Optimal Combination Policies for Error Exponent Maximization in Social Learning

Ping Hu
School of Engineering
EPFL
Lausanne, Switzerland
ping.hu@epfl.ch

Mert Kayaalp
School of Engineering
EPFL
Lausanne, Switzerland
mert.kayaalp@epfl.ch

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

***Abstract*—Distributed decision-making over networks involves multiple agents collaborating to achieve a common goal. In the social learning process, where agents aim at inferring an unknown state from a stream of local observations, the probability of error in their decisions converges to zero exponentially in the asymptotic regime. The rate of this convergence, known as the error exponent, is influenced by the combination policy employed by the network. This work addresses the challenge of identifying the optimal combination policies to maximize the error exponent. We establish an upper bound on the achievable error exponents by the social learning rule and provide the conditions for the combination policy to reach this upper bound. By implementing the optimized policy, we enhance the error exponent, leading to improved accuracy and efficiency in the distributed decision-making process.**

***Index Terms*—social learning, combination policy, large deviations, error exponent.**

I. Introduction

Social learning (SL) is a key paradigm in online distributed decision making, where a group of agents receives a stream of observations that depend on an unknown state of nature. The agents are tasked with inferring the true state, from a finite set of hypothesized states, that best explains their observations. This paradigm is valuable in applications such as sensor networks, robotics, and social networks, where global coordination is challenging or impractical [1]–[4].

In the SL framework, each agent's state is described by a belief vector that reflects the agent's confidence in each hypothesis. There are various SL algorithms in the literature, most of which are designed based on the following two principles: i) incorporating new information from the streaming observations using Bayes' rule, and ii) reaching a decision consensus by combining information from other agents through specific pooling rules [4], [5]. Different choices of pooling rules lead to distinct SL algorithms. Some representative pooling rules include the arithmetic rule [6]–[8] and the geometric rule [9]–[12]. All these rules provide a uniformly consistent learning strategy, ensuring that each agent's belief in the true state converges to 1 almost surely as the number of observations increases. The *convergence rate* of the belief vector serves as a measure of learning performance for different SL variants. For example, the geometric rule provides a larger convergence rate than the arithmetic rule, as shown in [13]. Particularly, it is revealed in [8]–[11] that the convergence rate is related to the *combination policy* employed by the agents when implementing the pooling rule. An important conclusion from these studies is that placing the agent with the most evidence from local observations in the most centralized position within the network is beneficial for learning.

However, the convergence rate of belief vectors is not the only performance measure for a learning strategy. In statistical decision theory [14], [15], the *probability of error* is a key criterion for evaluating learning performance. This criterion is considered in adaptive social learning under non-stationary environments [12], where the probability of error remains non-zero in the steady state. The effect of combination policies on the steady-state probability of error in the context of adaptive SL is explored in [16]. Here, we consider non-adaptive SL rules [6]–[11]. Since they are uniformly consistent, the probability of error will converge to zero. What is the decaying rate of the probability of error in this context? Is this decaying rate affected by the combination policy? If so, how does this influence differ from its effect on the convergence rate of belief vectors? Moreover, can we optimize the combination policy to improve the convergence rate of the probability of error? These are the main questions addressed in this paper.

In this work, we will focus on the geometric rule due to its strong behavioral foundation [4], [11], [12]. Using large deviations theory [17], [18], it has been shown in [19] that in the asymptotic regime, the probability of error under the geometric SL rule decays exponentially with a rate–known as the *error exponent* in the literature–associated with the *Perron vector* of the combination policy. Building on this, we establish lower and upper bounds on the error exponents that can be achieved by all combination policies. Furthermore, we provide the criterion for designing the combination policy to reach the error exponent's upper bound, should it be feasible for the given learning task. It turns out that the optimal Perron entries must satisfy specific proportionality conditions related to the agent's local likelihood models. This indicates that the optimal combination strategy for minimizing the probability of error strikes a balance among the centralities of *all* informative agents, rather than merely increasing the centrality of the *most* informative agent, as suggested in prior studies [8]–[11] aimed at improving the convergence speed of belief vectors.

