

Beliefs in Decision-Making Cascades

Daewon Seo , Ravi Kiran Raman , Joong Bum Rhim, Vivek K Goyal , and Lav R. Varshney 

Abstract—This work explores a social learning problem with agents having nonidentical noise variances and mismatched beliefs. We consider an N -agent binary hypothesis test in which each agent sequentially makes a decision based not only on a private observation, but also on preceding agents' decisions. In addition, the agents have their own beliefs instead of the true prior, and have nonidentical noise variances in the private signal. We focus on the Bayes risk of the last agent, where preceding agents are selfish. We first derive the optimal decision rule by recursive belief update and conclude, counterintuitively, that beliefs deviating from the true prior could be optimal in this setting. The effect of nonidentical noise levels in the two-agent case is also considered and analytical properties of the optimal belief curves are given. Next, we consider a predecessor selection problem wherein the subsequent agent of a certain belief chooses a predecessor from a set of candidates with varying beliefs. We characterize the decision region for choosing such a predecessor and argue that a subsequent agent with beliefs varying from the true prior often ends up selecting a suboptimal predecessor, indicating the need for a social planner. Lastly, we discuss an augmented intelligence design problem that uses a model of human behavior from cumulative prospect theory and investigate its near-optimality and suboptimality.

Index Terms—Social learning, cascading binary hypothesis test, cumulative prospect theory, augmented intelligence.

I. INTRODUCTION

TEAM decision-making typically involves individual decisions that are influenced by private observations and the opinions of the rest of the team. The *social learning* setting is one such context where decisions of individual agents are influenced by preceding agents in the team [3], [4]. We consider the setting in which individual agents are selfish and aim to minimize their perceived Bayes risk, according to their beliefs as reinforced by the decisions of preceding agents.

Social learning, also referred to as observational learning, has been widely studied and we provide a non-exhaustive listing

Manuscript received November 19, 2018; revised May 1, 2019 and July 2, 2019; accepted August 4, 2019. Date of publication August 28, 2019; date of current version September 5, 2019. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Soummya Kar. This work was supported by the National Science Foundation under Grants CCF-1101147 and CCF-1717530. This paper was presented in part at the IEEE International Symposium on Information Theory, Istanbul, Turkey, July 2013 [1] and in part at the IEEE International Conference on Acoustics, Speech, and Signal Processing, Calgary, AB, Canada, April 2018 [2]. (Corresponding author: Daewon Seo.)

D. Seo, R. K. Raman, and L. R. Varshney are with the Coordinated Science Laboratory and Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Urbana, IL 61801 USA (e-mail: dseo9@illinois.edu; rraman10@illinois.edu; varshney@illinois.edu).

J. B. Rhim is with the GroupM, New York, NY 10007 USA (e-mail: jbrhim@alum.mit.edu).

V. K. Goyal is with the Department of Electrical and Computer Engineering, Boston University, Boston, MA 02215 USA (e-mail: v.goyal@ieee.org).

Digital Object Identifier 10.1109/TSP.2019.2935865

of some of the relevant works. Aspects of conformism and “herding” were studied in [5]–[7], where an incorrect decision may cascade for the rest of the agents once agents at the beginning make incorrect decisions. The concept of herding is a consequence of boundedly informative private signals [8]. For example, assume the private signals are binary and give true or false information, each with positive probability. It can happen that a couple of the first agents receive false private signals and thus choose wrong actions. Then, the effect of these actions on the beliefs of subsequent agents can be so great as to cause them to ignore their private signals and follow their predecessors. The private signals are bounded so that they are not strong enough to overcome the effect of the wrong actions. Further convergence properties of actions taken under social learning have been explored under imperfect information [9]. The notion of sequential social learning has been generalized to learning from neighbors in networks [10], and explored in generality [11]. Social learning has also been explored under quantization of priors [12], and distributed detection with symmetric fusion [13].

Such social learning problems have also been studied as *distributed inference* or *learning*. The traditional setting assumes a central fusion node that aggregates all information from distributed nodes and makes the final decision [14], [15], where the links between distributed nodes and fusion center could be rate-limited [16] or imperfect [17]–[19]. It is also common to consider such learning over networks. The network setting could be the simplest tandem network (in particular, this is called serial detection) [20]–[23] and extended to a general network, in which all nodes can identify the hypothesis by repeatedly updating local beliefs without complete knowledge of network connectivity [24]–[26]. Independent works [27] and [28] propose similar update rules and convergence results for fixed networks and time-varying networks, respectively. In [29], binary hypothesis testing in the presence of Gaussian process noise is studied and minimal expected stopping times are derived. In [30], the setup where the entire hypotheses are locally indistinguishable, but globally identifiable by belief update is considered and convergence rate is provided.

This paper differs from the literature in the sense that we consider unbounded private signals so that there is no herding behavior. Unlike typical decision-making cascades (e.g., [31]) where all agents know the true prior, we assume agents may have *beliefs* that do not necessarily match the true prior. Further, private signal strengths of agents could be different, i.e., noise variances are not necessarily identical. Information is only propagated along the chain once so there is no iterative belief update. We focus largely on the effects of initial belief and private

signal strength. The decision-making of individual agents is also different in that agents make locally Bayes-optimal decisions, i.e., decisions that minimize their individual Bayes risk. This is different from the context of collectively optimizing the team's risk [13] or decision-making that maximizes a personal reward based on the last agent's decision [32].

We study cascading binary hypothesis testing (or sequential social learning following the notion of [10]) and characterize optimal beliefs of agents that minimize the Bayes risk of the last-acting agent. In general, it counterintuitively turns out that agents using beliefs that do not match the true prior are optimal, i.e., each agent has a perceived *belief* of the prior. For instance, in the two-agent system with equal noise levels, the optimal predecessor is one who overweights the belief for small prior, and underweights when it is large. On the other hand, the corresponding optimal last agent is one who behaves in the opposite way to the predecessor. For concise description, we refer to these two modes of operation as *open-minded* and *closed-minded*, respectively.¹ We describe analytical characteristics of the optimal beliefs and also show how the nature of such behaviors of agents change when noise levels differ in the private signals.

We are ultimately interested in the Bayes risk of the last-acting agent, and thus it is important that the last agent uses the correct set of preceding agents for the task. To this end, we consider a team construction problem for such cascading hypothesis testing, and characterize the criterion used for predecessor selection. We observe that self-organized teams may have suboptimal compositions, emphasizing the importance of a social planner that is aware of the true prior.

We also consider a collaborative decision-making system with human and AI (Artificial Intelligence or Augmented Intelligence). A cascading decision-making model captures the nature of collaboration in human-AI teams with either the AI system advising the human who makes the final decision or less typically a human advising an AI system that makes the final decision [34, p. 56]. Examples of the first kind include AI-assisted physicians, and of the second kind, human-in-the-loop AI systems such as crowdsourcing.

Note that in human-AI systems, human actions are affected by individual perceptions of the underlying context that cannot be tuned/controlled like machine agents (e.g. sensor nodes). Cumulative prospect theory [35]–[37] seeks to describe boundedly rational human behavior under risk by introducing probability reweighting functions. Among reweighting functions, the Prelec reweighting function [38] has significant empirical support and satisfies a majority of the axioms of prospect theory. We first show the Prelec model does not capture all patterns of optimal beliefs of agents in the case of diverse noise levels. Next we show that a team of suboptimal human-AI agents could outperform a team of standalone optimal human-AI agents, if it is well-composed.

The rest of this paper is organized as follows. Sec. II describes the cascading binary hypothesis testing problem. Sec. III proposes a recursive belief update equation that

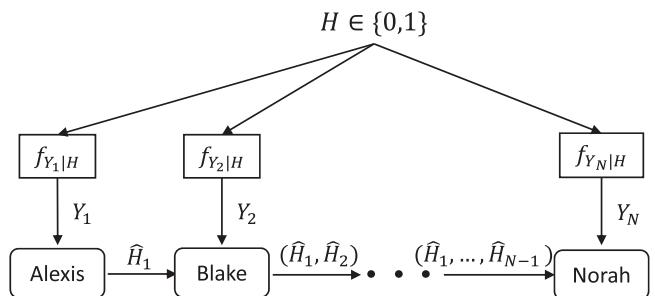


Fig. 1. A cascading decision making model with N agents.

transforms the cascading hypothesis testing problem into a single-agent binary hypothesis testing problem. Sec. IV shows the optimal beliefs that minimize the last agent's Bayes risk and Sec. V evaluates them for Gaussian likelihoods. Sec. VI considers a two-agent team construction problem and Sec. VII discusses design principles of AI-human collaboration systems. Sec. VIII concludes.

This cascading decision-making problem with identically noisy agents was first presented in [1] and in particular two-agent systems with varying noise levels were investigated in [2]. This paper integrates and generalizes our previous results, and also, significantly improves analytic understanding on optimal beliefs in cascading decision-making. In addition, we provide a novel interpretation of Prelec-like beliefs in terms of AI-human collaboration systems.

