Social Learning with Beliefs in a Parallel Network

Daewon Seo
University of Wisconsin-Madison

Ravi Kiran Raman
Analog Devices

Lav R. Varshney
University of Illinois at Urbana-Champaign

***Abstract*—Consider a social learning problem in a parallel network, where N distributed agents make independent selfish binary decisions, and a central agent aggregates them together with a private signal to make a final decision. In particular, all agents have private beliefs for the true prior, based on which they perform binary hypothesis testing. We focus on the Bayes risk of the central agent, and counterintuitively find that a collection of agents with incorrect beliefs could outperform a set of agents with correct beliefs. We also consider many-agent asymptotics (i.e., N is large) when distributed agents all have identical beliefs, for which it is found that the central agent's decision is polarized and beliefs determine the limit value of the central agent's risk. Moreover, it is surprising that when all agents believe a certain prior-agnostic constant belief, it achieves globally optimal risk as $N \to \infty$.**

I. Introduction

When individuals are asked to make a decision, they often consider the decisions made by others (e.g., online reviews) in addition to their own assessment, cf. [1]. With technology-mediated social influence becoming much more prevalent, there is growing interest in understanding *social wisdom* or *social learning* from a theoretical perspective. Social learning, often referred to as *observational learning*, is such a scenario where individuals interact and learn from others' decisions as well as their own private signal. Here, we study a Bayesian social learning problem in a parallel network, where N distributed agents make decisions to minimize their own Bayes risk, and the decisions are sent to the central agent. The central agent aggregates these N prior decisions and its own private signal to make a decision, e.g., whether to buy or not.

Social learning has been widely studied by many communities with different flavors. In economics, a seminal result is so-called *information cascade* [2]–[4] for a tandem of agents, where the agents observe the history of decisions made by preceding agents. As a result of Bayesian decision making, an information cascade occurs, i.e., agents after some point ignore their private signal and herd on the previous agent's (possibly incorrect) decision. Herding happens due to bounded informativeness of private signals (such as a binary signal), which are not sufficiently informative to counter biased prior decisions [5]. There are a variety of extensions to the basic social learning setting, for example, over networks [6] and with noisy history [7].

Another line of work is in *distributed inference*, where a central fusion agent collects local decisions from distributed agents and makes a final decision [8], [9]. The link between

This work was supported in part by the National Science Foundation under grant CCF-1717530.

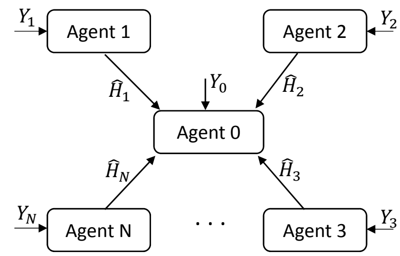

Fig. 1. The parallel network model.

the distributed nodes and the fusion center could be rate-limited [10], imperfect [11], [12], or with memory [13]. It is also common to consider learning behavior and study its convergence speed. The simplest setting is a tandem network, also called serial detection, [14]–[16]. For a general network, every vertex agent in a network can identify the unknown hypothesis by repeating local belief exchanges [17]–[21].

In our previous research [22], we studied a tandem of agents that have private prior beliefs on the hypothesis that are not necessarily identical to the true prior, i.e., each agent has a perceived belief of the prior. Focusing on the Bayes risk of the last agent of the tandem, one might have thought that beliefs identical to the prior would achieve the smallest Bayes risk, since prior decisions are locally Bayes-optimal and the last agent does not misunderstand them. However, we found that a certain combination of incorrect beliefs achieves smaller Bayes risk. Here we consider a parallel network with the same setting—each agent has a perceived belief of the prior, and the focus is on the Bayes risk of the central agent. As will be seen, a certain combination of incorrect beliefs outperforms the case of agents all having the true prior. Moreover, assuming homogeneous distributed agents with identical beliefs and focusing on asymptotics when $N \to \infty$, we further find that the central agent makes a certain decision with probability 1. Surprisingly, it is asymptotically optimal that all agents have the belief such that decision thresholds are exactly the middle between two hypotheses, e.g., all agents believe both hypotheses are equally likely if costs are equal.

II. Problem Description and Belief Update

A. Problem Model

Consider a parallel network, depicted in Fig. 1, consisting of N distributed agents and a single central agent, denoted as agent 0. The underlying binary hypothesis, $H \in \{0, 1\}$, follows the prior $\mathbb{P}[H=0]=p_0$ and $\mathbb{P}[H=1]=\bar{p}_0 \triangleq 1-p_0$, which is unknown to the agents. Instead of the unknown p_0, each agent $i \in \{0,1,\dots,N\}$ believes q_i is the true prior. Each agent receives the private signal $Y_i = H + Z_i$, where Z_i is taken as an independent standard Gaussian noise for brevity of presentation. We assume that correct decisions incur no cost and the costs for false alarm (or type I error, i.e., choosing $\widehat{H}=1$ when $H=0$) and missed detection (or type II error, i.e., choosing $\widehat{H}=0$ when $H=1$) are c_{FA} and c_{MD}, respectively. In addition, we assume that all agents share the same costs so they are a team in the sense of Radner [23]. Agents are Bayes-rational so make decisions that minimize perceived Bayes risk, i.e.,

$$R_{i,[i]} = c_{\mathsf{FA}} q_i p_{\widehat{H}_i|H}(1|0)_{[i]} + c_{\mathsf{MD}}(1-q_i) p_{\widehat{H}_i|H}(0|1)_{[i]}, \quad (1)$$

where subscript $[i]$ indicates quantities 'seen' by agent i as if q_i is the true prior. When the quantity does not have $[i]$, it implies the quantity seen by an oracle aware of $(p_0, q_0, q_1, \dots, q_N)$.