Notation: We use boldface fonts to denote random variables, and normal fonts for their realizations, e.g., $\boldsymbol{x}$ and x. $\mathbb{E}$ and $\mathbb{P}$ denote expectation and probability operators, respectively.

II. Problem Formulation

A. Observation Models

We consider a group of K agents, each receiving a stream of private observations over time. These observations are statistically governed by an *unknown* true state $\theta^\star$, which belongs to a finite set of M possible hypotheses. The sets of the K agents and the M hypotheses are denoted by $\mathcal{K} \triangleq \{1, \ldots, K\}$ and $\Theta \triangleq \{\theta_1, \ldots, \theta_M\}$, respectively. Specifically, at each time i, each agent k observes a private signal $\boldsymbol{\xi}_{k,i} \in \mathcal{X}_k$, which is an independent and identically distributed (i.i.d.) realization of a random variable described by the distribution $L_k(\cdot|\theta^\star)$. This distribution is drawn from a family of likelihood models $\{L_k(\cdot|\theta) : \theta \in \Theta\}$ specific to agent k, with all models defined on the same support $\mathcal{X}_k$. The agents are *heterogeneous* in that their signal spaces $\mathcal{X}_k$ and likelihood models $L_k(\cdot|\theta)$ may be different. We assume the observations from different agents to be independent conditioned on the true hypothesis $\theta^\star$. Without loss of generality, we assume that $\theta^\star = \theta_1$.

The goal of the agents is to infer the true hypothesis that underlies their observations. For each agent k, two distinct hypotheses θ_m and θ_n are called *locally indistinguishable* if $L_k(\cdot|\theta_m) = L_k(\cdot|\theta_n)$ almost everywhere. Equivalently, this can be expressed in terms of the Kullback-Leibler (KL) divergence [20] as $D_{\mathsf{KL}}(L_k(\cdot|\theta_m)\|L_k(\cdot|\theta_n)) = 0$. The learning task is infeasible for agent k if θ_1 is locally indistinguishable from any other θ. We impose the following identifiability condition.

Assumption 1. *For each pair of hypotheses (θ_1, θ), there exists at least one agent k such that $D_{\mathsf{KL}}(L_k(\cdot|\theta_1)\|L_k(\cdot|\theta)) > 0$.* □

This assumption guarantees that θ_1 is *globally* distinguishable for the network. However, the agents may not be able to learn the true hypothesis on their own due to potential local infeasibility. This necessitates collaboration among agents using certain learning rules, which we describe next.

B. Social Learning Rule

At each time i, each agent k maintains a local belief vector $\boldsymbol{\mu}_{k,i}$ and an intermediate belief vector $\boldsymbol{\psi}_{k,i}$, both of which are probability mass functions over the hypothesis space Θ. The entry $\boldsymbol{\mu}_{k,i}(\theta)$ of the belief vector represents agent k's confidence that θ is the true hypothesis. The social learning rule consists of two alternating steps. In the adaptation step, each agent k updates its intermediate belief $\boldsymbol{\psi}_{k,i}$ according to Bayes' rule using the new observation $\boldsymbol{\xi}_{k,i}$:

$$\boldsymbol{\psi}_{k,i}(\theta) = \frac{\boldsymbol{\mu}_{k,i-1}(\theta) L_k(\boldsymbol{\xi}_{k,i}|\theta)}{\sum_{\theta' \in \Theta} \boldsymbol{\mu}_{k,i-1}(\theta') L_k(\boldsymbol{\xi}_{k,i}|\theta')}, \quad \forall \theta \in \Theta. \tag{1}$$