II. PROBLEM DESCRIPTION

Consider an N -agent cascading decision making problem, as illustrated in Fig. 1. The underlying hypothesis, $H \in \{0, 1\}$, is a binary signal with prior $\mathbb{P} H = 0 = p_0$ and $\mathbb{P} H = 1 = 1 - p_0$. There are N agents that sequentially detect the state in a predetermined order. The n th agent has a *private* signal Y_n generated according to the likelihood $f_{Y_n|H}$, which is not necessarily identical for all n . Let the decision made by the n th agent be \hat{H}_n . In addition to the private signal, the n th agent also observes the decisions made by preceding agents, $\{\hat{H}_1, \dots, \hat{H}_{n-1}\}$, to make a decision \hat{H}_n .

However, the n th agent believes the prior probability of the null hypothesis is $q_n \in (0, 1)$ as against the true prior probability p_0 . We call this the *belief* of the agent in order to distinguish it from the prior. Agent n is also aware of her own likelihood $f_{Y_n|H}$ that defines her private signal. However, she also perceives the likelihoods and beliefs of the other agents to be the same as hers, i.e., she thinks $f_{Y_j|H} = f_{Y_n|H}$, $q_j = q_n$ for all $j \neq n$, even though they could be different and unknown to her. We assume that the likelihood ratio of each agent is an increasing function in y ,² i.e., for all agents

$$\mathcal{L}_n(y) := \frac{f_{Y_n|H}(y|1)}{f_{Y_n|H}(y|0)}$$

is an increasing function of y .

²This property is particularly useful in uniformly most powerful (UMP) tests.

¹As far as we know, these terms were first introduced in [33].

Several numerical examples are given for private signals defined with independent additive Gaussian noise. The desired monotonicity also holds for many non-additive models, such as exponential distribution with mean H^{-1} , $H \in \mathbb{R}_+$, binomial distribution with success probability $H \in [0, 1]$, and Poisson distribution with rate $H \in \mathbb{R}_+$ are members of such family, where H could take two values.

Our performance analysis focuses on the last agent (N th agent, Norah) and her decision \hat{H}_N . Upon observing her private signal Y_N and the $(N-1)$ preceding decisions, she determines her decision rule. The relative importance of correct decisions and errors can be abstracted as a cost function. For simplicity, we assume correct decisions have zero cost and use the shorthand notations $c_{10} = c(1, 0)$ as the cost for false alarm or Type I error (choosing $\hat{H} = 1$ when $H = 0$), and $c_{01} = c(0, 1)$ as the cost for missed detection or Type II error (choosing $\hat{H} = 0$ when $H = 1$). In addition, we assume that agents have the same costs; they are a team in the sense of Radner [33]. Then the Bayes risk is

$$R_N = c_{10}p_0p_{\hat{H}_N|H}(1|0) + c_{01}(1-p_0)p_{\hat{H}_N|H}(0|1). \quad (1)$$

As \hat{H}_n depends on the previous decisions, the computation of (1) also depends on $(\hat{H}_1, \dots, \hat{H}_{N-1})$, and the Bayes risk can be expanded as

$$\begin{aligned} R_N = & \sum_{\hat{h}_1, \dots, \hat{h}_{N-1}} c_{10}p_0p_{\hat{H}_N, \hat{H}_{N-1}, \dots, \hat{H}_1|H}(1, \hat{h}_{N-1}, \dots, \hat{h}_1|0) \\ & + c_{01}(1-p_0)p_{\hat{H}_N, \hat{H}_{N-1}, \dots, \hat{H}_1|H}(0, \hat{h}_{N-1}, \dots, \hat{h}_1|1). \end{aligned} \quad (2)$$

We determine the optimal set of beliefs of the agents $\{q_n^*\}_{n=1}^N$ that minimize (2).

In our model, the n th agent minimizes her *perceived* Bayes risk, which is the Bayes risk with prior probability p_0 replaced by her belief q_n . In other words, for all $n = 1, \dots, N$, the n th agent adopts the decision rule that minimizes her perceived Bayes risk R_n , and her decision is revealed to other agents as a public signal. The decisions $\{\hat{H}_1, \dots, \hat{H}_{n-1}\}$ of the earlier-acting agents reveal information about H and thus should be incorporated into the decision-making process by agent n . As mentioned earlier, since she believes q_n is the true prior, she aggregates information under the assumption that $q_1 = q_2 = \dots = q_n$.

It is important to note that every agent is selfish and rational; the agents do not adjust their decision rules for Norah's sake. The novelty in the model (and hence in the conclusions) comes from agent n having the limitation of using a private initial belief q_n in place of the true prior probability p_0 .

A. Prospect Theory

Let us also formally introduce the Prelec reweighting function from cumulative prospect-theoretic models of human behavior. It spans a family of open- and closed-minded beliefs (will be clarified later) and thus the optimal beliefs that emerge in following sections could be approximated by a function in the Prelec family.

Definition 1 ([38]): For $\alpha, \beta > 0$, the Prelec reweighting function $w : [0, 1] \mapsto [0, 1]$ is

$$w(p; \alpha, \beta) = \exp(-\beta(-\log p)^\alpha).$$

The function satisfies several properties such as:

- 1) $w(p; \alpha, \beta)$ is strictly increasing;
- 2) has a unique fixed point $w(p; \alpha, \beta) = p$ at $p^* = \exp(-\exp(\log \beta/(1 - \alpha)))$; and
- 3) spans a class of open-minded beliefs when $\alpha < 1$, i.e., overweights (underweights) small (high) probability, and vice versa when $\alpha > 1$.

A more generic form, termed composite Prelec weighting function, has been defined in [39].

B. Notations

Throughout the paper, we use f for continuous probability density functions and p for discrete probability mass functions. All logarithms are natural logarithms. We use $\mathcal{N}(\mu, \sigma^2)$ to denote a Gaussian distribution with mean μ and variance σ^2 , and $\phi(x; \mu, \sigma^2)$ to denote its density function, i.e.,

$$\phi(x; \mu, \sigma^2) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}.$$

Also in the case of the standard Gaussian, $\phi(x) := \phi(x; 0, 1)$ for simplicity. $Q(x)$ is defined as the complementary cumulative distribution function of the standard Gaussian,

$$Q(x) = \int_x^\infty \phi(t) dt.$$

III. BELIEF UPDATE AND SEQUENTIAL DECISION MAKING

Our model assumes unbounded private signals. Thus, unlike in [5], [6], it is always possible that a subsequent agent may not follow previous decisions; that is, herding happens with arbitrarily low probability. We now discuss using both a decision history and private signals for Bayesian binary hypothesis testing. The decision rule can be interpreted as each agent updating her posterior belief based on the decision history and then applying a likelihood ratio test to her private signal.

A. Alexis, the First Agent

Since Alexis has no prior decision history, she follows usual binary hypothesis testing. She uses the following likelihood ratio test with her initial belief q_1 , with ties broken arbitrarily:

$$\mathcal{L}_1(y_1) = \frac{f_{Y_1|H}(y_1|1)}{f_{Y_1|H}(y_1|0)} \begin{cases} \hat{H}_1=1 \\ \hat{H}_1=0 \end{cases} \gtrless \frac{c_{10}q_1}{c_{01}(1-q_1)}. \quad (3)$$

Since we assume the likelihood ratio is increasing in y_1 , the rule simplifies to comparing the private signal with an appropriate decision threshold:

$$y_1 \begin{cases} \hat{H}_1=1 \\ \hat{H}_1=0 \end{cases} \gtrless \lambda_1(q_1), \quad (4)$$

where $\lambda_i(q)$ denotes the decision threshold λ that satisfies

$$\mathcal{L}_i(\lambda) = \frac{f_{Y_i|H}(\lambda|1)}{f_{Y_i|H}(\lambda|0)} = \frac{c_{10}q}{c_{01}(1-q)}. \quad (5)$$

B. Blake, the Second Agent

Blake observes Alexis's decision $\widehat{H}_1 = \widehat{h}_1$ and evaluates the likelihood ratio for (\widehat{H}_1, Y_2) , using his initial belief q_2 as

$$\frac{f_{Y_2, \widehat{H}_1|H}(y_2, \widehat{h}_1|1)}{f_{Y_2, \widehat{H}_1|H}(y_2, \widehat{h}_1|0)} \underset{\widehat{H}_2=1}{\gtrless} \frac{c_{10}q_2}{c_{01}(1-q_2)}. \quad (6)$$

The private signals Y_1 and Y_2 are independent conditioned on H , so \widehat{H}_1 and Y_2 are also independent conditioned on H . Hence, the left side of (6) is

$$f_{Y_2, \widehat{H}_1|H}(y_2, \widehat{h}_1|h) = f_{Y_2|H}(y_2|h)p_{\widehat{H}_1|H}(\widehat{h}_1|h).$$

So we can rewrite (6) as³

$$\frac{f_{Y_2|H}(y_2|1)}{f_{Y_2|H}(y_2|0)} \underset{\widehat{H}_2=0}{\gtrless} \frac{c_{10}q_2}{c_{01}(1-q_2)} \frac{p_{\widehat{H}_1|H}(\widehat{h}_1|0)[2]}{p_{\widehat{H}_1|H}(\widehat{h}_1|1)[2]}. \quad (7)$$

