findings, providing insights into the ASL performance in dynamic settings.**

***Index Terms*—Adaptive social learning, large deviations, Markov chain, hidden Markov model, opinion formation.**

I. Introduction and Background

Social learning allows a set of agents over a network to form opinions by observing the local environment and exchanging information with their neighbors [1], [2]. The growing literature on the topic presents a rich family of algorithms relying on Bayesian and graph theories to address problems such as hypothesis testing, classification, and opinion formation over social networks [1]–[8].

Social learning strategies are particularly useful in solving decentralized inference problems. These strategies enable a set of K cooperating agents, each collecting streaming observations about a common phenomenon, to learn the hypothesis that best describes the observations. This hypothesis, denoted by $\theta^\star$, is referred to as the *true state*, and is chosen from a set of plausible hypotheses $\Theta = \{\theta_1, \dots, \theta_H\}$. In traditional social learning, each agent k starts from an initial belief vector $\mu_{k,0}$, representing a probability vector over the set of plausible hypotheses Θ. At every time instant $i \geq 1$, each agent k receives a random observation $\boldsymbol{x}_{k,i} \in \mathcal{X}_k$ (we use bold font for random quantities), distributed according to the likelihood $L_k(\cdot|\theta^\star)$. By leveraging its set of likelihood models $\{L_k(\cdot|\theta)\}_{\theta\in\Theta}$, agent k performs a *local* Bayesian update step to incorporate the received observation into its (past) belief vector $\boldsymbol{\mu}_{k,i-1}$ and to construct an *intermediate* belief vector $\psi_{k,i}$ (*self-learning step*). Subsequently, each agent k combines the intermediate beliefs of its neighbors (by means of a weighted geometric or arithmetic averaging rule) to compute the *private* belief vector $\boldsymbol{\mu}_{k,i}$ (*combination step*). Under a stationary setting where the true state does not change over time, repeated application of the aforementioned two steps leads the agents to truth learning [1]–[3], [5], [7].

However, applications with *nonstationary* environments are common, where the true state undergoes continuous and unpredictable drifts over time. This setting requires carefully modifying the social learning algorithm to endow it with *adaptation* abilities. To this end, the work in [9] proposed the *adaptive social learning* (ASL) strategy and characterized its performance both in the steady state and transient regimes. The ASL algorithm takes the form of a two-step recursion that iterates over time as follows:

$$\boldsymbol{\psi}_{k,i}(\theta) \propto L_k(\boldsymbol{x}_{k,i}|\theta)\boldsymbol{\mu}_{k,i-1}^{1-\delta}(\theta) \qquad \text{(self-learning)} \qquad (1)$$

$$\boldsymbol{\mu}_{k,i}(\theta) \propto \prod_{\ell=1}^{K} \left[\boldsymbol{\psi}_{\ell,i}(\theta)\right]^{a_{\ell k}} \qquad \text{(combination)} \qquad (2)$$

where the proportionality symbol $\propto$ indicates that the entries of $\boldsymbol{\mu}_{k,i}$ and $\boldsymbol{\psi}_{k,i}$ are normalized to add up to 1. The quantity $a_{\ell k}$ is a nonnegative weight assigned by agent k to the information received from neighbor ℓ satisfying the following conditions:

$$0 \leq a_{\ell k} \leq 1, \quad \sum_{\ell=1}^{K} a_{\ell k} = 1, \quad a_{\ell k} = 0 \text{ for } \ell \notin \mathcal{N}_k, \qquad (3)$$

where $\mathcal{N}_k$ denotes the neighborhood of agent k (which includes agent k itself). The positive scalar $\delta \in (0,1)$ is an adaptation parameter used to tune the network's ability to adapt in view of changes. This parameter controls the importance of old information (i.e., of the prior belief $\boldsymbol{\mu}_{k,i-1}$) in the update rule (1). Specifically, the smaller the weight assigned to the prior is, the better the algorithm will adapt to drifts.