In the combination step, agent k combines the intermediate beliefs from other agents to update its belief vector $\boldsymbol{\mu}_{k,i}$:

$$\boldsymbol{\mu}_{k,i}(\theta) = \frac{\prod_{\ell \in \mathcal{N}_k} \boldsymbol{\psi}_{\ell,i}(\theta)^{a_{\ell k}}}{\sum_{\theta' \in \Theta} \prod_{\ell \in \mathcal{N}_k} \boldsymbol{\psi}_{\ell,i}(\theta')^{a_{\ell k}}}, \quad \forall \theta \in \Theta \tag{2}$$

where $a_{\ell k}$ is the combination weight assigned to agent ℓ by agent k, and $\mathcal{N}_k$ denotes the set of neighboring agents that send information to agent k. The communication network among agents can be represented by a graph, which is assumed to be strongly connected.

Assumption 2. *There exist paths with positive weights between any two distinct agents in both directions (the two paths need not be the same), and at least one agent k has a self-loop ($a_{kk} > 0$).* □

The *combination policy* used by the agents is described by matrix $A = [a_{\ell k}]$, which is assumed to be left-stochastic with

$$A^\top \mathbb{1} = \mathbb{1}, \quad a_{\ell k} > 0, \ \forall \ell \in \mathcal{N}_k \text{ and } a_{\ell k} = 0, \ \forall \ell \notin \mathcal{N}_k \tag{3}$$

where $\mathbb{1}$ denotes the K-dimensional vector of all ones. Under Assumption 2, the matrix A is primitive so that the Perron eigenvector π exists according to the Perron-Frobenius theorem [21], which can be normalized to have strictly positive entries:

$$\pi^\top \mathbb{1} = 1, \quad \pi_k > 0, \ \forall k \in \mathcal{K}. \tag{4}$$

To avoid trivial cases, we assume that the initial belief of each agent on each hypothesis is positive.

Assumption 3. *For each $k \in \mathcal{K}$ and $\theta \in \Theta$, $\boldsymbol{\mu}_{k,0}(\theta) > 0$.* □

C. Decision Making Rule

At each time i, the agents need to make a decision about the true hypothesis based on their current belief vectors. One natural choice for them is to select the hypothesis with the highest belief. According to this rule, a decision error occurs at agent k if $\boldsymbol{\mu}_{k,i}(\theta_1)$ is either not the maximum or is the maximum but not unique. The instantaneous *probability of error* of agent k at time i, denoted by $p_{k,i}$, is defined as

$$p_{k,i} \triangleq \mathbb{P}\big(\exists \theta \neq \theta_1 : \boldsymbol{\mu}_{k,i}(\theta) \geq \boldsymbol{\mu}_{k,i}(\theta_1)\big). \tag{5}$$

The convergence behavior of $p_{k,i}$ obviously depends on that of the belief vectors. From [9], [10], we know that $\boldsymbol{\mu}_{k,i}(\theta_1) \to 1$ almost surely with the exponential convergence rate:

$$I(\pi, \Theta) \triangleq \min_{\theta \neq \theta_1} \sum_{k=1}^{K} \pi_k D_{\mathsf{KL}}(L_k(\cdot|\theta_1)\|L_k(\cdot|\theta)). \tag{6}$$

Due to Assumption 1, it holds $I(\pi, \Theta) > 0$ for any π satisfying (4). This implies from (5) that $p_{k,i} \to 0$ as i grows. It is clear from (6) that the Perron centrality of each agent plays a crucial role in the convergence rate of the belief vectors. Particularly, a combination policy that places more centrality on the most influential agent (measured by the KL divergence) is preferred as it boosts the rate of convergence.

However, an increase in the convergence speed of the belief vector $\boldsymbol{\mu}_{k,i}$ does not equate to a decrease in the probability of error $p_{k,i}$, which is another important performance measure in statistical decision theory [14], [15]. The main goal of this work is to examine the effect of the combination policy on the convergence rate of the probability of error $p_{k,i}$, with particular focus on the exponential decaying rate–referred to as the *error exponent*–in the asymptotic regime as i becomes large.