The likelihood ratio test (7) can be interpreted as Blake updating his initial belief upon observing Alexis's decision \widehat{H}_1 . Combined with q_2 , his initial belief is updated according to $p_{\widehat{H}_1|H}(\widehat{h}_1|h)[2]$, from q_2 to $q_2^{\widehat{h}_1}$:

$$\frac{q_2^{\widehat{h}_1}}{1 - q_2^{\widehat{h}_1}} = \frac{q_2}{1 - q_2} \frac{p_{\widehat{H}_1|H}(\widehat{h}_1|0)[2]}{p_{\widehat{H}_1|H}(\widehat{h}_1|1)[2]}. \quad (8)$$

The posterior belief is

$$\begin{aligned} q_2^{\widehat{h}_1} &= \frac{q_2 p_{\widehat{H}_1|H}(\widehat{h}_1|0)[2]}{q_2 p_{\widehat{H}_1|H}(\widehat{h}_1|0)[2] + (1 - q_2) p_{\widehat{H}_1|H}(\widehat{h}_1|1)[2]} \\ &= \frac{p_{\widehat{H}_1|H}(\widehat{h}_1, 0)[2]}{p_{\widehat{H}_1|H}(\widehat{h}_1, 0)[2] + p_{\widehat{H}_1|H}(\widehat{h}_1, 1)[2]} \\ &= p_{H|\widehat{H}_1}(0|\widehat{h}_1)[2]. \end{aligned} \quad (9)$$

It should be noted that the true $p_{\widehat{H}_1|H}(\widehat{h}_1|h)$ is given by

$$\begin{aligned} p_{\widehat{H}_1|H}(0|h) &= p_{\widehat{H}_1|H}(0|h)[1] = \mathbb{P}[Y_1 \leq \lambda_1(q_1)|H = h] \\ &= \int_{-\infty}^{\lambda_1(q_1)} f_{Y_1|H}(y|h) dy, \\ p_{\widehat{H}_1|H}(1|h) &= \int_{\lambda_1(q_1)}^{\infty} f_{Y_1|H}(y|h) dy. \end{aligned}$$

³The subscript [2] in $p_{\widehat{H}_1|H}(\widehat{h}_1|h)[2]$ indicates the value of $p_{\widehat{H}_1|H}(\widehat{h}_1|h)$ that Blake (the second agent) thinks. We specify this because Blake does not know Alexis's belief q_1 . Thus, he interprets her decision based on his belief q_2 . The value is different from the true value of $p_{\widehat{H}_1|H}(\widehat{h}_1|h) = p_{\widehat{H}_1|H}(\widehat{h}_1|h)[1]$. Of course, it will also be different from what Chuck, the third agent, perceives, which is denoted by $p_{\widehat{H}_1|H}(\widehat{h}_1|h)[3]$. This will be explained in the next subsection.

But Blake evaluates Alexis's decision \widehat{H}_1 as if it were made based on q_2 and the likelihood $f_{Y_2|H}(\cdot)$, as against q_1 , $f_{Y_1|H}(\cdot)$ respectively. Thus the probability $p_{\widehat{H}_1|H}(\widehat{h}_1|h)$ is computed based on $\lambda_2(q_2)$, instead of $\lambda_1(q_2)$:

$$p_{\widehat{H}_1|H}(0|h)[2] = \int_{-\infty}^{\lambda_2(q_2)} f_{Y_2|H}(y|h) dy, \quad (10a)$$

$$p_{\widehat{H}_1|H}(1|h)[2] = \int_{\lambda_2(q_2)}^{\infty} f_{Y_2|H}(y|h) dy. \quad (10b)$$

An interesting observation is that Alexis's belief q_1 does not affect Blake's belief update as observed in (9) and (10). That is, for any belief q_1 that Alexis might hold, Blake, who does not know this belief, presumes that the conditional probabilities are computed according to (10) and updates his belief as in (9) which depends only on Blake's initial belief and Alexis's decision.

However, Alexis's initial belief implicitly affects Blake's performance since her biased belief changes the resulting decisions whose probabilities are embedded in the probability of Blake's decision:

$$\begin{aligned} p_{\widehat{H}_2|H}(\widehat{h}_2|h) &= \sum_{\widehat{h}_1 \in \{0, 1\}} p_{\widehat{H}_2, \widehat{H}_1|H}(\widehat{h}_2, \widehat{h}_1|h) \\ &= p_{\widehat{H}_2|\widehat{H}_1, H}(\widehat{h}_2|0, h)[2] \times p_{\widehat{H}_1|H}(0|h)[1] \\ &\quad + p_{\widehat{H}_2|\widehat{H}_1, H}(\widehat{h}_2|1, h)[2] \times p_{\widehat{H}_1|H}(1|h)[1]. \end{aligned}$$

Thus, Alexis's biased belief changes the probability of not only her decision but also of Blake's decision.

C. Chuck, the Third Agent

Chuck's detection process is similar to Blake's. He observes both Alexis's and Blake's decisions and also updates his initial belief q_3 like in (8):

$$\begin{aligned} \frac{q_3^{\widehat{h}_1, \widehat{h}_2}}{1 - q_3^{\widehat{h}_1, \widehat{h}_2}} &= \frac{q_3}{1 - q_3} \frac{p_{\widehat{H}_2, \widehat{H}_1|H}(\widehat{h}_2, \widehat{h}_1|0)[3]}{p_{\widehat{H}_2, \widehat{H}_1|H}(\widehat{h}_2, \widehat{h}_1|1)[3]} \\ &= \left(\frac{q_3}{1 - q_3} \frac{p_{\widehat{H}_1|H}(\widehat{h}_1|0)[3]}{p_{\widehat{H}_1|H}(\widehat{h}_1|1)[3]} \right) \frac{p_{\widehat{H}_2|\widehat{H}_1, H}(\widehat{h}_2|\widehat{h}_1, 0)[3]}{p_{\widehat{H}_2|\widehat{H}_1, H}(\widehat{h}_2|\widehat{h}_1, 1)[3]}. \end{aligned} \quad (11)$$

Note that \widehat{H}_1 and \widehat{H}_2 are not conditionally independent given H as Blake's decision \widehat{H}_2 depends on Alexis's decision \widehat{H}_1 .

Chuck's belief update can be understood as a two-step process. The first step is to update his belief according to Alexis's decision:

$$\frac{q_3^{\widehat{h}_1}}{1 - q_3^{\widehat{h}_1}} = \frac{q_3}{1 - q_3} \frac{p_{\widehat{H}_1|H}(\widehat{h}_1|0)[3]}{p_{\widehat{H}_1|H}(\widehat{h}_1|1)[3]}. \quad (12)$$

The second step is to update it from $q_3^{\widehat{h}_1}$ based on Blake's decision:

$$\frac{q_3^{\widehat{h}_1, \widehat{h}_2}}{1 - q_3^{\widehat{h}_1, \widehat{h}_2}} = \frac{q_3^{\widehat{h}_1}}{1 - q_3^{\widehat{h}_1}} \frac{p_{\widehat{H}_2|\widehat{H}_1, H}(\widehat{h}_2|\widehat{h}_1, 0)[3]}{p_{\widehat{H}_2|\widehat{H}_1, H}(\widehat{h}_2|\widehat{h}_1, 1)[3]}. \quad (13)$$

Again, Chuck is not aware of neither Alexis's nor Blake's initial beliefs or likelihoods. Thus, Chuck computes all probabilities based on his own belief q_3 and likelihood $f_{Y_3|H}$, which is indicated by the subscript [3] in (12) and (13).

Details of computations of (12) and (13) are as follows:

$$p_{\hat{H}_1|H}(0|h)_{[3]} = \int_{-\infty}^{\lambda_3(q_3)} f_{Y_3|H}(y|h) dy,$$

$$p_{\hat{H}_1|H}(1|h)_{[3]} = \int_{\lambda_3(q_3)}^{\infty} f_{Y_3|H}(y|h) dy.$$

Similar to Blake (8), Chuck computes $q_3^{\hat{h}_1}$ for $\hat{H}_1 = 0$ and $\hat{H}_1 = 1$ respectively as:

$$q_3^0 = \frac{q_3}{q_3 + (1 - q_3) \frac{\int_{-\infty}^{\lambda_3(q_3)} f_{Y_3|H}(y|1) dy}{\int_{-\infty}^{\lambda_3(q_3)} f_{Y_3|H}(y|0) dy}}, \quad (14a)$$

$$q_3^1 = \frac{q_3}{q_3 + (1 - q_3) \frac{\int_{\lambda_3(q_3)}^{\infty} f_{Y_3|H}(y|1) dy}{\int_{\lambda_3(q_3)}^{\infty} f_{Y_3|H}(y|0) dy}}. \quad (14b)$$

Then,

$$p_{\hat{H}_2|\hat{H}_1,H}(0|\hat{h}_1, h)_{[3]} = \int_{-\infty}^{\lambda_3(q_3^{\hat{h}_1})} f_{Y_3|H}(y|h) dy, \quad (15a)$$

$$p_{\hat{H}_2|\hat{H}_1,H}(1|\hat{h}_1, h)_{[3]} = \int_{\lambda_3(q_3^{\hat{h}_1})}^{\infty} f_{Y_3|H}(y|h) dy. \quad (15b)$$

Even though the value of \hat{h}_1 does not appear in (15), it is implicit in $q_3^{\hat{h}_1}$ and affects the computation results. Chuck's posterior belief $q_3^{\hat{h}_1, \hat{h}_2}$ is obtained by substituting (14) and (15) in (13).