To simplify notation, we use $x^N = (x_1, \dots, x_N)$ to denote a tuple of length N, and $x_{-i}^N = (x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_N)$ to denote the tuple excluding the ith element. All logarithms are natural logarithms. We use p, f to denote probability mass functions and probability density functions, respectively. $Q(x)$ is defined to be the complementary cumulative distribution function of the standard Gaussian,

$$Q(x) = \int_x^\infty \phi(t;0)dt,$$

where $\phi(\cdot;\mu)$ is the probability density function of Gaussian with mean μ and unit variance.

B. Belief Update

It is easy to see that the likelihood ratio test (LRT) as if q_i is the true prior minimizes (1), that is, for $i \in \{1,\dots,N\}$, the following test minimizes $R_{i,[i]}$:

$$\frac{f_{Y_i|H}(y_i|1)}{f_{Y_i|H}(y_i|0)} \underset{\widehat{H}_i=0}{\overset{\widehat{H}_i=1}{\gtrless}} \frac{c_{\mathsf{FA}} q_i}{c_{\mathsf{MD}}(1-q_i)}. \quad (2)$$

Noting that $f_{Y_i|H}(y_i|h)$ is Gaussian with mean h and unit variance, (2) can be simplified to decision threshold $\lambda_i \triangleq \lambda(q_i)$,

$$y_i \underset{\widehat{H}_i=0}{\overset{\widehat{H}_i=1}{\gtrless}} \lambda(q_i) \triangleq \frac{1}{2} + \log\left(\frac{c_{\mathsf{FA}} q_i}{c_{\mathsf{MD}}(1-q_i)}\right). \quad (3)$$

Therefore for distributed agents, the conditional error probabilities are

$$p_{\widehat{H}_i|H}(1|0) = \int_{\lambda_i}^\infty \phi(t;0)dt = Q(\lambda_i),$$

$$p_{\widehat{H}_i|H}(0|1) = \int_{-\infty}^{\lambda_i} \phi(t;1)dt = 1 - Q(\lambda_i - 1) = Q(1-\lambda_i),$$

where the last equality follows from the property that $Q(x) = 1 - Q(-x)$.

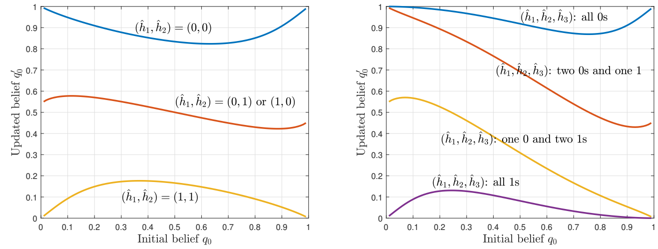

(a) $N=2$ (b) $N=3$

Fig. 2. Updated belief for possible decisions.

The central agent with belief q_0 has access to all decisions made by distributed agents, so its LRT, given $(y_0, \hat{h}_1, \dots, \hat{h}_N)$ is

$$\frac{f_{Y_0,\widehat{H}^N|H}(y_0,\hat{h}^N|1)}{f_{Y_0,\widehat{H}^N|H}(y_0,\hat{h}^N|0)} \underset{\widehat{H}_0=0}{\overset{\widehat{H}_0=1}{\gtrless}} \frac{c_{\mathsf{FA}} q_0}{c_{\mathsf{MD}}(1-q_0)}.$$

Since $Y_0, \widehat{H}_1, \dots, \widehat{H}_N$ are independent conditioned on H,

$$f_{Y_0,\widehat{H}^N|H}(y_0,\hat{h}^N|h) = f_{Y_0|H}(y_0|h) \prod_{i=1}^N p_{\widehat{H}_i|H}(\hat{h}_i|h).$$

Here $p_{\widehat{H}_i|H}$ is a function of q_i only, however, the central agent recognizes q_0 is the prior. Hence, the agent computes $p_{\widehat{H}_i|H}$ as if distributed agents performed hypothesis testing (2) with q_0. It leads to the following LRT[1]

$$\frac{f_{Y_0|H}(y_0|1)}{f_{Y_0|H}(y_0|0)} \underset{\widehat{H}_0=0}{\overset{\widehat{H}_0=1}{\gtrless}} \frac{c_{\mathsf{FA}} q_0}{c_{\mathsf{MD}}(1-q_0)} \prod_{i=1}^N \frac{p_{\widehat{H}_i|H}(\hat{h}_i|0)_{[0]}}{p_{\widehat{H}_i|H}(\hat{h}_i|1)_{[0]}}. \quad (4)$$