III. MAIN RESULTS

For our subsequent analysis of the error exponent, we need to introduce some useful variables. For each agent k and each time i, we define $\boldsymbol{x}_{k,i}(\theta)$ and $\boldsymbol{\lambda}_{k,i}(\theta)$ as the log-likelihood and log-belief ratios associated with the true hypothesis θ_1 and an alternative hypothesis $\theta \neq \theta_1$:

$$\boldsymbol{x}_{k,i}(\theta) \triangleq \log \frac{L_k(\boldsymbol{\xi}_{k,i}|\theta_1)}{L_k(\boldsymbol{\xi}_{k,i}|\theta)}, \quad \boldsymbol{\lambda}_{k,i}(\theta) \triangleq \log \frac{\boldsymbol{\mu}_{k,i}(\theta_1)}{\boldsymbol{\mu}_{k,i}(\theta)}. \quad (7)$$

We also introduce the logarithmic moment generating function (LMGF) of $\boldsymbol{x}_{k,i}(\theta)$:

$$\Lambda_{k,\theta}(t) \triangleq \log \mathbb{E} e^{t\boldsymbol{x}_{k,i}(\theta)} = \log \mathbb{E} \exp\left\{ t \log \frac{L_k(\boldsymbol{\xi}_{k,i}|\theta_1)}{L_k(\boldsymbol{\xi}_{k,i}|\theta)} \right\}. \quad (8)$$

We make the following assumption on $\Lambda_{k,\theta}(t)$, which is known as Cramér's condition [4], [22].

Assumption 4. *For each $k \in \mathcal{K}$, $\Lambda_{k,\theta}(t) < \infty, \forall t \in \mathbb{R}$ for all $\theta \neq \theta_1$.* □

The network LMGF $\Lambda_{\mathsf{ave},\theta}(\pi, t)$ is defined as the average of the individual LMGFs $\Lambda_{k,\theta}(t)$ weighted by the Perron vector π:

$$\Lambda_{\mathsf{ave},\theta}(\pi, t) \triangleq \sum_{k=1}^{K} \Lambda_{k,\theta}(\pi_k t). \quad (9)$$

Using the well-known Gärtner-Ellis theorem [17], [18], it was established in [19] that the probability of error $p_{k,i}$ satisfies the large deviations principle (LDP) with the rate function given by the Fenchel-Legendre transform $\phi_{\mathsf{ave},\theta}(\pi, x)$ of $\Lambda_{\mathsf{ave},\theta}(\pi, t)$:

$$\phi_{\mathsf{ave},\theta}(\pi, x) = \sup_{t\in\mathbb{R}} \Big[tx - \Lambda_{\mathsf{ave},\theta}(\pi, t) \Big], \quad \forall x \in \mathbb{R}. \quad (10)$$

Lemma 1 ([19])**.** *Under Assumptions 1–4, the probabilities of error of all agents satisfy the LDP with the error exponent*

$$\Phi(\pi) \triangleq \min_{\theta\neq\theta_1} \phi_{\mathsf{ave},\theta}(\pi, 0) > 0 \quad (11)$$

for any given combination policy with Perron vector π. □

According to LDP theory, this means that the probability of error $p_{k,i}$ can be characterized as

$$p_{k,i} \doteq e^{-i\Phi(\pi)}, \quad \text{as } i \to \infty \quad (12)$$

where the notation $\doteq$ means equality to the leading order (namely, i) in the exponent. Therefore, a larger error exponent $\Phi(\pi)$ indicates a faster decaying rate of $p_{k,i}$ in the asymptotic regime. Since $\Phi(\pi)$ depends on the Perron vector π, a pertinent question is whether we can find an optimal Perron vector π that maximizes the error exponent. This leads to the following optimization problem:

$$\max_{\pi} \Phi(\pi) \quad \text{s.t. (4)}. \quad (13)$$

Before solving this problem, we first derive bounds on the achievable error exponents for the SL rule (1)–(2).