D. Norah, the N th Agent

Norah, the N th agent, observes Y_N and $\{\hat{H}_1, \dots, \hat{H}_{N-1}\}$. Paralleling the arguments in the preceding subsections, her initial belief update is a function of q_N as well as $\{\hat{H}_1, \dots, \hat{H}_{N-1}\}$, but not of $\{q_1, \dots, q_{N-1}\}$. Generalizing (11), we have

$$\frac{q_N^{\hat{h}_1, \dots, \hat{h}_{N-1}}}{1 - q_N^{\hat{h}_1, \dots, \hat{h}_{N-1}}} = \frac{q_N}{1 - q_N} \frac{p_{\hat{H}_1|H}(\hat{h}_1|0)_{[N]}}{p_{\hat{H}_1|H}(\hat{h}_1|1)_{[N]}} \times \prod_{n=2}^{N-1} \frac{p_{\hat{H}_n|\hat{H}_{n-1}, \dots, \hat{H}_1, H}(\hat{h}_n|\hat{h}_{n-1}, \dots, \hat{h}_1, 0)_{[N]}}{p_{\hat{H}_n|\hat{H}_{n-1}, \dots, \hat{H}_1, H}(\hat{h}_n|\hat{h}_{n-1}, \dots, \hat{h}_1, 1)_{[N]}}. \quad (16)$$

Combining all observations, we obtain the following theorem. Define the initial belief update function for N th agent, U_N as

$$q_N^{\hat{h}_1, \dots, \hat{h}_{N-1}} = U_N(q_N, \hat{h}_1, \hat{h}_2, \dots, \hat{h}_{N-1}; N).$$

Theorem 1: The function $U_n, n \leq N$ yielding the posterior belief of N th agent has the following recurrence relation:

- For $n = 1$, $U_1(q; N) = q$.
- For $n > 1$,

$$U_n(q, \hat{h}_1, \dots, \hat{h}_{n-1}, 0; N) = \frac{\tilde{q}}{\tilde{q} + (1 - \tilde{q}) \frac{\int_{-\infty}^{\lambda_N(\tilde{q})} f_{Y_N|H}(y|1) dy}{\int_{-\infty}^{\lambda_N(\tilde{q})} f_{Y_N|H}(y|0) dy}}, \quad (17a)$$

$$U_n(q, N, \hat{h}_1, \dots, \hat{h}_{n-1}, 1; N) = \frac{\tilde{q}}{\tilde{q} + (1 - \tilde{q}) \frac{\int_{\lambda_N(\tilde{q})}^{\infty} f_{Y_N|H}(y|1) dy}{\int_{\lambda_N(\tilde{q})}^{\infty} f_{Y_N|H}(y|0) dy}}, \quad (17b)$$

where $\tilde{q} = U_{n-1}(q, \hat{h}_1, \dots, \hat{h}_{n-1}; N)$.

Note that capital N in (17a) and (17b) indicate the recursive updates are computed from the value that the N th agent thinks.

Fig. 2 depicts the function $U_4(q_4, \hat{h}_1, \hat{h}_2, \hat{h}_3; 4)$ for $N = 4$ for eight possible combinations of Alexis's, Blake's, and Chuck's decisions $(\hat{h}_1, \hat{h}_2, \hat{h}_3)$. An interesting property of U_N is that the posterior belief is much more dependent on the most recent decision \hat{h}_{N-1} than on the earlier decisions $(\hat{h}_1, \dots, \hat{h}_{N-2})$. This is because later acting agents consider more previous decisions, and hence more information than the first agents, their decisions should carry more weight. In this sense, we can say that recent decisions give more information than earlier decisions. This is especially the case when the $(N-1)$ th agent has not followed precedent. This is because the N th agent rationally concludes that the $(N-1)$ th agent observed strong evidence to justify a deviation from precedent. For example, if the decision history of the first five agents is $(0, 0, 0, 0, 1)$ then the sixth agent takes the last decision 1 seriously even though the first four agents chose 0. A reversal of an arbitrarily long precedent sequence may occur because we assume unbounded private signals; if private signals are bounded [5], [6], then the influence of the precedent can reach a point where agents cannot receive a signal strong enough to justify a decision running counter to precedent. Another interesting point is that smaller noise variance changes beliefs more. It is clear from (17), but also reasonable that when the variance is smaller, the N th agent trusts and is more inclined towards previous decisions. Note even though the prior updates of Norah in Fig. 2 do not depend on $\{q_1, \dots, q_{N-1}\}$ and their corresponding likelihoods, the probability of prior decisions depends on them and implicitly, so does Norah's decision.

As we can see in Fig. 2, the dominant previous decision for agent N is the decision of agent $(N-1)$. We can prove that observing the $(N-1)$ th agent's decision 0 (or decision 1), the N th agent's posterior belief becomes larger (or smaller), which in turn implies that the decision threshold of N th agent becomes larger (or smaller) so that she is more likely to declare decision 0 (or 1) as well.

Theorem 2: Suppose that noises are independent and additive, and have continuous densities. Fix some prior decisions $\{\hat{h}_1, \dots, \hat{h}_{N-2}\}$ and let $\tilde{q}_N, \tilde{q}_N^0, \tilde{q}_N^1$ denote the posterior beliefs of the N th agent given the $(N-2)$ decisions only, the $(N-2)$ decisions with $\hat{h}_{N-1} = 0$, and the $(N-2)$ decisions with $\hat{h}_{N-1} = 1$. Then,

$$\tilde{q}_N^1 < \tilde{q}_N < \tilde{q}_N^0.$$

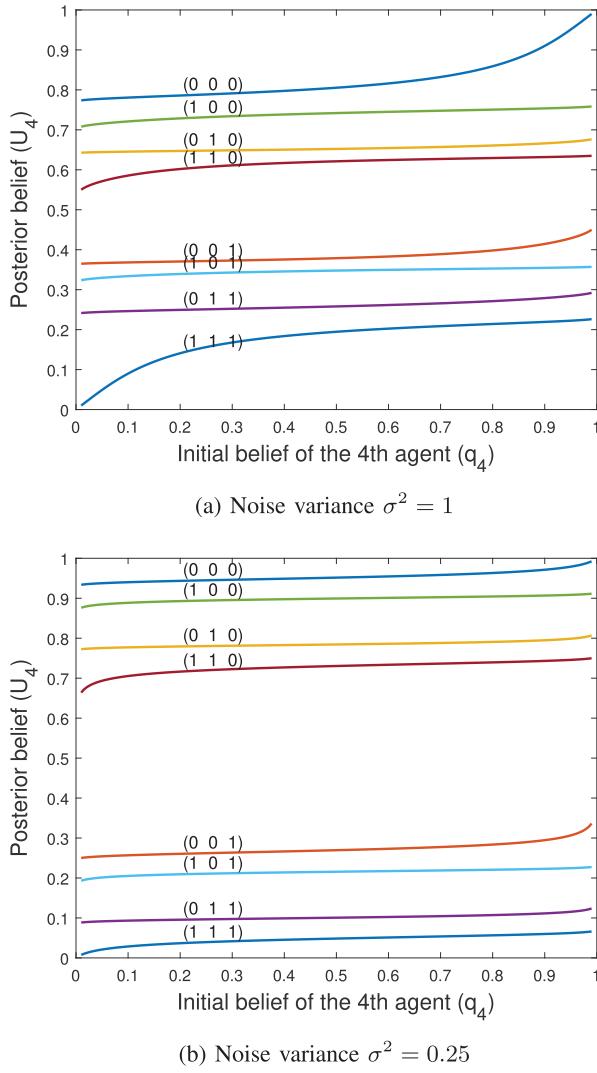


Fig. 2. The function $U_4(q_4, \hat{h}_1, \hat{h}_2, \hat{h}_3; 4)$ —posterior belief of the fourth agent ($q_4^{\hat{h}_1, \hat{h}_2, \hat{h}_3}$)—for each possible combination of Alexis's, Blake's, and Chuck's decisions $[\hat{h}_1, \hat{h}_2, \hat{h}_3]$ when $c_{10} = c_{01} = 1$ and private signals are distorted by additive Gaussian noise with two noise levels. The posterior belief is mostly dependent on Chuck's decision; the top four curves are for $\hat{h}_3 = 0$ and the bottom four curves are for $\hat{h}_3 = 1$.