It is shown in [9] that the adaptation (or transient) time is inversely proportional to δ, while the steady-state error probability vanishes as δ approaches zero. This is one instance of the learning/adaptation trade-off: a larger δ ensures faster adaptation to changes at the expense of a higher steady-state error probability. These findings are established without accounting for the nature of the underlying process that governs the drifts. In fact, the analysis in [9] is conducted under the assumption that the true state remains constant for sufficiently long intervals, allowing the agents to learn it before it changes.

This work evaluates the performance of ASL by modeling the true state as a time-varying stochastic process, thereby accounting for its inherent nonstationary nature. Specifically, the *true state* is modeled as a random process $\boldsymbol{\theta}_i^\star$ that we refer to as the *state process*. The bold font highlights the random nature of the true state $\boldsymbol{\theta}_i^\star$, while its dependence on i emphasizes its nonstationary nature. Formally, we assume that the *state process* $\boldsymbol{\theta}_i^\star$ is a discrete-time homogeneous Markov chain taking values in the state space Θ with the following transition probabilities:

$$\mathbb{P}\left[\boldsymbol{\theta}_i^\star = \theta' | \boldsymbol{\theta}_{i-1}^\star = \theta\right] = \begin{cases} 1 - \varepsilon q_{\theta\theta} & \text{if } \theta = \theta' \\ \varepsilon q_{\theta\theta'} & \text{if } \theta \neq \theta' \end{cases}, \quad \theta, \theta' \in \Theta, \tag{4}$$

where $0 < \varepsilon < 1$ represents the drift parameter, $q_{\theta\theta'} \geq 0$ for $\theta' \neq \theta$, and $\sum_{\theta' \neq \theta} q_{\theta\theta'} = q_{\theta\theta}$ for each $\theta \in \Theta$. We assume that the Markov chain $\boldsymbol{\theta}_i^\star$ is irreducible and aperiodic, thereby guaranteeing the existence of its stationary distribution, which we denote by $p_s = [p_s(1), \ldots, p_s(H)]^\mathsf{T}$.

Before proceeding with our result, we introduce a set of assumptions in the following section.

II. Assumptions

The observations and likelihood functions considered in this work satisfy the following assumptions.

Assumption 1 (Bounded log-likelihood ratios) *There exists a positive constant B such that*

$$\max_{k \in \{1,\ldots,K\}} \max_{\theta,\theta' \in \Theta} \sup_{x \in \mathcal{X}_k} \left| \log \frac{L_k(x|\theta)}{L_k(x|\theta')} \right| \leq B. \tag{5}$$

□

Assumption 1 is automatically satisfied when the observations are discrete random variables with the same finite support.

Assumption 2 (Statistical model) *Conditioned on the current value of the state process, the current observations are i) independent of any past states or observations; and ii) independent over space. Formally, let $\boldsymbol{x}_i \triangleq \{\boldsymbol{x}_{k,i}\}_{k=1}^K$ collect all observations from across the agents at time i. Then, the joint likelihood at time i satisfies*

$$L(\boldsymbol{x}_i | \boldsymbol{\theta}_i^\star, \ldots, \boldsymbol{\theta}_1^\star, \boldsymbol{x}_{i-1}, \ldots, \boldsymbol{x}_1) = L(\boldsymbol{x}_i | \boldsymbol{\theta}_i^\star) \tag{6a}$$

$$= \prod_{k=1}^K L_k(\boldsymbol{x}_{k,i} | \boldsymbol{\theta}_i^\star). \tag{6b}$$

□

Assumption 3 (Global Identifiability) *For each pair $\theta, \theta' \in \Theta$ such that $\theta \neq \theta'$, there exists at least one agent k such that the Kullback-Leibler (KL) divergence [10] between $L_k(.|\theta)$ and $L_k(.|\theta')$ is positive, which reads as*

$$D_k(\theta, \theta') \triangleq \mathbb{E}_\theta \left[\log \frac{L_k(\boldsymbol{x}|\theta)}{L_k(\boldsymbol{x}|\theta')} \right] > 0, \tag{7}$$

where the subscript on the expectation operator means that the expectation is computed under $L_k(\cdot|\theta)$. □

The global identifiability assumption requires that, for any two distinct hypotheses in Θ, at least one agent must be capable to distinguish between them.