For each hypothesis $\theta \neq \theta_1$, we denote by $\mathcal{K}_I(\theta)$ the set of agents for whom θ is distinguishable from θ_1, i.e.,

$$\mathcal{K}_I(\theta) \triangleq \big\{ k \in \mathcal{K} : D_{\mathsf{KL}}(L_k(\cdot|\theta_1) \| L_k(\cdot|\theta)) > 0 \big\}. \quad (14)$$

Under Assumption 1, we have $\mathcal{K}_I(\theta) \neq \emptyset$ for all $\theta \neq \theta_1$. Let $\Phi_k(\theta)$ denote the error exponent associated with hypothesis θ for agent k in the non-cooperative learning scenario where

$$\Theta^{\mathsf{nc}} \triangleq \{\theta_1, \theta\}, \quad \mathcal{K}^{\mathsf{nc}} \triangleq \{k\}. \quad (15)$$

By replacing $(\Theta, \mathcal{K})$ with $(\Theta^{\mathsf{nc}}, \mathcal{K}^{\mathsf{nc}})$ in Lemma 1, we obtain

$$\Phi_k(\theta) \triangleq \phi_{k,\theta}(0) = -\inf_{t\in\mathbb{R}} \Lambda_{k,\theta}(t), \quad (16)$$

where $\phi_{k,\theta}(x)$ is the conjugate of $\Lambda_{k,\theta}(t)$ defined in a similar way to (10). Moreover, we have $\Phi_k(\theta) > 0$ if $k \in \mathcal{K}_I(\theta)$ and $\Phi_k(\theta) = 0$ otherwise. A bound on the error exponent $\Phi(\pi)$ can be established using the individual $\Phi_k(\theta)$ as follows.

Theorem 1. *The error exponent $\Phi(\pi)$ is bounded as*

$$\min_{\theta\neq\theta_1} \min_{k\in\mathcal{K}_I(\theta)} \Phi_k(\theta) \leq \Phi(\pi) \leq \min_{\theta\neq\theta_1} \sum_{k=1}^{K} \Phi_k(\theta). \quad (17)$$

Proof. Omitted for brevity, however, the argument follows by using similar tools as in [16]. □

This theorem shows that the maximum error exponent achievable for the optimization problem (13) is upper bounded by the right-hand side term of (17), which we denote as $\bar{\Phi}$. Also, we denote the set of hypotheses corresponding to $\bar{\Phi}$ by

$$\Theta^* \triangleq \arg\min_{\theta\neq\theta_1} \sum_{k=1}^{K} \Phi_k(\theta). \quad (18)$$

This set contains all hypotheses that are hardest to distinguish from the true hypothesis θ_1 for the entire network (in terms of the error exponents). The subsequent question is whether we can attain this upper bound by optimizing π. We will answer this question in the affirmative by showing that the optimal Perron vector $\pi^\star$ needs to satisfy a proportionality condition related to the quantity $t_k(\theta)$, which is defined as the critical t that determines the individual error exponent $\Phi_k(\theta)$ in (16):

$$t_k(\theta) \in \arg\min_{t\in\mathbb{R}} \Lambda_{k,\theta}(t). \quad (19)$$

Using the convexity of the LMGF $\Lambda_{k,t}(\theta)$, we can prove that $t_k(\theta)$ is unique and negative for all $k \in \mathcal{K}_I(\theta)$, that is,

$$t_k(\theta) \triangleq \arg\min_{t\in\mathbb{R}} \Lambda_{k,\theta}(t) < 0, \quad \forall k \in \mathcal{K}_I(\theta). \quad (20)$$

By definition of $\mathcal{K}_I(\theta)$ in (14), we know $\Lambda_{k,\theta}(t) = 0, \forall t \in \mathbb{R}$ for each $k \notin \mathcal{K}_I(\theta)$. In this case, $t_k(\theta)$ can be any value.