Proof: We know that $\tilde{q}_N, \tilde{q}_N^0, \tilde{q}_N^1$ differ only by the last multiplicative term of (16). Since $\frac{q}{1-q}$ is monotone increasing, the statement is equivalent to showing:

$$\frac{\int_{\lambda_N(\tilde{q}_N)}^{\infty} f_{Y_N|H}(y|0)dy}{\int_{\lambda_N(\tilde{q}_N)}^{\infty} f_{Y_N|H}(y|1)dy} < 1 < \frac{\int_{-\infty}^{\lambda_N(\tilde{q}_N)} f_{Y_N|H}(y|0)dy}{\int_{-\infty}^{\lambda_N(\tilde{q}_N)} f_{Y_N|H}(y|1)dy}.$$

Since the noise is independent and additive, $f_{Y_N|H}(y|1) = f_{Y_N|H}(y-1|0)$ so the term on the left side

$$\begin{aligned} \frac{\int_{\lambda_N(\tilde{q}_N)}^{\infty} f_{Y_N|H}(y|0)dy}{\int_{\lambda_N(\tilde{q}_N)}^{\infty} f_{Y_N|H}(y|1)dy} &= \frac{\int_{\lambda_N(\tilde{q}_N)}^{\infty} f_{Y_N|H}(y|0)dy}{\int_{\lambda_N(\tilde{q}_N)-1}^{\infty} f_{Y_N|H}(y|0)dy} \\ &= \frac{\int_{\lambda_N(\tilde{q}_N)}^{\infty} f_{Y_N|H}(y|0)dy}{\int_{\lambda_N(\tilde{q}_N)-1}^{\infty} f_{Y_N|H}(y|0)dy + \int_{\lambda_N(\tilde{q}_N)}^{\infty} f_{Y_N|H}(y|0)dy} < 1. \end{aligned}$$

The right inequality can be shown similarly. \blacksquare

Considering the complicated relationships that individual decisions have on the evolution of initial beliefs, it is also important to verify if the belief evolution preserves the ordering, given the same set of subsequent decisions. That is, given two beliefs $q_L < q_R$ at some point of the recursive update and the same sequence of following d decisions, then it is important to characterize the likelihoods for which the ordering is preserved in the resulting posterior beliefs, given the sequence of decisions, which is described in the following theorem.

Theorem 3: Suppose that noise is independent and additive, and has a continuous density. Consider two beliefs $q_L < q_R$. Then, for any given later-acting decisions d , the posterior belief satisfies $q_L^d < q_R^d$ if and only if

$$g_1(q) := \frac{q}{1-q} \frac{\int_{-\infty}^{\lambda_N(q)} f_{Y_N|H}(y|0)dy}{\int_{-\infty}^{\lambda_N(q)} f_{Y_N|H}(y|1)dy}, \quad (18)$$

$$g_2(q) := \frac{q}{1-q} \frac{\int_{\lambda_N(q)}^{\infty} f_{Y_N|H}(y|0)dy}{\int_{\lambda_N(q)}^{\infty} f_{Y_N|H}(y|1)dy} \quad (19)$$

are both increasing in q .

Proof: Note that once observing decision 0, beliefs are updated as

$$\begin{aligned} \frac{q_L^0}{1-q_L^0} &= \frac{q_L}{1-q_L} \frac{\int_{-\infty}^{\lambda_N(q_L)} f_{Y_N|H}(y|0)dy}{\int_{-\infty}^{\lambda_N(q_L)} f_{Y_N|H}(y|1)dy}, \\ \frac{q_R^0}{1-q_R^0} &= \frac{q_R}{1-q_R} \frac{\int_{-\infty}^{\lambda_N(q_R)} f_{Y_N|H}(y|0)dy}{\int_{-\infty}^{\lambda_N(q_R)} f_{Y_N|H}(y|1)dy}, \end{aligned}$$

and so if (18) holds, $q_L^0 < q_R^0$. Similarly, (19) can be shown by updating after decision 1. \blacksquare

Let us state some properties of Mills ratio [40], [41], which is about Gaussian distribution, and we will see that $g_1(q), g_2(q)$ are both increasing if likelihood is Gaussian.

Lemma 1 [41]: Define $\eta(x) := \phi(x)/Q(x)$, the inverse of Mills ratio. Then, for any $x \in \mathbb{R}$, it is true that $0 < \eta'(x) < 1$ and $\eta''(x) > 0$.

Corollary 1: Consider a Gaussian likelihood, i.e., $Y_N = H + Z_N$, where Z_N are independent and identically drawn from $\mathcal{N}(0, \sigma^2)$, for some $\sigma^2 > 0$. Then $g_1(q), g_2(q)$ are both increasing in q .

Proof: Let us consider $g_2(q)$ first. For the binary hypothesis test with Gaussian noise, we know that the decision threshold for the likelihood ratio test is given by

$$\lambda_N(q) = \frac{1}{2} + \sigma^2 \log \left(\frac{c_{10}q}{c_{01}(1-q)} \right).$$

Then, we have

$$g_2(q) = \frac{q}{1-q} \frac{Q\left(\frac{\lambda_N(q)}{\sigma}\right)}{Q\left(\frac{\lambda_N(q)-1}{\sigma}\right)}.$$

Letting $x := \log \frac{c_{10}q}{c_{01}(1-q)}$, it is sufficient to show that

$$\begin{aligned}\tilde{g}(x) &:= \log \left(\frac{c_{10}}{c_{01}} g_2(q) \right) \\ &= x + \log \left(Q \left(\sigma x + \frac{1}{2\sigma} \right) \right) - \log \left(Q \left(\sigma x - \frac{1}{2\sigma} \right) \right),\end{aligned}$$

is increasing in x since c_{10}, c_{01} are positive constants, $\log(\cdot)$ is a monotonically increasing function, and x is a strictly increasing function of q .

The first derivative of \tilde{g} is given by

$$\tilde{g}'(x) = 1 - \sigma \eta \left(\sigma x + \frac{1}{2\sigma} \right) + \sigma \eta \left(\sigma x - \frac{1}{2\sigma} \right). \quad (20)$$

Since $\eta(\cdot)$ is a continuous function, using the mean value theorem, there exists $y \in (\sigma x - \frac{1}{2\sigma}, \sigma x + \frac{1}{2\sigma})$, such that

$$\sigma \eta \left(\sigma x + \frac{1}{2\sigma} \right) - \sigma \eta \left(\sigma x - \frac{1}{2\sigma} \right) = \sigma \eta'(y) \frac{1}{\sigma} = \eta'(y). \quad (21)$$

From the first property of Lem. 1, $0 < \eta'(y) < 1$, we have

$$\eta \left(\sigma x + \frac{1}{2\sigma} \right) - \eta \left(\sigma x - \frac{1}{2\sigma} \right) < 1.$$

Thus, from (20), it follows that $\tilde{g}'(x) > 0$ for all x , indicating that $\tilde{g}(\cdot)$ is an increasing function of x . This in turn implies that $g_2(\cdot)$ is also an increasing function.

To prove the result for g_1 , it is sufficient to observe that by the symmetry of error probabilities:

$$g_1(q) = \frac{1}{g_2(1-q)}. \quad \blacksquare$$

IV. OPTIMAL BELIEF

We described the initial belief evolution and decision-making model in Sec. III. In this section, we investigate the set of initial beliefs that minimize the Bayes risk. We consider the case of two agents for analytical tractability although the broad nature of the arguments extend to multi-agent systems. Note that the Bayes risk of the system with $N = 2$ is the same as Blake's Bayes risk because his decision is adopted as the final decision.

Let us recapitulate the computation of Blake's Bayes risk. Alexis chooses her decision threshold as $\lambda_1 := \lambda_1(q_1)$. Her probabilities of error are given by

$$\begin{aligned}P_{e,1}^I &= p_{\hat{H}_1|H}(1|0) = \int_{\lambda_1}^{\infty} f_{Y_1|H}(y|0) dy, \\ P_{e,1}^{II} &= p_{\hat{H}_1|H}(0|1) = \int_{-\infty}^{\lambda_1} f_{Y_1|H}(y|1) dy.\end{aligned}$$

Blake however presumes Alexis uses the decision threshold $\lambda_{1,[2]} := \lambda_2(q_2)$ and computes her probabilities of error accordingly⁴:

$$\begin{aligned}P_{e,1,[2]}^I &= p_{\hat{H}_1|H}(1|0)_{[2]} = \int_{\lambda_{1,[2]}}^{\infty} f_{Y_2|H}(y|0) dy, \\ P_{e,1,[2]}^{II} &= p_{\hat{H}_1|H}(0|1)_{[2]} = \int_{-\infty}^{\lambda_{1,[2]}} f_{Y_2|H}(y|1) dy.\end{aligned}$$

⁴Recall that the subscript [2] denotes the quantity 'seen by' Blake.