Next, we introduce the following assumption on the initial beliefs of the agents.

Assumption 4 (Positive initial beliefs) *The initial beliefs of all agents are positive, i.e., $\mu_{k,0}(\theta) > 0$ for all $k \in \{1, \ldots, K\}$ and $\theta \in \Theta$.* □

Finally, we introduce some standard assumptions on the weights $a_{\ell k}$.

Assumption 5 (Connected network) *The network graph is assumed to be connected, which means that, given any pair of distinct nodes (ℓ, k), a path with nonzero weights exists between ℓ and k in both directions (the forward and backward paths need not be the same).* □

It is useful to define a *combination matrix* A whose (ℓ, k) entry is $a_{\ell k}$. Under conditions (3), A is a *left-stochastic* matrix. Moreover, under Assumption 5, the matrix A is irreducible. Therefore, in view of the Perron-Frobenius theorem [11], [12], matrix A has spectral radius equal to 1 and a single eigenvalue equal to 1. This eigenvalue is associated with an eigenvector π, referred to as the Perron vector, which can be scaled to have all positive entries that add up to 1, namely,

$$A\pi = \pi, \quad \sum_{\ell=1}^K \pi_\ell = 1, \quad \pi_\ell > 0 \text{ for all } \ell = 1, 2, \ldots, K. \tag{8}$$

Assumption 6 (Primitive matrix) *Under Assumption 5, the combination matrix is further said to to be primitive when the only eigenvalue on the unit circle is the real eigenvalue at 1. Note that, a sufficient condition for the matrix to be primitive is that the network is strongly connected, which means that, in addition to being connected, it has at least one self-loop, i.e., there exists at least one agent k with $a_{kk} > 0$.*

III. ASL and Markov Chains

To evaluate the performance of ASL in tracking the Markov chain $\boldsymbol{\theta}_i^\star$, we consider the natural performance metric of error probability. Specifically, according to the maximum-a-posteriori (MAP) criterion, each agent k makes a mistake whenever its belief is not maximized at the true hypothesis. Therefore, the instantaneous error probability of agent k at time instant i can be expressed as

$$p_{k,i} \triangleq \mathbb{P}\left[\arg\max_{\theta \in \Theta} \boldsymbol{\mu}_{k,i}(\theta) \neq \boldsymbol{\theta}_i^\star\right]. \tag{9}$$

In the theory of adaptation, the learning performance is characterized by evaluating the steady-state performance of this error probability as $i \to \infty$ [13]. With reference to the ASL algorithm under a stationary setting, this translates into the *consistency* conclusion that the steady-state error probability converges to 0 as the adaptation parameter δ tends to 0. In comparison, in the theory of Bayesian filtering and smoothing [14] [15], consistency is examined for the rare transition regime when the drift parameter ε vanishes. In our study, we blend these two approaches by incorporating both parameters into the error probability: the *drift* parameter ε (related to the environment) and the *adaptation* parameter δ (related to the social learning algorithm).

The work in [16] is a relevant study that also considers these two parameters, albeit for a different inferential problem, namely, for *parameter estimation* or *regression*. The authors analyze the tracking properties of the (centralized) least-mean-squares (LMS) algorithm when the underlying parameter evolves according to a finite-state Markov chain with infrequent jumps. They establish a bound on the mean-square error involving both ε and the LMS step-size parameter. Then, by imposing a certain asymptotic relationship between the LMS step-size parameter and ε, they establish additional guarantees on the tracking algorithm.

Returning to the decision making problem that is of interest in social learning, in the next theorem, we establish a sufficient condition for the parameters ϵ and δ to ensure that each network agent tracks the drifting hypothesis with vanishing error probability.