Theorem 2. *Let $\Phi^\star$ be the maximum error exponent in the optimization problem* (13) *and $\Pi^\star$ be the set of optimal Perron vectors. If the upper bound in* (17) *is achievable, i.e., $\Phi^\star = \bar{\Phi}$, then for each Perron vector $\pi^\star \in \Pi^\star$, the following relation*

$$\frac{\pi_k^\star}{\pi_\ell^\star} = \frac{t_k(\theta^*)}{t_\ell(\theta^*)}, \quad \forall k, \ell \in \mathcal{K}_I(\theta^*) \quad (21)$$

holds for all hypotheses $\theta^* \in \Theta^*$.

Proof. This proof is based on the nice properties of the LMGF $\Lambda_{\mathsf{ave},\theta}(\pi, t)$ and the rate function $\phi_{\mathsf{ave},\theta}(\pi, x)$, proceeding in a manner similar to the analysis conducted by [16]. □

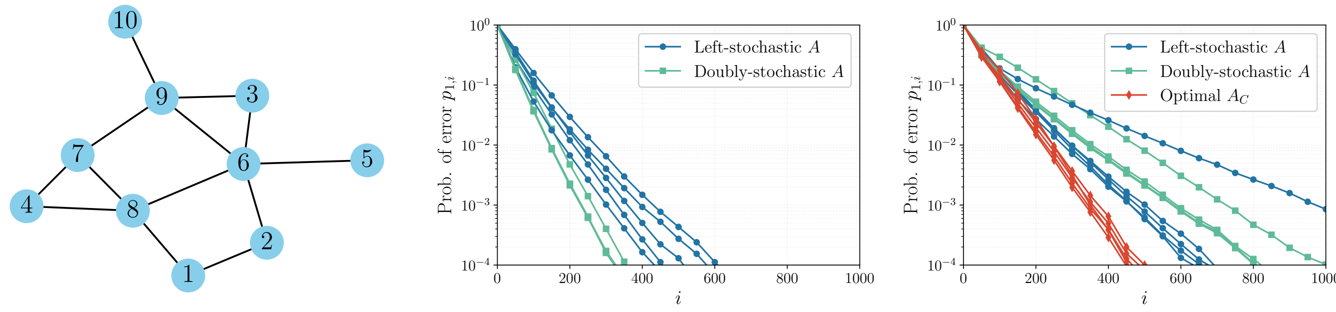

(a) Network structure (b) Error probability (noiseless scenario) (c) Error Probability (noisy scenario)

Fig. 1: Probabilities of error under different combination policies in the two learning scenarios on a given network.

TABLE I: Parameters of Gaussian models

Agent k	$\mathsf{m}_k(\theta_1)$	$\mathsf{m}_k(\theta_2)$	$\mathsf{m}_k(\theta_3)$	σ_k^2	ε_k
1–3	0	0.1	0.1	1	10
4–7	0	0.2	0	1	0.1
8–10	0	0	0.3	1	0.01

Remark 1. *Assume $\Theta^* = \{\theta^*\}$ and that the upper bound $\bar{\Phi}$ is attainable. There are two special cases of $\Pi^\star$ when $\mathcal{K}_I(\theta^*)$ takes two extreme values: $\mathcal{K}_I(\theta^*) = \mathcal{K}$ or $\mathcal{K}_I(\theta^*) = \{k\}$. In the first case, all agents are informative in distinguishing θ^* from θ_1. From* (20) *and* (21)*, we know that the optimal Perron vector $\pi^\star$ is unique, i.e., $\Pi^\star = \{\pi^\star\}$, with*

$$\pi_k^\star = \frac{t_k(\theta^*)}{\sum_{\ell\in\mathcal{K}} t_\ell(\theta^*)}, \quad \forall k \in \mathcal{K}. \tag{22}$$