When Alexis decides $\hat{H}_1 = 0$, Blake updates his belief q_2 to the posterior q_2^0 :

$$\begin{aligned}\frac{q_2^0}{1 - q_2^0} &= \frac{q_2}{1 - q_2} \frac{1 - P_{e,1,[2]}^I}{P_{e,1,[2]}^{II}} \\ \implies q_2^0 &= \frac{q_2(1 - P_{e,1,[2]}^I)}{q_2(1 - P_{e,1,[2]}^I) + (1 - q_2)P_{e,1,[2]}^{II}},\end{aligned} \quad (22)$$

his decision threshold is $\lambda_2^0 := \lambda_2(q_2^0)$, and the probabilities of error are

$$\begin{aligned}P_{e,2}^{I_0} &= p_{\hat{H}_2|\hat{H}_1,H}(1|0,0) = \int_{\lambda_2^0}^{\infty} f_{Y_2|H}(y|0) dy, \\ P_{e,2}^{II_0} &= p_{\hat{H}_2|\hat{H}_1,H}(0|0,1) = \int_{-\infty}^{\lambda_2^0} f_{Y_2|H}(y|1) dy.\end{aligned}$$

Likewise, when Alexis decides $\hat{H}_1 = 1$, Blake updates his belief q_2 to the posterior q_2^1 :

$$\begin{aligned}\frac{q_2^1}{1 - q_2^1} &= \frac{q_2}{1 - q_2} \frac{P_{e,1,[2]}^I}{1 - P_{e,1,[2]}^{II}} \\ \implies q_2^1 &= \frac{q_2 P_{e,1,[2]}^I}{q_2 P_{e,1,[2]}^I + (1 - q_2)(1 - P_{e,1,[2]}^{II})},\end{aligned} \quad (23)$$

his decision threshold is $\lambda_2^1 := \lambda_2(q_2^1)$, and the probabilities of error are

$$\begin{aligned}P_{e,2}^{I_1} &= p_{\hat{H}_2|\hat{H}_1,H}(1|1,0) = \int_{\lambda_2^1}^{\infty} f_{Y_2|H}(y|0) dy, \\ P_{e,2}^{II_1} &= p_{\hat{H}_2|\hat{H}_1,H}(0|1,1) = \int_{-\infty}^{\lambda_2^1} f_{Y_2|H}(y|1) dy.\end{aligned}$$

Now we compute the system's Bayes risk (or Blake's Bayes risk) R_2 :

$$\begin{aligned}R_2 &= c_{10} p_{\hat{H}_2,H}(1,0) + c_{01} p_{\hat{H}_2,H}(0,1) \\ &= c_{10} \sum_{\hat{h}_1 \in \{0,1\}} p_{\hat{H}_2|\hat{H}_1,H}(1|\hat{h}_1,0) p_{\hat{H}_1|H}(\hat{h}_1|0) p_H(0) \\ &\quad + c_{01} \sum_{\hat{h}_1 \in \{0,1\}} p_{\hat{H}_2|\hat{H}_1,H}(0|\hat{h}_1,1) p_{\hat{H}_1|H}(\hat{h}_1|1) p_H(1) \\ &= c_{10} [P_{e,2}^{I_0}(1 - P_{e,1}^I) + P_{e,2}^{I_1} P_{e,1}^I] p_0 \\ &\quad + c_{01} [P_{e,2}^{II_0} P_{e,1}^{II} + P_{e,2}^{II_1}(1 - P_{e,1}^{II})] (1 - p_0).\end{aligned} \quad (24)$$

Note that the Bayes risk R_2 in (24) is a function of q_1 and q_2 . One might think that R_2 is minimum at $q_1 = q_2 = p_0$ as Alexis makes the best decision for the true prior and Blake does not misunderstand her decision. Surprisingly, however, this turns out to not be true. We prove this by studying Alexis's optimal belief q_1^* that minimizes R_2 .

Theorem 4: Alexis's and Blake's optimal beliefs q_1^*, q_2^* that minimize R_2 satisfy

$$\frac{q_1^*}{1 - q_1^*} = \frac{p_0(P_{e,2}^{I_1} - P_{e,2}^{I_0})}{(1 - p_0)(P_{e,2}^{II_0} - P_{e,2}^{II_1})}. \quad (25)$$

Before proceeding to the proof, note that error probability terms in the right-side are dependent on q_2 , but not on q_1 . Furthermore, the value of $(P_{e,2}^{I_1} - P_{e,2}^{I_0})/(P_{e,2}^{II_0} - P_{e,2}^{II_1})$ is generally not 1, i.e., in general $q_1 = q_2 = p_0$ is not the optimal belief. For example, for the additive Gaussian noise model considered in the next section, the ratio is not equal to 1 except when $p_0 = c_{01}/(c_{01} + c_{10})$.

Proof of Thm. 4: Let us consider the first derivative of (24) with respect to q_1 :

$$\begin{aligned} \frac{\partial R_2}{\partial q_1} &= c_{10}p_0(P_{e,2}^{I_1} - P_{e,2}^{I_0}) \frac{\partial P_{e,1}^I}{\partial q_1} \\ &+ c_{01}(1 - p_0)(P_{e,2}^{II_0} - P_{e,2}^{II_1}) \frac{\partial P_{e,1}^{II}}{\partial q_1}. \end{aligned}$$

We want to find q_1 that minimizes R_2 , i.e., q_1 makes the first derivative zero. Using

$$\begin{aligned} \frac{dP_{e,1}^I}{dq_1} &= \frac{dP_{e,1}^I}{d\lambda_1} \frac{d\lambda_1}{dq_1} = -f_{Y_1|H}(\lambda_1|0) \frac{d\lambda_1}{dq_1}, \\ \frac{dP_{e,1}^{II}}{dq_1} &= \frac{dP_{e,1}^{II}}{d\lambda_1} \frac{d\lambda_1}{dq_1} = f_{Y_1|H}(\lambda_1|1) \frac{d\lambda_1}{dq_1}; \end{aligned}$$

this occurs when

$$\begin{aligned} &c_{10}p_0(P_{e,2}^{I_1} - P_{e,2}^{I_0})f_{Y_1|H}(\lambda_1|0) \\ &= c_{01}(1 - p_0)(P_{e,2}^{II_0} - P_{e,2}^{II_1})f_{Y_1|H}(\lambda_1|1) \\ \Leftrightarrow &\frac{f_{Y_1|H}(\lambda_1|1)}{f_{Y_1|H}(\lambda_1|0)} = \frac{c_{10}p_0(P_{e,2}^{I_1} - P_{e,2}^{I_0})}{c_{01}(1 - p_0)(P_{e,2}^{II_0} - P_{e,2}^{II_1})}. \end{aligned} \quad (26)$$

Note that $\lambda_1 = \lambda_1(q_1)$ is the solution to (4),

$$\frac{f_{Y_1|H}(\lambda_1|1)}{f_{Y_1|H}(\lambda_1|0)} = \frac{c_{10}q_1}{c_{01}(1 - q_1)}. \quad (27)$$

Equating (26) and (27) completes the proof. \blacksquare

The theorem considers general continuous likelihoods $\{f_{Y_n|H}\}$ with the monotonicity assumption on $\lambda(q)$. It is interesting to evaluate the optimal beliefs in the case of Gaussian likelihoods (i.e., additive Gaussian noise) and obtain insights into optimality in the sequential decision-making problem.

V. GAUSSIAN LIKELIHOODS

We now focus on Gaussian likelihoods and study their optimal beliefs in this section. Suppose the n th agent receives the signal $Y_n = H + Z_n$, where Z_n is an independent additive Gaussian noise with zero mean and variance $\sigma_n^2 > 0$. Thus, the received signal probability densities for $H = h$ are

$$f_{Y_n|H}(y_n|h) = \phi(y_n; h, \sigma_n^2).$$

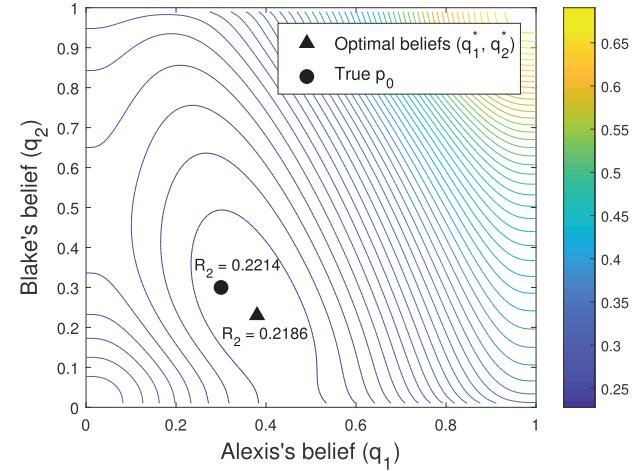


Fig. 3. The Bayes risk for $q_1, q_2 \in (0, 1)$ with $p_0 = 0.3$, $c_{10} = c_{01} = 1$, and additive standard Gaussian noise. The pair of optimal beliefs (\blacktriangle) yields $R_2 = 0.2186$, while the true prior (\bullet) yields $R_2 = 0.2214$.

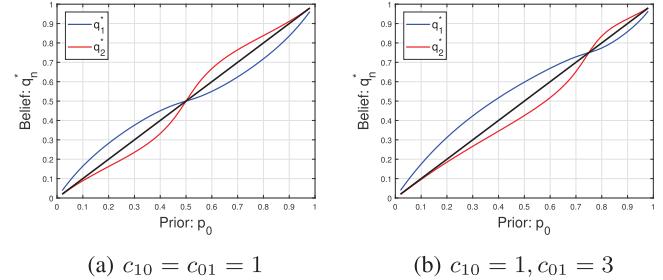


Fig. 4. The trend of the optimal beliefs for $N = 2$ (Alexis, Blake). Z_1, Z_2 are standard Gaussian.