Since $x/(1-x)$ is monotonically increasing in $x \in (0,1)$, we can interpret (4) as a new LRT with updated belief q_0',

$$\frac{f_{Y_0|H}(y_0|1)}{f_{Y_0|H}(y_0|0)} \underset{\widehat{H}_0=0}{\overset{\widehat{H}_0=1}{\gtrless}} \frac{c_{\mathsf{FA}} q_0'}{c_{\mathsf{MD}}(1-q_0')}, \quad (5)$$

where q_0' is defined so that

$$\frac{q_0'}{1-q_0'} = \frac{q_0}{1-q_0} \prod_{i=1}^N \frac{p_{\widehat{H}_i|H}(\hat{h}_i|0)_{[0]}}{p_{\widehat{H}_i|H}(\hat{h}_i|1)_{[0]}}. \quad (6)$$

Finally, the true Bayes risk of the central agent is

$$R_0 = c_{\mathsf{FA}} p_0 p_{\widehat{H}_0|H}(1|0) + c_{\mathsf{MD}} \bar{p}_0 p_{\widehat{H}_0|H}(0|1), \quad (7)$$

with

$$p_{\widehat{H}_0|H}(\hat{h}_0|h) = \sum_{\hat{h}^N} p_{\widehat{H}^N,\widehat{H}_0|H}(\hat{h}^N,\hat{h}_0|h).$$

III. Results for Finite N

A. Belief Update

As stated, the central agent adopts the new LRT based on the updated belief q_0' as in (5). Fig. 2 depicts the updated belief q_0' in (6) for possible decisions for $N=2,3$. The

[1] Again, the subscript [0] denotes the value that the central agent thinks.

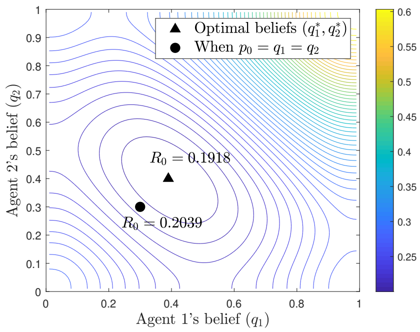

Fig. 3. Risk contour for $N = 2$ at $p_0 = 0.3$, $q_0 = 0.7372$, $c_{\mathsf{FA}} = c_{\mathsf{MD}} = 1$.

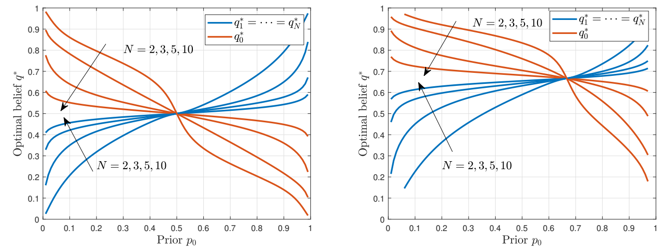

(a) $c_{\mathsf{FA}} = c_{\mathsf{MD}} = 1$

(b) $c_{\mathsf{FA}} = 1$, $c_{\mathsf{MD}} = 2$

Fig. 4. Optimal beliefs that minimize R_0 for several N. The curves for $N = 2, 3$ are found by exhaustive search, and curves for $N = 5, 10$ are by assuming $q_1 = q_2 = \cdots = q_N$.

curves indicate how observing local decisions changes the central agent's belief. In addition, q_0 changes significantly when local agent decisions differ from what the central agent expects. For example in Fig. 2(b), when q_0 is small the central agent believes H is highly likely to be 1. However, observing $(\widehat{h}_1, \widehat{h}_2, \widehat{h}_3) = (0, 0, 0)$, his updated belief approaches 1 so he now believes H is highly likely to be 0. On the other hand, observing $(\widehat{h}_1, \widehat{h}_2, \widehat{h}_3) = (1, 1, 1)$ he confirms the small q_0 and enhances it so $q_0' < q_0$ after $(\widehat{h}_1, \widehat{h}_2, \widehat{h}_3) = (1, 1, 1)$.

It is noteworthy that the updated belief curves are not monotonic in q_0 for each set of prior decisions. In a tandem network [22, Fig. 2 and Thm. 3], it is shown that the update equation (6) for $N = 1$ preserves the ordering of beliefs, i.e., the updated belief is always monotonically increasing in q_0. However, this is no longer true in the parallel case as illustrated in Fig. 2 when multiple local decisions are taken into account. This is because $q_0/(1 - q_0)$ is increasing in q_0, whereas $p_{\widehat{H}_i|H}(\widehat{h}_i|0)_{[0]}/p_{\widehat{H}_i|H}(\widehat{h}_i|1)_{[0]}$ is decreasing in q_0 for both $\widehat{h}_i = 0, 1$. So the reversal of ordering takes place when the multiplicative terms in the right side of (6) are strong enough to counter the increment of $q_0/(1 - q_0)$ term.

B. Optimal Beliefs

Following the LRTs (2) and (4), agents declare decisions that appear in R_0 according to (7). Clearly R_0 is a function of $(q_0, q_1, \ldots, q_N)$ for given p_0 and costs. One might think that R_0 achieves its minimum when each agents knows the true prior, i.e., at $p_0 = q_0 = q_1 = \cdots = q_N$, since distributed agents make the best decisions and the central agent does not misunderstand them. However, this turns out to be false.