In the second case, only agent k is able to exclude the wrong hypothesis θ^* based on its likelihood models. The condition (21) *holds regardless of the Perron centrality π_k. Therefore, $\Pi^\star$ includes all Perron vectors π that attain the upper bound. It is worth noting that this conclusion contradicts the existing results in [9]–[11], which support making π_k as large as possible for improving the convergence rate of belief vectors.*

IV. Numerical Simulations

In the simulations, we consider a network of 10 agents shown in Fig. 1a, where each edge is generated with probability 0.1. The network is constructed to be undirected to facilitate the design of the combination policy for a given Perron vector. Each agent is assumed to have a self-loop, which is not shown in the figure. We cluster the 10 agents into 3 groups such that the agents in the same group share the same likelihood models.

We consider 3 hypotheses and assume that each agent owns a family of Gaussian models, i.e., $L_k(\cdot|\theta) = G(m_k(\theta), \sigma_k^2(\theta))$, where $G(m, \sigma^2)$ denotes the Gaussian distribution with mean m and variance σ^2. The group assignment and the parameters of the Gaussian models for each group are provided in Table I. For two Gaussian models $G(m_k(\theta_1), \sigma_k^2)$ and $G(m_k(\theta), \sigma_k^2)$ with the same variance σ_k^2, the critical t in (20) is computed as $t_k(\theta) = -\frac{1}{2}$. Since for any $\theta \neq \theta_1$, $t_k(\theta)$ is identical for all $k \in \mathcal{K}_I(\theta)$, the uniform Perron vector satisfies (21). We can prove that it is an optimal Perron vector. To validate its optimality, we compare five doubly-stochastic A with a uniform Perron vector to five left-stochastic A with sub-optimal Perron vectors. All matrices are generated randomly following the procedure in [16]. Under different A, the probability of error $p_{1,i}$ of agent 1 is shown in Fig. 1b. It can be observed that doubly-stochastic matrices provide a larger error exponent.

In the second scenario, we assume that the observations $\boldsymbol{\xi}_{k,i}$ are corrupted by some noise $\boldsymbol{n}_{k,i} \sim G(0, \nu_k^2)$. The noise level ε_k at agent k is defined as the ratio between ν_k^2 and σ_k^2, and its value is given in Table I. In this noisy scenario, we have

$$t_k(\theta) = -\frac{1}{2(1+\varepsilon_k)}, \; \Phi_k(\theta) = \frac{(m_k(\theta_1) - m_k(\theta))^2}{8\sigma_k^2(1+\varepsilon_k)}. \tag{23}$$

From Table I, we obtain $\Theta^* = \{\theta_2\}$ and $\mathcal{K}_I(\theta_2) = \{1, \dots, 7\}$. The upper bound $\bar{\Phi}$ of the error exponents is achievable for this noisy learning setup by the Perron vector constructed in (22) with $t_k(\theta_2)$ given by (23). Since the proportionality condition (21) is imposed only on agents belonging to $\mathcal{K}_I(\theta_2)$, the set $\Pi^\star$ may not be a singleton. To illustrate this, we assume

$$t_k(\theta_2) = -C < 0, \quad \forall k \notin \mathcal{K}_I(\theta_2), \tag{24}$$

and consider the Perron vector π_C given by (22) using (24). It can be proved that $\pi_C \in \Pi^\star$ for any choice of $C \geq 0.045$. In the simulations, we construct five combination policies A_C with Perron vectors π_C by selecting C to be 0.2, 0.4, 0.6, 0.8, and 1. Together with the ten combination policies in Fig. 1b, we plot the evolution of $p_{1,i}$ under these fifteen combination policies in Fig. 1c. The benefit of employing optimal combination policies is evident, as the curves corresponding to each A_C lie below all other curves in Fig. 1c.