For a belief q_n , the decision threshold is then determined by the likelihood ratio test,

$$\mathcal{L}_n(y_n) = \frac{f_{Y_n|H}(y_n|1)}{f_{Y_n|H}(y_n|0)} \begin{cases} \hat{H}_1=1 & \geq \frac{c_{10}q_n}{c_{01}(1 - q_n)}, \\ \hat{H}_1=0 & \end{cases}$$

that simplifies to the following simple threshold condition for Gaussian likelihoods:

$$y_n \begin{cases} \hat{H}_1=1 & \geq \lambda_n(q_n) = \frac{1}{2} + \sigma_n^2 \log \left(\frac{c_{10}q_n}{c_{01}(1 - q_n)} \right). \\ \hat{H}_1=0 & \end{cases} \quad (28)$$

Here the index n represents the n th agent in the system, as the belief and variance of the agent varies along the chain.

Using the recursive update in Sec. III and decision threshold (28), it is possible to obtain the Bayes risk of Blake (i.e., $N = 2$) for given beliefs q_1, q_2 . Fig. 3 depicts Blake's Bayes risk for $q_1, q_2 \in (0, 1)$, and explicitly shows that knowing true prior probability is not optimal. The social learning problem with Bayes costs $c_{10} = c_{01} = 1$, prior $p_0 = 0.3$, and additive Gaussian noise with zero mean and unit variance results in a Bayes risk that is minimum when Alexis's belief is 0.38 and Blake's belief is 0.23, shown in the figure (triangle) and is compared to the true prior (circle).

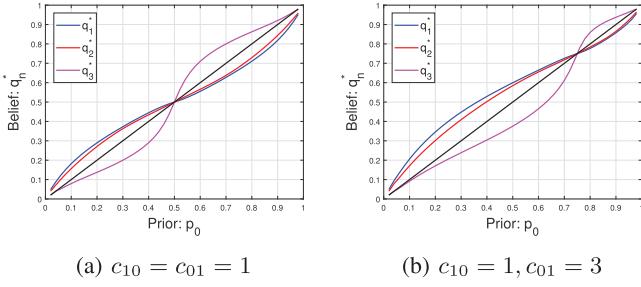


Fig. 5. The trend of the optimal beliefs for $N = 3$ (Alexis, Blake, and Chuck). Z_1, Z_2, Z_3 are standard Gaussian.

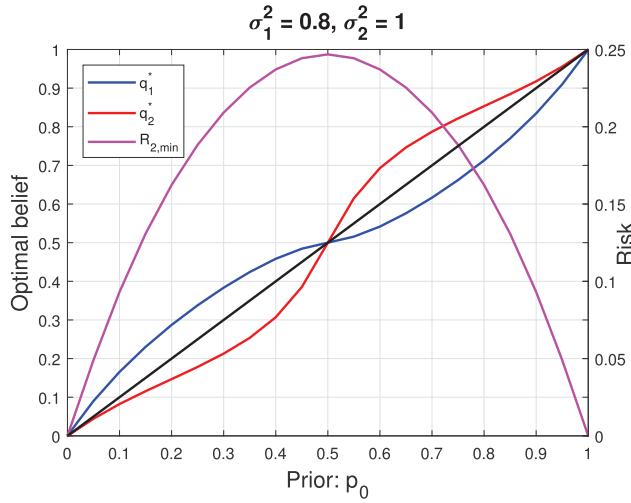


Fig. 6. Optimal beliefs when the preceding agent has smaller noise, where $\sigma_1^2 = 0.8$ and $\sigma_2^2 = 1$.

Figs. 4 and 5 show the trend of optimal belief pair that minimizes the last agent's Bayes risk, when all agents have the same noise levels for the case of two and three agents respectively. We can observe several common characteristics. First, the non-terminal agents (i.e., Alexis for $N = 2$ and Alexis and Blake for $N = 3$) overweight their beliefs if p_0 is small and underweight it if p_0 is large. We call this *open-minded* behavior as it enhances less likely events. Second, the last agent (i.e., Blake for $N = 2$ and Chuck for $N = 3$) underweights the belief if p_0 is small and overweights it if p_0 is large, implicitly compensating for the biases of the preceding agents. Such behavior is referred to as being *closed-minded* as it represents a cautious outlook to the decision-making problem. Lastly, there is a unique, non-trivial prior, $p_0 \in (0, 1)$, where all agents' optimal beliefs are identical to the true prior.

However, the case of nonidentical noise variances of agents results in a very different behavior of optimal beliefs, especially when the last agent has smaller noise. The optimal beliefs for $N = 2$ and the case of the preceding agent having smaller noise, and that of the last agent having smaller noise respectively are shown in Figs. 6 and 7. As can be observed, the optimal belief curves are markedly different when the last agent has smaller noise, and we now derive some analytical properties of q_1^*, q_2^* .

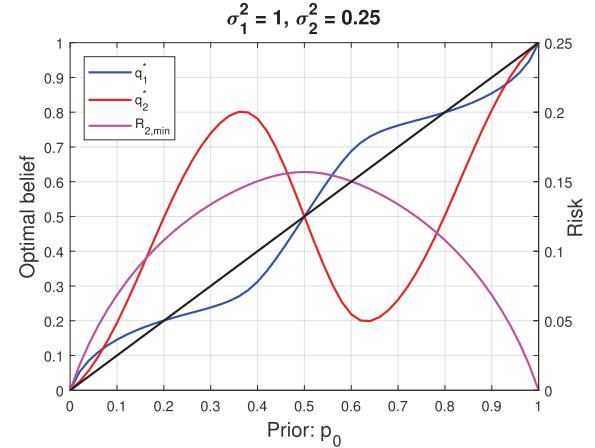


Fig. 7. Optimal beliefs when the later-acting agent has smaller noise, where $\sigma_1^2 = 1$ and $\sigma_2^2 = 0.25$.

Theorem 5: For any σ_1^2 and σ_2^2 , q_1^* and q_2^* satisfy:

- 1) for $p_0 \in (0, 1)$, $q_1^* \leq p_0$ if and only if $q_2^* \geq \frac{c_{01}}{c_{01} + c_{10}}$, with equality for $q_2^* = \frac{c_{01}}{c_{01} + c_{10}}$.
- 2) $p_0 = q_1^* = q_2^*$ if and only if $p_0 \in \{0, \frac{c_{01}}{c_{01} + c_{10}}, 1\}$.

Proof: Given in App. A. \blacksquare

Thm. 5 highlights the fact that if the last agent believes the null hypothesis is more likely, then the ideal predecessor underweights the prior, and vice versa. Additionally, for p_0 near zero (near one) the optimal predecessor overweights (underweights) the prior.

In particular, let us consider two cases separately. First, let the predecessor have smaller noise. Then the curves for optimal beliefs and the corresponding Bayes risk are as shown in Fig. 6. The behavior here is similar to the case with equal noise, indicating that the reducing noise of the predecessor does not alter the overall behaviors of beliefs, as the last agent is unaware of this improved signal quality.

On the other hand, when the last agent has smaller noise, we notice that the nature of curves changes, as shown in Fig. 7. The behavior of the ideal agents indicates that when the predecessor has significantly larger noise than the last agent, the last agent stays open-minded. In addition, q_1^* has multiple crossings with p_0 , but q_2^* has a single crossing at $q_2^* = c_{01}/(c_{01} + c_{10})$.

As expected, the ideal predecessor is open-minded for near-deterministic priors (p_0 close to zero or one). However, when the prior uncertainty in the hypotheses is high (p_0 near 1/2), we note that the ideal last agent with less noise favors the less likely hypothesis. This can be attributed to the fact that the last agent stays open-minded to the less likely hypothesis when the predecessor with larger noise is more likely to make errors. To further understand the nature of such an predecessor, we characterize the crossings of the optimal belief curve with the prior $q_1^* = p_0$.

Theorem 6: The set of all p_0 such that $q_1^* = p_0, q_2^* = \frac{c_{01}}{c_{01} + c_{10}}$ is given by the solutions to

$$e^x = \frac{1 - \beta Q(-\alpha + \sigma_1 x)}{1 - \beta Q(-\alpha - \sigma_1 x)}, \quad (29)$$

where

$$x = \log \left(\frac{c_{10}p_0}{c_{01}(1-p_0)} \right), \quad \alpha = \frac{1}{2\sigma_1}, \quad \beta = 1 - \frac{Q(1/2\sigma_2)}{Q(-1/2\sigma_2)}.$$

Proof: Given in App. B. \blacksquare

We note that $p^* = \frac{c_{01}}{c_{01}+c_{10}}$ is always a solution to (29). The case of multiple solutions to (29) is of particular interest and a sufficient condition is given in the following corollary.

Corollary 2: If

$$\frac{2\beta\sigma_1\phi(\alpha)}{1-\beta Q(-\alpha)} > 1, \quad (30)$$

then, (29) has at least 3 solutions in $(0,1)$.