Theorem 1 (Consistency for small-ε) *Under Assumptions 1-6, let $\delta = \delta_\varepsilon$ be the adaptation parameter used for a given drift parameter ε. If the following conditions are satisfied:*

$$\lim_{\varepsilon \to 0} \delta_\varepsilon = 0, \tag{10}$$

$$\lim_{\varepsilon \to 0} \frac{\varepsilon}{\delta_\varepsilon} = 0, \tag{11}$$

then each agent $k \in \{1, \dots, K\}$ consistently learns the truth as ε goes to zero, which is formally expressed as

$$\lim_{\varepsilon \to 0} \limsup_{i \to \infty} p_{k,i} = 0. \tag{12}$$

Proof: Due to space limitations, we provide a sketch of the proof. In what follows, the notation $f(\varepsilon) = O(g(\varepsilon))$ means that $\lim_{\varepsilon \to 0} \frac{f(\varepsilon)}{g(\varepsilon)}$ is finite.

Let $T_\varepsilon > 0$ be an integer function of ε. We can write the instantaneous error probability as follows:

$$\begin{aligned} p_{k,i} &= \mathbb{P}\left[\operatorname*{argmax}_{\theta \in \Theta} \boldsymbol{\mu}_{k,i}(\theta) \neq \boldsymbol{\theta}_i^\star, \text{no jumps in } [i - T_\varepsilon, i]\right] \\ &+ \mathbb{P}\left[\operatorname*{argmax}_{\theta \in \Theta} \boldsymbol{\mu}_{k,i}(\theta) \neq \boldsymbol{\theta}_i^\star, \text{one or more jumps in } [i - T_\varepsilon, i]\right] \\ &\leq \mathbb{P}\left[\operatorname*{argmax}_{\theta \in \Theta} \boldsymbol{\mu}_{k,i}(\theta) \neq \boldsymbol{\theta}_i^\star, \text{no jumps in } [i - T_\varepsilon, i]\right] \\ &+ \mathbb{P}\Big[\text{one or more jumps in } [i - T_\varepsilon, i]\Big], \end{aligned} \tag{13}$$

where a jump in the Markov chain $\boldsymbol{\theta}_i^\star$ is observed at instant j if $\boldsymbol{\theta}_j^\star \neq \boldsymbol{\theta}_{j-1}^\star$. We see from (13) that $p_{k,i}$ is upper bounded by the sum of two terms that depend on T_ε. We would like to suitably design T_ε such that the probability of one or more jumps of the Markov chain in $[i - T_\varepsilon, i]$ is "small". Formally, we can show that if T_ε verifies the following two conditions

$$\lim_{\varepsilon \to 0} \varepsilon T_\varepsilon = 0, \tag{14}$$

$$\lim_{\varepsilon \to 0} T_\varepsilon = +\infty, \tag{15}$$

then

$$\mathbb{P}\Big[\text{one or more jumps in } [i - T_\varepsilon, i]\Big] = O(\varepsilon T_\varepsilon). \tag{16}$$

From (13) and (16), we see that upper bounding $p_{k,i}$ reduces to upper bounding the dominant term $\mathbb{P}\left[\operatorname{argmax}_{\theta \in \Theta} \boldsymbol{\mu}_{k,i}(\theta) \neq \boldsymbol{\theta}_i^\star, \text{ no jumps in } [i - T_\varepsilon, i]\right]$. This term exhibits a desirable property because it involves an event where the state process remains constant in the time interval $[i - T_\varepsilon, i]$. This allows us to handle this term with the techniques exploited in [9, Appendix F] to establish the following upper bound on the steady-state error probability