Recall that local decisions are independent conditioned on H, which implies that $P_{\widehat{H}_0|H}(\widehat{h}_0|h)$ in (7) can be rewritten as

$$p_{\widehat{H}_0|H}(\widehat{h}_0|h) = \sum_{\widehat{h}^N} \left(\prod_{i=1}^{N} p_{\widehat{H}_i|H}(\widehat{h}_i|h) \right) p_{\widehat{H}_0|H,\widehat{H}^N}(\widehat{h}_0|h, \widehat{h}^N).$$

Therefore (7) can be expressed as

$$R_0 = c_{\mathsf{FA}} p_0 \sum_{\widehat{h}^N} \left(\prod_{i=1}^{N} p_{\widehat{H}_i|H}(\widehat{h}_i|0) \right) p_{\widehat{H}_0|H,\widehat{H}^N}(1|0, \widehat{h}^N)$$
$$+ c_{\mathsf{MD}} \bar{p}_0 \sum_{\widehat{h}^N} \left(\prod_{i=1}^{N} p_{\widehat{H}_i|H}(\widehat{h}_i|1) \right) p_{\widehat{H}_0|H,\widehat{H}^N}(0|1, \widehat{h}^N). \quad (8)$$

Theorem 1. *Let $(q_0^*, q_1^*, \ldots, q_N^*)$ be the optimal belief tuple that minimizes R_0. Then, the following necessary condition holds: $(q_0^*, q_1^*, \ldots, q_N^*)$ is the solution to*

$$\frac{q_j}{1 - q_j} = \frac{p_0}{1 - p_0} \frac{A_1^{(j)} - A_0^{(j)}}{B_0^{(j)} - B_1^{(j)}}, \quad \forall j \in \{1, \ldots, N\} \quad (9)$$

where

$$A_h^{(j)} = \sum_{\widehat{h}_{-j}^N} \left(p_{\widehat{H}_i|H}(\widehat{h}_i|0) \right) p_{\widehat{H}_0|H,\widehat{H}_{-j}^N,\widehat{H}_j}(1|0, \widehat{h}_{-j}^N, h),$$
$$B_h^{(j)} = \sum_{\widehat{h}_{-j}^N} \left(p_{\widehat{H}_i|H}(\widehat{h}_i|1) \right) p_{\widehat{H}_0|H,\widehat{H}_{-j}^N,\widehat{H}_j}(0|1, \widehat{h}_{-j}^N, h).$$

Proof: Differentiating (8) with respect to decision threshold λ_j and rearranging terms give the claim. See [24] for details. ■

Quantities $A_h^{(j)}, B_h^{(j)}$ are the false alarm and missed detection probabilities of the central agent conditioned on $\widehat{h}_j = h$, therefore independent of q_j. Thm. 1 can be thought of as a balance condition that the optimal initial beliefs must satisfy between error probabilities. Clearly, the value $\frac{A_1^{(j)} - A_0^{(j)}}{B_0^{(j)} - B_1^{(j)}}$ is not 1 in general, thus, $q_i^* \neq p_0$ in general. Fig. 3 illustrates optimal beliefs for $N = 2$, $c_{\mathsf{FA}} = c_{\mathsf{MD}} = 1$, and $p_0 = 0.3$. The central agent's initial belief is given by $q_0 = 0.7372$, at which R_0 attains its minimum from Fig. 4. Clearly, biased beliefs $q_1 = q_2 = 0.3960$ with $R_0 = 0.1918$ outperforms context-aware distributed agents $p_0 = q_1 = q_2$ with $R_0 = 0.2039$. Also note that when $p_0 = q_1 = q_2 = q_0$ (not shown in Fig. 3), it gives $R_0 = 0.1976$, strictly worse. Another interesting implication of Fig. 4 is that optimal beliefs become closer to $\frac{c_{\mathsf{MD}}}{c_{\mathsf{FA}} + c_{\mathsf{MD}}}$ as N grows for the entire range of p_0. It suggests that setting $q_i = \frac{c_{\mathsf{MD}}}{c_{\mathsf{FA}} + c_{\mathsf{MD}}}$ for all $i \in \{0, 1, \ldots, N\}$ would be asymptotically optimal as N grows. This will be rigorously proven in Thm. 5 in the risk exponent sense.

The global optimization problem for R_0 belongs to neither a convex class nor any analytically solvable classes, as far as we know. A popular numerical approach for this is the person-by-person optimization (PBPO) that optimizes only one variable at a time with other variables being fixed, e.g., [25], [26]. It is also applicable for our setting. Before stating an algorithm, note the coordinate-wise convexity of R_0.

Lemma 1. *R_0 is strictly convex in $p_{\widehat{H}_j|H}(1|0), j \in \{0, 1, \ldots, N\}$ when other quantities are fixed.*

Proof: Focusing on agent $j \neq 0$ and rearranging (8) in terms of $p_{\widehat{H}_j|H}(1|0)$ and $p_{\widehat{H}_j|H}(1|1)$,

$$R_0 = c_{\mathsf{FA}} p_0 p_{\widehat{H}_j|H}(1|0) \left(A_1^{(j)} - A_0^{(j)} \right) + c_{\mathsf{FA}} p_0 A_0^{(j)} - c_{\mathsf{MD}} \bar{p}_0 p_{\widehat{H}_j|H}(1|1) \left(B_0^{(j)} - B_1^{(j)} \right) + c_{\mathsf{MD}} \bar{p}_0 B_0^{(j)}.$$