V. Conclusions

We investigate the influence of the combination policies on the error exponent of the probability of error in social learning. The bound on the error exponent is derived, followed by an error exponent maximization problem. Our results show that by employing a combination policy with an optimal Perron vector, which satisfies certain proportionality conditions related to the local likelihood models, the agents can learn the true hypothesis more efficiently.

REFERENCES

[1] D. Acemoglu and A. Ozdaglar, "Opinion dynamics and learning in social networks," *Dynamic Games and Applications*, vol. 1, pp. 3–49, 2011.

[2] C. Chamley, A. Scaglione, and L. Li, "Models for the diffusion of beliefs in social networks: An overview," *IEEE Signal Processing Magazine*, vol. 30, no. 3, pp. 16–29, 2013.

[3] 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 Processing Magazine*, vol. 30, no. 3, pp. 43–57, 2013.

[4] V. Bordignon, V. Matta, and A. H. Sayed, *Social Learning: Opinion Formation over Graphs*. NOW-EURASIP Open Access Book Series on Information and Learning Sciences, 2024.

[5] G. Koliander, Y. El-Laham, P. M. Djurić, and F. Hlawatsch, "Fusion of probability density functions," *Proceedings of the IEEE*, vol. 110, no. 4, pp. 404–453, 2022.

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

[7] X. Zhao and A. H. Sayed, "Learning over social networks via diffusion adaptation," in *Proc. Asilomar Conference on Signals, Systems and Computers*, Pacific Grove, USA, 2012, pp. 709–713.

[8] A. Jadbabaie, P. Molavi, and A. Tahbaz-Salehi, "Information heterogeneity and the speed of learning in social networks," *Columbia Business School Research Paper*, no. 13-28, 2013.

[9] A. Lalitha, T. Javidi, and A. D. Sarwate, "Social learning and distributed hypothesis testing," *IEEE Transactions on Information Theory*, vol. 64, no. 9, pp. 6161–6179, 2018.

[10] A. Nedic, A. Olshevsky, and C. A. Uribe, "Fast convergence rates for distributed non-Bayesian learning," *IEEE Transactions on Automatic Control*, vol. 62, no. 11, pp. 5538–5553, 2017.

[11] P. Molavi, A. Tahbaz-Salehi, and A. Jadbabaie, "A theory of non-Bayesian social learning," *Econometrica*, vol. 86, no. 2, pp. 445–490, 2018.

[12] V. Bordignon, V. Matta, and A. H. Sayed, "Adaptive social learning," *IEEE Transactions on Information Theory*, vol. 67, no. 9, pp. 6053–6081, 2021.

[13] M. Kayaalp, Y. İnan, E. Telatar, and A. H. Sayed, "On the arithmetic and geometric fusion of beliefs for distributed inference," *IEEE Transactions on Automatic Control*, vol. 69, no. 4, pp. 2265–2280, 2024.

[14] J. O. Berger, *Statistical Decision Theory and Bayesian Analysis*. Springer New York, NY, 2013.

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

[16] P. Hu, V. Bordignon, S. Vlaski, and A. H. Sayed, "Optimal aggregation strategies for social learning over graphs," *IEEE Transactions on Information Theory*, vol. 69, no. 9, pp. 6048–6070, 2023.

[17] A. Dembo and O. Zeitouni, *Large Deviations Techniques and Applications*. Springer, 1998.

[18] F. Den Hollander, *Large Deviations*. American Mathematical Society, 2008.

[19] V. Bordignon, M. Kayaalp, V. Matta, and A. H. Sayed, "Social learning with non-Bayesian local updates," in *2023 31st European Signal Processing Conference (EUSIPCO)*, 2023, pp. 1–5.

[20] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. NJ, USA: Wiley, 1991.

[21] A. H. Sayed, "Adaptation, learning, and optimization over networks," *Foundations and Trends in Machine Learning*, vol. 7, no. 4-5, pp. 311–801, 2014.

[22] Y. Polyanskiy and Y. Wu, *Information Theory: From Coding to Learning*. Cambridge University Press, 2024.