$$\begin{aligned} \limsup_{i \to \infty} p_{k,i} &\leq \sum_{\theta'} \sum_{\theta \neq \theta'} \exp\Bigg(\frac{1}{\delta}\Big(- \Phi(\theta', \theta) \\ &+ |t^\star_{\theta',\theta}|(1-\delta)^{T_\varepsilon}\big(D(\theta', \theta) + B\big) + O(\delta)\Big)\Bigg) p_s(\theta') \\ &+ O(\varepsilon T_\varepsilon), \end{aligned} \tag{17}$$

where $\Phi(\theta', \theta) > 0$ and $t^\star_{\theta',\theta} < 0$ can be respectively computed as $\Phi(\theta)$ and $t^\star_\theta$ in [9, Appendix F, Lemma 2] (with θ' and θ in our notation representing respectively θ and θ_0 in [9]). Furthermore, $D(\theta', \theta) \triangleq \sum_{\ell=1}^K \pi_\ell D_\ell(\theta', \theta)$.

To guarantee that the steady-state error probability vanishes as ε goes to zero, we investigate a design for δ that leads the upper bound in (17) to vanish as ε goes to zero. If the adaptation parameter $\delta = \delta_\varepsilon$ verifies (10) and (11), then it is always possible to choose $T_\varepsilon = \alpha/\delta_\varepsilon$, for $\alpha > 0$, which indeed satisfies (14) and (15). Now, with this choice of T_ε, we have

$$\lim_{\varepsilon \to 0} (1 - \delta_\varepsilon)^{\frac{\alpha}{\delta_\varepsilon}} = \lim_{\delta_\varepsilon \to 0} (1 - \delta_\varepsilon)^{\frac{\alpha}{\delta_\varepsilon}} = \exp(-\alpha). \tag{18}$$

Therefore, there exists α small enough such that

$$\lim_{\varepsilon \to 0} \limsup_{i \to \infty} p_{k,i} = 0. \tag{19}$$

■

Theorem 1 shows that the agents can achieve consistent learning if the adaptation parameter of ASL satisfies conditions (10) and (11). The first condition is necessary because, as is known from the analysis of adaptive social learning [9], if the adaptation parameter does not vanish, the steady-state error probability does not vanish, even in the limiting case where no transitions occur ($\varepsilon = 0$). We also note that the limiting case $\delta_\epsilon = \epsilon = 0$ brings us back to traditional social learning designed for stationary environments.

Condition (11) requires that $\varepsilon/\delta_\varepsilon$ goes to zero as ε vanishes. This condition admits a useful interpretation in terms of the characteristic times of the system. In fact, it is shown in [9] that

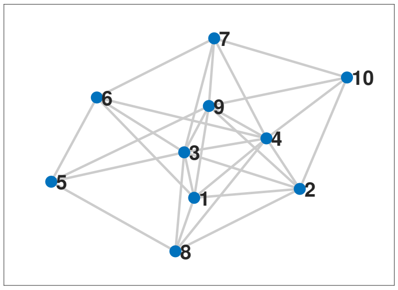

Fig. 1: Strongly connected network with $K = 10$ agents.

the adaptation time of ASL, denoted by $\mathrm{T}_{\mathrm{ASL}}$, is proportional to $1/\delta_\varepsilon$. On the other hand, the average time between two consecutive jumps in the Markov chain, denoted by T_{MC}, can be shown to be proportional to $1/\varepsilon$. Therefore, Eq. (11) can be translated into the following condition

$$\lim_{\varepsilon \to 0} \frac{\mathrm{T}_{\mathrm{ASL}}}{\mathrm{T}_{\mathrm{MC}}} = 0. \tag{20}$$

In other words, condition (11) imposes that the adaptation time of ASL ($\mathrm{T}_{\mathrm{ASL}}$) is smaller than the Markov chain time (T_{MC}) that governs the model drifts. This ensures that the ASL process reacts sufficiently fast to track the drifts, ultimately guaranteeing consistency of the social learning algorithm.