Now recall what $A_h^{(j)}, B_h^{(j)}$ stand for—$A_h^{(j)}$ (or $B_h^{(j)}$) is the false alarm (or missed detection) probability of the central agent conditioned on $\widehat{h}_j = h$. Also recall that conditioning on $\widehat{h}_j = 0$ increases the central agent's initial belief, which in turn implies the decision threshold also does, whereas conditioning on $\widehat{h}_j = 1$ decreases the decision threshold. Since the false alarm probability is decreasing in the decision threshold, we can conclude that $A_1^{(j)} - A_0^{(j)}$ is nonnegative always. A similar argument for missed detection shows $B_0^{(j)} - B_1^{(j)}$ is nonnegative. Finally the fact from the property of a receiver operating curve [27] that $p_{\widehat{H}_j|H}(1|1)$ is strictly concave in $p_{\widehat{H}_j|H}(1|0)$ yields the convexity in $p_{\widehat{H}_j|H}(1|0)$.

For $p_{\widehat{H}_0|H}(1|0)$, it can be shown similarly. ■

Therefore, a convex optimization algorithm with respect to $\{p_{\widehat{H}_j|H}(1|0)\}_j$ numerically finds the PBPO solution $\{p_{\widehat{H}_j|H}(1|0)\}_j$, which in turn implies the PBPO solution $(q_0, q_1, \ldots, q_N)$ since they are continuous bijection. Coordinated gradient descent with Gauss-Seidel update in the following solves the PBPO:

1) Initialize q_i for $i = 0, 1, \ldots, N$ arbitrarily.
2) For each $i \in \{1, \ldots, N\}$, update $p_{\widehat{H}_i|H}(1|0)$ to $p_{\widehat{H}_i|H}(1|0)'$ (so update q_i to q_i' as well) assuming $\{p_{\widehat{H}_j|H}(1|0)'\}_{j=1}^{i-1}$, $\{p_{\widehat{H}_j|H}(1|0)\}_{j=i+1}^{N}$, q_0 are all fixed. Then, update q_0 to q_0' assuming $\{p_{\widehat{H}_j|H}(1|0)'\}_{j \le N}$ are fixed.
3) Repeat 2) until $(q_0, q_1, \ldots, q_N)$ converges.

The algorithm exhibits monotone decreasing R_0 over each iteration in Step 2), hence, converges. Once the convergence occurs, there is no decrease in R_0 along any ith direction and, therefore, attains either a local minimum or a saddle point [28]. Repeating Step 1)–3) a number of times and selecting the solution that yields the least risk yields the global solution.

Since distributed agents' observations are i.i.d., the assumption of identical distributed beliefs is often made. Although this does not in general guarantee global optimality [29], it greatly simplifies numerical computation. Note that the fixed point in Fig. 4, i.e., $p_0 = q_0^* = q_1^* = \ldots = q_N^*$, is at $\frac{c_{\mathsf{MD}}}{c_{\mathsf{FA}} + c_{\mathsf{MD}}}$. Restricting to identical distributed beliefs, we can prove that this is globally optimal. Before stating it, note a useful property of (9). Due to page limitation, proof is provided in [24].

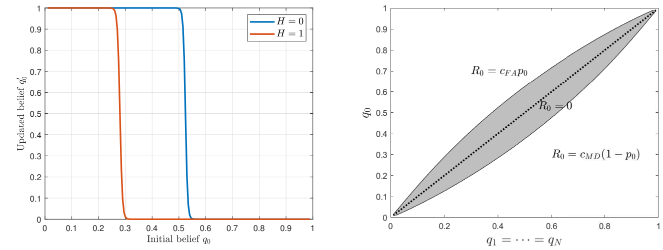

(a) (b)

Fig. 5. (a) Belief polarization (10) for $N = 100$ with $q_1 = \cdots = q_{100} = 0.4$. (b) Beliefs partition by limiting value of $R_0 \in \{0, c_{\mathsf{FA}} p_0, c_{\mathsf{MD}} \bar{p}_0\}$. The optimal points for large N suggested by Fig. 4, i.e., $\left(\frac{c_{\mathsf{MD}}}{c_{\mathsf{FA}} + c_{\mathsf{MD}}}, \ldots, \frac{c_{\mathsf{MD}}}{c_{\mathsf{FA}} + c_{\mathsf{MD}}}, \frac{c_{\mathsf{MD}}}{c_{\mathsf{FA}} + c_{\mathsf{MD}}} \right)$, are drawn in dotted line.

Lemma 2. *The right side of* (9) *is strictly decreasing in q_i, $i \in \{1, \ldots, N\}$ with other parameters being fixed.*

Then, the fixed point theorem follows.