IV. Illustrative Examples

We consider the strongly connected network of $K = 10$ agents depicted in Fig. 1, where all agents are assumed to have a self-loop that is not displayed. We design the combination matrix A using the uniform-averaging combination rule [12], [13], which results in a left-stochastic matrix satisfying Assumptions 5-6. We assume that the Markov chain $\boldsymbol{\theta}_i^\star$ takes on values in the state space $\Theta = \{0.1, 0.2, 0.3\}$ with the following transition probabilities

$$\mathbb{P}\left[\boldsymbol{\theta}_i^\star = \theta | \boldsymbol{\theta}_{i-1}^\star = \theta'\right] = \begin{cases} 1-\varepsilon & \text{if } \theta = \theta' \\ \frac{\varepsilon}{H-1} & \text{if } \theta \neq \theta' \end{cases}, \quad \theta, \theta' \in \Theta. \tag{21}$$

Note that, in this setting, we can show that $\mathrm{T}_{\mathrm{MC}} = 1/\varepsilon$. The likelihood models of the agents belong to a family of binomial distributions with number of trials $n = 5$ and probability of success $\theta \in \Theta$, which are given by the following probability mass function:

$$f_\theta(x) = \binom{5}{x} \theta^x (1-\theta)^{5-x}, \ x \in \{1, \dots, 5\}, \tag{22}$$

where $\binom{5}{x} = \frac{5!}{x!(5-x)!}$ and $f_\theta(x)$ represents the probability of x successes in n trials. The parameter θ takes on values in the set $\Theta = \{0.1, 0.2, 0.3\}$. Thus, in this setting, the state process $\boldsymbol{\theta}_i^\star$ represents the success probability. It is easily seen that the choice of likelihood models in (22) satisfies Assumption 1.

Furthermore, we assume that the observations $\boldsymbol{x}_{k,i}$ are independent across the agents (Assumption 2). We also assume

TABLE I: Identifiability setup of the agents.

Agent k	**Likelihood model**: $L_k(\cdot\|\theta)$		
	$\theta = 0.1$	$\theta = 0.2$	$\theta = 0.3$
1, 5, 10	$f_{0.1}$	$f_{0.2}$	$f_{0.3}$
2, 6	$f_{0.1}$	$f_{0.2}$	$f_{0.2}$
3, 7	$f_{0.1}$	$f_{0.1}$	$f_{0.3}$
4, 8	$f_{0.1}$	$f_{0.2}$	$f_{0.1}$

that some agents cannot locally differentiate all pairs of hypotheses, as it can be verified from the identifiability setup in Table I. The setup in the Table I satisfies Assumption 3.

Theorem 1 provides sufficient conditions for the adaptation parameter to guarantee consistency. Specifically, the adaptation parameter δ_ε should be a function of ε verifying (10)-(11). In order to illustrate this result, we consider the following choices of δ_ε:

$$\delta_\varepsilon = \varepsilon^{\frac{2}{3}}, \ \delta_\varepsilon = \varepsilon^{\frac{1}{2}}, \ \delta_\varepsilon = \frac{1}{\log \frac{1}{\varepsilon}}. \tag{23}$$

We can see that each of the above choices of δ_ε adhere to both conditions (10) and (11).

In order to illustrate the performance of ASL under the different choices of δ_ε, we consider 20 values of ε (uniformly spaced in the log domain) in the interval $[0.005, 0.1]$. Then, for each value of ε, we run the ASL algorithm considering the three different choices of δ_ε in (23). For each of these three configurations, we consider 1000 iterations of the ASL algorithm, which we assume are sufficient to reach the steady state. For each time sample i, we execute the ASL algorithm over 25000 Monte Carlo runs to estimate the error probability at instant i. We plot the steady-state error probability as a function of $T_{\mathrm{MC}} = 1/\varepsilon$ in Fig. 2a. Then, we set $\varepsilon = 0.005$ and plot the instantaneous error probability as a function of the time index i in Fig. 2b.

In Fig. 2a, we can see that the different choices of the adaptation parameter in (23) result in a vanishing steady-state error probability as the average drift time T_{MC} diverges to infinity. This behavior confirms the predictions of Theorem 1. In addition, we see that distinct choices of the adaptation parameter δ_ε lead to distinct levels of steady-state error probability, suggesting the existence of an optimal configuration that maximizes the performance.

To gain further insight into the performance of ASL under different adaptation parameters, we examine the time evolution of the instantaneous error probability for a fixed $\varepsilon = 0.005$, illustrated in Fig. 2b. The results reveal discernible variations in performance across different adaptation parameters. Specifically, setting $\delta_\varepsilon = \frac{1}{\log \frac{1}{\varepsilon}} = 0.1887$ yields the lowest instantaneous error probability, while $\delta_\varepsilon = \varepsilon^{\frac{1}{2}} = 0.0707$ results in a higher error probability. Moreover, selecting $\delta_\varepsilon = \varepsilon^{\frac{2}{3}} = 0.0292$ notably increases the error probability compared to the first two choices, indicating that ASL becomes too slow to track the Markov chain as δ_ε decreases.

Next, we examine the belief evolution of ASL over time, for $\delta_\varepsilon = \frac{1}{\log \frac{1}{\varepsilon}}$, which exhibits the best performance among the curves in Fig. 2. Simulations are performed with $\varepsilon = 0.005$

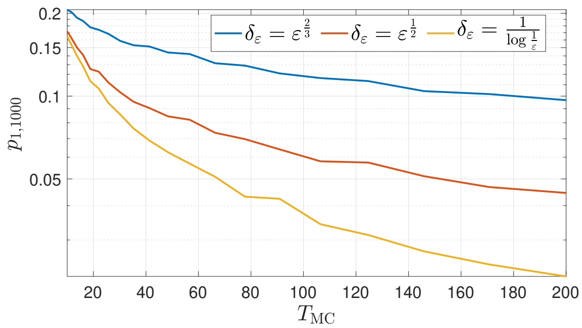

(a) Steady-state error probability.

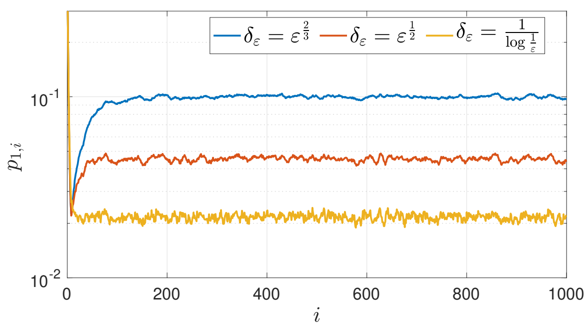

(b) Instantaneous error probability for $\varepsilon = 0.005$.

Fig. 2: Error probability of agent 1 as a function of the drift time (Fig. 2a) and as a function of the number of iterations (Fig. 2b), for different choices of δ_ε.

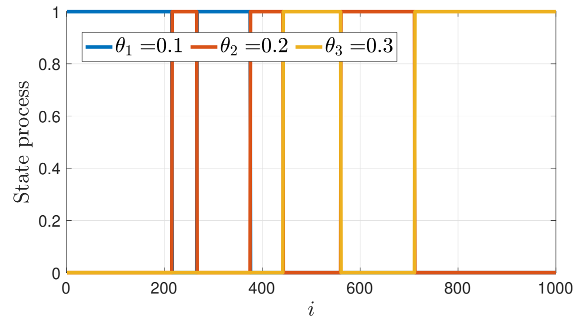

(a) Time evolution of the state process.

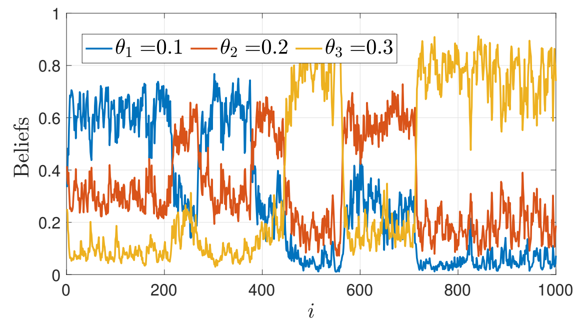

(b) Time evolution of the belief of agent 1.