Theorem 2. *Being aware of the true prior attains the globally minimal R_0 when $p_0 \in \{0, \frac{c_{\mathsf{MD}}}{c_{\mathsf{FA}} + c_{\mathsf{MD}}}, 1\}$, i.e., $p_0 = q_0^* = q_1^* = \cdots = q_N^*$ when $p_0 \in \{0, \frac{c_{\mathsf{MD}}}{c_{\mathsf{FA}} + c_{\mathsf{MD}}}, 1\}$.*

Proof: The cases $p_0 \in \{0, 1\}$ are trivial so focus on $p_0 = \frac{c_{\mathsf{MD}}}{c_{\mathsf{FA}} + c_{\mathsf{MD}}}$. At this q_i^*, each agent takes initial decision threshold $\lambda_i = 1/2$ by (3). It implies by symmetry that

$$p_{\widehat{H}_i|H}(1|0) = p_{\widehat{H}_i|H}(0|1) \text{ and } p_{\widehat{H}_i|H}(0|0) = p_{\widehat{H}_i|H}(1|1).$$

Furthermore, the central agent's initial threshold is also $1/2$ so that

$$p_{\widehat{H}_0|H, \widehat{H}_{-j}^N, \widehat{H}_j}(1|0, \widehat{h}_{-j}^N, h) = p_{\widehat{H}_0|H, \widehat{H}_{-j}^N, \widehat{H}_j}(0|1, (\widehat{h}_{-j}^N)', h'),$$

where $(\cdot)'$ stands for a flip of decision. Hence, $A_1^{(j)} = B_0^{(j)}, A_0^{(j)} = B_1^{(j)}$, and (9) hold. Since the right side of (9) is decreasing along the $q_1 = \cdots = q_N$ direction, the solution is unique. ■

IV. Many Homogeneous Distributed Agents

In this section, we consider asymptotic global optimality when $N \to \infty$. To this end, we begin with homogeneous distributed agents assumption, i.e., $q_1 = \cdots = q_N$, but q_0 being arbitrary, and will end up with $q_0 = q_1 = \cdots = q_N$. Recall the results for finite N that optimal beliefs q_0^*, q_i^* are dependent on p_0 so the system designer must be context-aware to attain the least Bayes risk. However, unlike finite N, simply setting $q_0 = q_1 = \cdots = q_N = \frac{c_{\mathsf{MD}}}{c_{\mathsf{FA}} + c_{\mathsf{MD}}}$ without knowledge of p_0 is asymptotically optimal, among all (possibly nonidentical) beliefs as in Thm. 5.

Theorem 3. *When $q_1 = \cdots = q_N$, the updated belief* (6) *of the central agent approaches either 0 or 1 almost surely as $N \to \infty$.*

TABLE I
ASYMPTOTIC RISK OF THE CENTRAL AGENT AS A FUNCTION OF INITIAL BELIEFS.

	$z_1 z_2^{Q(\lambda_1)}$	$z_1 z_2^{Q(\lambda_1-1)}$	Resulting R_0
CASE 1	>1	<1	0
CASE 2	<1	<1	$c_{\mathsf{FA}}p_0$
CASE 3	>1	>1	$c_{\mathsf{MD}}\bar{p}_0$
CASE 4	<1	>1	impossible

Proof. Consider the belief update formula (6) for $(\widehat{h}_1, \ldots, \widehat{h}_N)$ and define a random variable r_1 to be the ratio of 1 decisions in $\widehat{h}^N$, i.e., $r_1 \triangleq \frac{\#\text{ of ones in } \widehat{h}^N}{N}$. Then, by algebra, we can write

$$\frac{q_0'}{1-q_0'} = \frac{q_0}{1-q_0} \prod_{i=1}^{N} \frac{p_{\widehat{H}_i|H}(\widehat{h}_i|0)_{[0]}}{p_{\widehat{H}_i|H}(\widehat{h}_i|1)_{[0]}} = \frac{q_0}{1-q_0} \left(z_1 z_2^{r_1}\right)^N,$$

where

$$z_1 \triangleq \frac{p_{\widehat{H}_i|H}(0|0)_{[0]}}{p_{\widehat{H}_i|H}(0|1)_{[0]}}, z_2 \triangleq \frac{p_{\widehat{H}_i|H}(0|1)_{[0]}}{p_{\widehat{H}_i|H}(0|0)_{[0]}} \cdot \frac{p_{\widehat{H}_i|H}(1|0)_{[0]}}{p_{\widehat{H}_i|H}(1|1)_{[0]}}.$$

Here z_1, z_2 are dependent only on q_0 since they are perceived quantities by the central agent. In addition, $\widehat{h}_1, \ldots, \widehat{h}_N$ are N i.i.d. copies of Bernoulli random variable with $\mathbb{P}[\widehat{h}_i = 1|H = 0] = Q(\lambda_1)$ if $H = 0$, and $\mathbb{P}[\widehat{h}_i = 1|H = 1] = Q(\lambda_1 - 1)$ if $H = 1$. This implies that

$$r_1 \to \begin{cases} Q(\lambda_1) & \text{if } H = 0, \\ Q(\lambda_1 - 1) & \text{if } H = 1, \end{cases}$$

almost surely as N grows. Hence, the right side converges to

$$\frac{q_0'}{1-q_0'} = \begin{cases} \frac{q_0}{1-q_0}\left(z_1 z_2^{Q(\lambda_1)}\right)^N & \text{if } H = 0 \\ \frac{q_0}{1-q_0}\left(z_1 z_2^{Q(\lambda_1-1)}\right)^N & \text{if } H = 1 \end{cases} \quad (10)$$

almost surely. Depending on the value to be exponentiated, the right side approaches either 0 or ∞. Therefore we can conclude that the updated belief is polarized using the fact that $x/(1-x) : (0,1) \mapsto (0,\infty)$ is monotonic in $x \in (0,1)$. □

Thm. 3 reveals an interesting fact that when N is large, the central agent makes a decision either 0 or 1 almost surely, in other words, the decision is asymptotically deterministic, as a function of q_1 and q_0 no matter what value the private signal takes. Updating the belief, the central agent could make a correct decision always if $q_0' = 1$ when $h = 0$ and $q_0' = 0$ when $h = 1$. Tab. I summarizes, and corresponding regions are depicted in Fig. 5(b) with limiting values of R_0. The shaded region in Fig. 5(b) achieves $R_0 = 0$ asymptotically for all p_0. Clearly the shaded region contains $\frac{c_{\mathsf{MD}}}{c_{\mathsf{FA}}+c_{\mathsf{MD}}} = q_0 = q_1 = \cdots = q_N$ for any $c_{\mathsf{FA}}, c_{\mathsf{MD}}$, at which R_0 is asymptotically minimized regardless of p_0 as suggested numerically by Fig. 4.

Finally, we can also derive the speed of risk convergence to its limiting value in Fig. 5(b). To explicitly denote dependency on N, let $R_0^{(N)}$ be the risk of the central agent with N distributed agents and $R_0^{(\infty)} \triangleq \lim_{N\to\infty} R_0^{(N)} \in \{0, c_{\mathsf{FA}}p_0, c_{\mathsf{MD}}\bar{p}_0\}$. Then, the next theorem shows that $R_0^{(N)} \to R_0^{(\infty)}$ exponentially fast in N, that is,

$$\beta \triangleq -\lim_{N\to\infty} \frac{1}{N} \log\left(R_0^{(N)} - R_0^{(\infty)}\right)$$

is strictly positive.

Theorem 4. *Suppose (q_0, q_1) satisfies* CASE 1, 2, *or* 3, *that is, (q_0, q_1) strictly belongs to one of the regions in Fig. 5(b). Then, β is strictly positive and finite.*

Proof: Proof is mainly based on the concentration inequality of i.i.d. Bernoulli random variables and Chernoff approximation of the Q function. Details are in [24]. ■

Returning to the result of finite N in Fig. 4, two important observations can be made. The first is that we cannot achieve the smallest Bayes risk if p_0 is unknown since the optimal beliefs are functions of p_0. The other is that the optimal beliefs converges to $\frac{c_{\mathsf{MD}}}{c_{\mathsf{FA}}+c_{\mathsf{MD}}}$ as N grows, although curves for $N = 5, 10$ in Fig. 4 are drawn under $q_1 = \cdots = q_N$ assumption.

Denote the optimal risk exponent over all possible decision rules, not necessarily homogeneous nor LRTs, by

$$\beta^* \triangleq \sup\left(-\lim_{N\to\infty} \frac{1}{N} \log R_0^{(N)}\right),$$

where the supremum is over all decision rules. Unlike observations for finite N, the next theorem states that β^* can be attained simply by setting $q_0 = q_1 = \cdots = q_N = \frac{c_{\mathsf{MD}}}{c_{\mathsf{FA}}+c_{\mathsf{MD}}}$ for any p_0.

Theorem 5. *Suppose $q_0 = q_1 = \cdots = q_N = \frac{c_{\mathsf{MD}}}{c_{\mathsf{FA}}+c_{\mathsf{MD}}}$. Then the LRT asymptotically achieves β^* as $N \to \infty$. Furthermore,*

$$\beta^* = C(\mathit{Bern}(Q(0.5)), \mathit{Bern}(Q(-0.5))) \approx 0.0793,$$

where $C(\cdot,\cdot)$ is the Chernoff information.

Proof: The result of [29] is the mainstay of this proof. We show the claims by careful comparison with our model. Details are in [24]. ■

The scaling trick used in the proof of Thm. 5 also yields the following corollary.

Corollary 1. *The optimal risk exponent, β^*, is independent of values of $p_0, c_{\mathsf{FA}}, c_{\mathsf{MD}}$ and the presence of Y_0.*

V. CONCLUSION

This work investigates a social learning problem in a parallel network and mainly focuses on optimal beliefs. As in a tandem network [22], the optimal belief tuple that minimizes the central agent's Bayes risk is in general different from the tuple of true priors. The setting of many homogeneous distributed agents is also investigated. As suggested from Fig. 4 for finite number of agents, the optimal beliefs are asymptotically $\frac{c_{\mathsf{MD}}}{c_{\mathsf{FA}}+c_{\mathsf{MD}}}$ as $N \to \infty$ no matter what the true prior is. Also from the fact that the central agent's decision polarizes, belief partition depending on limit of Bayes risk is depicted in Fig. 5(b). It is also shown that the risk converges to its limiting value exponentially fast.

REFERENCES

[1] I. Lobel, E. Sadler, and L. R. Varshney, "Customer referral incentives and social media," *Manage. Sci.*, vol. 63, no. 10, pp. 3514–3529, Sep. 2016.

[2] S. Bikhchandani, D. Hirshleifer, and I. Welch, "A theory of fads, fashion, custom, and cultural change as informational cascades," *J. Polit. Econ.*, vol. 100, no. 5, pp. 992–1026, Oct. 1992.

[3] A. V. Banerjee, "A simple model of herd behavior," *Quart. J. Econ.*, vol. 107, no. 3, pp. 797–817, Aug. 1992.