Fig. 3: ASL tracking ability.

over 1000 time samples, and the resulting belief evolution is displayed in Fig. 3. Fig. 3a illustrates the state process evolution over time. We can see that the belief evolution of ASL, depicted in Fig. 3b, indicates that ASL effectively tracks the state process evolution.

V. Concluding Remarks

ASL is a social learning strategy enabling opinion formation and consistent truth learning within dynamic environments. Existing studies on the ASL performance are conducted without accounting for the process that governs the drifts in the underlying phenomenon of interest. This work addresses the nonstationary nature of the learning environment by modeling the true state as a time-varying stochastic process. Specifically, we assess the ASL performance in scenarios where the true state follows a Markov chain. Although this approach introduces complexity into the error probability analysis, it provides deeper insight into the fundamental interplay between the adaptation offered by ASL and the drift rate in the underlying Markov process.

References

[1] A. Jadbabaie, P. Molavi, A. Sandroni, and A. Tahbaz-Salehi, "Non-Bayesian social learning," *Games Econ. Behav.*, vol. 76, no. 1, pp. 210–225, Sep 2012.

[2] X. Zhao and A. H. Sayed, "Learning over social networks via diffusion adaptation," in Proc. *Asilomar Conf. Signals, Syst. and Comput.*, 2012, pp. 709–713.

[3] A. Nedić, A. Olshevsky, and C. A. Uribe, "Fast convergence rates for distributed non-bayesian learning," *IEEE Trans. Autom. Control*, vol. 62, no. 11, pp. 5538–5553, Mar. 2017.

[4] M. Pooya, T.-S. Alireza, and J. Ali, "A theory of non-bayesian social learning," *Econometrica*, vol. 86, no. 2, pp. 445–490, Mar. 2018.

[5] 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, May 2018.

[6] H. Salami, B. Ying, and A. H. Sayed, "Social learning over weakly connected graphs," *IEEE Trans. Signal and Inf. Process. over Networks*, vol. 3, no. 2, pp. 222–238, Feb. 2017.

[7] P. Molavi, K. R. Rad, A. Tahbaz-Salehi, and A. Jadbabaie, "On consensus and exponentially fast social learning," in *2012 Amer. Control Conf. (ACC)*, pp. 2165–2170.

[8] 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.

[9] V. Bordignon, V. Matta, and A. H. Sayed, "Adaptive social learning," *IEEE Trans. Inf. Theory*, vol. 67, no. 9, pp. 6053–6081, Jul. 2021.

[10] T. Cover and J. Thomas, Elements of Information Theory, 2nd ed. Wiley, 2006.

[11] R. A. Horn and C. R. Johnson, *Matrix Analysis*. Cambridge Univ. Press, 1985.

[12] A. H. Sayed, *Inference and Learning from Data*. Cambridge University Press, 2022.

[13] ——, "Adaptation, learning, and optimization over networks," *Found. Trends Mach. Learn.*, vol. 7, no. 4–5, p. 311–801, jul 2014.

[14] R. Khasminskii and O. Zeitouni, "Asymptotic filtering for finite state Markov chains," *Stochastic Processes and their Applications*, vol. 63, 1996.

[15] L. Shue, B. Anderson, and F. De Bruyne, "Asymptotic smoothing errors for hidden Markov models," *IEEE Trans. Signal Process.*, vol. 48, no. 12, pp. 3289–3302, Dec. 2000.

[16] G. Yin and V. Krishnamurthy, "LMS algorithms for tracking slow Markov chains with applications to hidden Markov estimation and adaptive multiuser detection," *IEEE Trans. Inf. Theory*, vol. 51, no. 7, pp. 2475–2490, 2005.