[4] V. Bala and S. Goyal, "Conformism and diversity under social learning," *Econ. Theor.*, vol. 17, no. 1, pp. 101–120, Jan. 2001.

[5] L. Smith and P. Sørensen, "Pathological outcomes of observational learning," *Econometrica*, vol. 68, no. 2, pp. 371–398, Mar. 2000.

[6] D. Gale and S. Kariv, "Bayesian learning in social networks," *Games Econ. Behav.*, vol. 45, no. 2, pp. 329–346, Nov. 2003.

[7] T. N. Le, V. G. Subramanian, and R. A. Berry, "Information cascades with noise," *IEEE Trans. Signal Inf. Process. Netw.*, vol. 3, no. 2, pp. 239–251, Jun. 2017.

[8] R. Viswanathan and P. K. Varshney, "Distributed detection with multiple sensors: Part I—fundamentals," *Proc. IEEE*, vol. 85, no. 1, pp. 54–63, Jan. 1997.

[9] P. K. Varshney, *Distributed Detection and Data Fusion*. New York: Springer-Verlag, 1997.

[10] T. Berger, Z. Zhang, and H. Viswanathan, "The CEO problem," *IEEE Trans. Inf. Theory*, vol. 42, no. 3, pp. 887–902, May 1996.

[11] V. Saligrama, M. Alanyali, and O. Savas, "Distributed detection in sensor networks with packet losses and finite capacity links," *IEEE Trans. Signal Process.*, vol. 54, no. 11, pp. 4118–4132, Nov. 2006.

[12] Z. Zhang, E. K. P. Chong, A. Pezeshki, and W. Moran, "Hypothesis testing in feedforward networks with broadcast failures," *IEEE J. Sel. Topics Signal Process.*, vol. 7, no. 5, pp. 797–810, Oct. 2013.

[13] T. M. Cover, "Hypothesis testing with finite statistics," *Ann. Math. Stat.*, vol. 40, no. 3, pp. 828–835, 1969.

[14] M. E. Hellman and T. M. Cover, "Learning with finite memory," *Ann. Math. Stat.*, vol. 41, no. 3, pp. 765–782, 1970.

[15] Z.-B. Tang, K. R. Pattipati, and D. L. Kleinman, "Optimization of detection networks: Part I–tandem structures," *IEEE Trans. Syst., Man, Cybern.*, vol. 21, no. 5, pp. 1044–1059, Sept.-Oct. 1991.

[16] W. P. Tay, J. N. Tsitsiklis, and M. Z. Win, "On the subexponential decay of detection error probabilities in long tandems," *IEEE Trans. Inf. Theory*, vol. 54, no. 10, pp. 4767–4771, Oct. 2008.

[17] M. Alanyali, S. Venkatesh, O. Savas, and S. Aeron, "Distributed Bayesian hypothesis testing in sensor networks," in *Proc. Am. Contr. Conf. (ACC 2004)*, vol. 6, June-July 2004, pp. 5369–5374.

[18] K. R. Rad and A. Tahbaz-Salehi, "Distributed parameter estimation in networks," in *Proc. 49th IEEE Conf. Decision Control*, Dec. 2010, pp. 5050–5055.

[19] D. Acemoglu, M. A. Dahleh, I. Lobel, and A. Ozdaglar, "Bayesian learning in social networks," *Rev. Econ. Stud.*, vol. 78, no. 4, pp. 1201–1236, Oct. 2011.

[20] A. K. Sahu and S. Kar, "Distributed sequential detection for Gaussian shift-in-mean hypothesis testing," *IEEE Trans. Signal Process.*, vol. 64, no. 1, pp. 89–103, Jan. 2016.

[21] 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.

[22] D. Seo, R. K. Raman, J. B. Rhim, V. K. Goyal, and L. R. Varshney, "Beliefs in decision-making cascades," *IEEE Trans. Signal Process.*, vol. 67, no. 19, pp. 5103–5117, Oct. 2019.

[23] R. Radner, "Team decision problems," *Ann. Math. Stat.*, vol. 33, no. 3, pp. 857–881, Sep. 1962.

[24] D. Seo, R. K. Raman, and L. R. Varshney, "Social learning with beliefs in a parallel network," arXiv:1912.12284 [cs.IT], Dec. 2019.

[25] I. Y. Hoballah and P. K. Varshney, "Distributed Bayesian signal detection," *IEEE Trans. Inf. Theory*, vol. 35, no. 5, pp. 995–1000, Sep. 1989.

[26] Z.-B. Tang, K. R. Pattipati, and D. L. Kleinman, "An algorithm for determining the decision thresholds in a distributed detection problem," *IEEE Trans. Syst., Man, Cybern.*, vol. 21, no. 1, pp. 231–237, Jan.-Feb. 1991.

[27] H. V. Poor, *An Introduction to Signal Detection and Estimation*. Springer Science & Business Media, 1988.

[28] D. P. Bertsekas, *Constrained Optimization and Lagrange Multiplier Methods*. New York, USA: Academic Press, 1982.

[29] J. N. Tsitsiklis, "Decentralized detection by a large number of sensors," *Math. Control Signals, Syst.*, vol. 1, no. 2, pp. 167–182, Jun. 1988.