Proof: Since x is a monotonic function of p_0 , it is sufficient to show that (29) has at least 3 solutions in x . From the symmetry in (29), since $x = 0$ is always a root, it suffices to show the existence of at least one more root in $x > 0$. First note the ranges of variables, $x \in (-\infty, \infty)$, $\alpha \in (0, \infty)$, $\beta \in (0, 1)$.

Letting $r(x)$ be the right side of (29), since $0 \leq Q(\cdot) \leq 1$, we have

$$1 - \beta \leq r(x) := \frac{1 - \beta Q(-\alpha + \sigma_1 x)}{1 - \beta Q(-\alpha - \sigma_1 x)} \leq \frac{1}{1 - \beta},$$

indicating that $r(x) \in [1 - \beta, \frac{1}{1-\beta}]$. However, note that e^x monotonically increases in $(1, \infty)$ for $x > 0$. Since $e^x, r(x)$ coincide at $x = 0$, it follows that they cross at least once on $(0, \infty)$ and also on $(-\infty, 0)$, if $r'(x) > \frac{d}{dx}e^x$ at $x = 0$ by the intermediate value theorem. Thus, the sufficient condition follows:

$$r'(0) = \frac{2\sigma_1\beta\phi(\alpha; 0, 1)}{1 - \beta Q(-\alpha)} > 1 = \frac{d}{dx}e^x \Big|_{x=0}.$$

Cor. 2 provides a sufficient condition on the noise level of agents under which there exists multiple crossings of the curves $q_1^*(p_0)$ and p_0 . The range of standard deviations of the additive Gaussian noise of the preceding and last agents that satisfy the sufficient condition of Cor. 2 is shown in Fig. 8. Note from the figure that the area below the red dotted contour in Fig. 8 has multiple solutions to $q_1^* = p_0$, i.e., when the last agent has comparatively smaller than the preceding agent.

This is important as the crossings indicate a change in the perceived bias of the predecessor and also indicates the regions in which the last agent overweights the unlikely hypothesis as in Fig. 7.

VI. TEAM CONSTRUCTION CRITERION

Having studied the mathematical conditions for optimal reweighting of initial beliefs, we now investigate team selection for social learning. Naturally, a social planner who is aware of the context p_0 can pick the optimal agent pairs to minimize Bayes risk. However, it is not clear if agents are capable of organizing themselves into ideal teams in the absence of contextual knowledge. Thus, we now identify the criterion for the last agent to identify the optimal predecessors among a set of given predecessors.

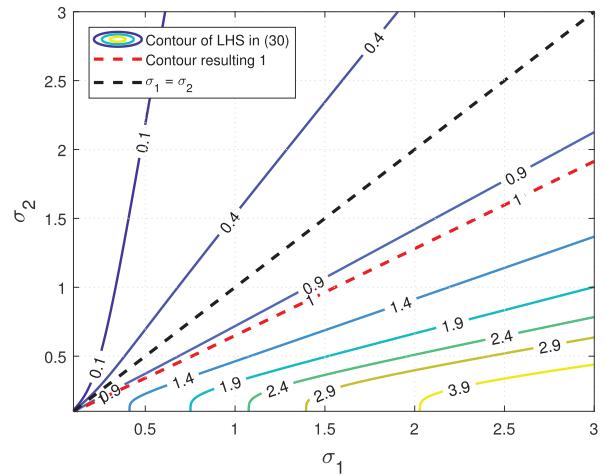


Fig. 8. Contour plot of (30) with values for various σ_1, σ_2 . The red dotted contour shows the contour that results in 1 so that the area below it satisfies (30) and therefore has multiple solutions to (29).

Theorem 7: Consider two predecessors with $q_1 < q_{1'}$. Let $\lambda_1, \lambda_{1'}$ be the decision thresholds of the respective predecessors. Then, the predecessor with belief q_1 is the optimal choice if and only if

$$\frac{\mathbb{P}_1[Y_1 \in [\lambda_1, \lambda_{1'}], Y_2 \in [\lambda_2^1, \lambda_2^{0'}]]}{\mathbb{P}_0[Y_1 \in [\lambda_1, \lambda_{1'}], Y_2 \in [\lambda_2^1, \lambda_2^{0'}]]} \geq \frac{c_{10}p_0}{c_{01}(1-p_0)}. \quad (31)$$

Proof: Given in App. C. \blacksquare

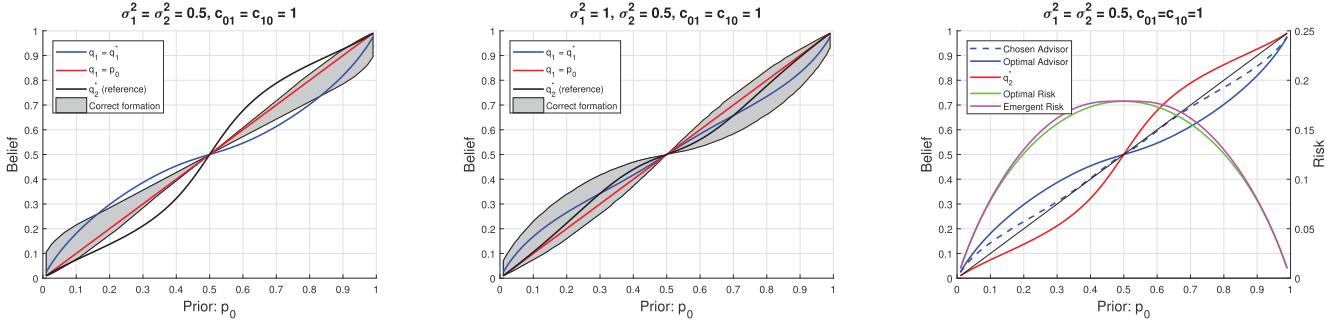
In other words, by rewriting (31) in a likelihood ratio form, we observe that the criterion for picking the predecessor with a smaller belief is given by the likelihood ratio test

$$\mathcal{L} \left[\widehat{H}_1 = \widehat{H}_2 = 1, \widehat{H}_{1'} = \widehat{H}_{2'} = 0 \right] \geq \frac{c_{10}p_0}{c_{01}(1-p_0)},$$

where $\widehat{H}_{2'}$ is the decision made by the last agent following the decision of the predecessor with belief $q_{1'}$.

Thus selecting an ideal predecessor requires a social planner who is aware of the context p_0 . Without this, the last agent selects an predecessor according to his personal belief q_2 . That is, the last agent verifies condition (31) by replacing p_0 by q_2 . Such a choice of predecessor might not always conform to the optimal choice when the belief of the last agent deviates significantly from the prior. To illustrate, we consider the problem of choosing between two predecessors with beliefs $q_1(p_0) = q_1^*(p_0)$ and $q_{1'}(p_0) = p_0$. Let $q(p_0, q_2)$ be the belief of the optimal predecessor choice for a given pair (p_0, q_2) . We identify the region of correct selection by shading, $\mathcal{S} = \{(p_0, q_2) : q(p_0, q_2) = q(q_2, q_2)\}$.

First, when noise levels are equal, the region in which the last agent picks the correct preceding agent is shown in Fig. 9(a). We note that the correct region is relatively small and does not include q_2^* . In particular, the last agent with optimal belief chooses the wrong predecessor always, whereas a suboptimal last agent with beliefs in the shaded region picks the correct one.



(a) Predecessor selection under equal noise. The optimal later-acting agent fails to recognize the optimal predecessor and makes mistakes often. The trend is similar in cases with smaller noise predecessor.

(b) Predecessor selection under noise diversity. The optimal later-acting agent selects the optimal predecessor in this case.

(c) Context-unaware predecessor selection. The last agent chooses predecessor using (32) without p_0 and the Bayes risk increases as a result.

Fig. 9. Context-unaware team selection.

On the other hand, when the last agent has smaller noise than the predecessor, the corresponding region is as shown in Fig. 9(b). Here we note that the last agent with optimal belief picks the correct preceding agent always.

Thus, we note that knowledge of the mathematically optimal beliefs does not guarantee selection of the right preceding agent. Further, we also observe that the diversity of noise levels may increase the feasibility of selecting the right preceding agent when the last agent has optimal belief.

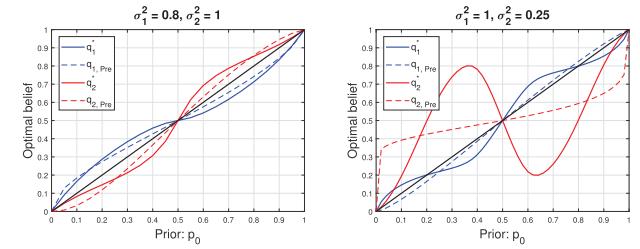
We also explore the optimal choice of predecessor for the given optimal last agent in the absence of knowledge of the prior probability. From (25), the belief of the optimal preceding agent, \tilde{q}_1 chosen by an last agent, in the absence of context (prior probability p_0) satisfies

$$\frac{\tilde{q}_1}{1 - \tilde{q}_1} = \frac{p_0}{1 - p_0} \frac{P_{e,2}^{I_1} - P_{e,2}^{I_0}}{P_{e,2}^{II_0} - P_{e,2}^{II_2}}. \quad (32)$$

The last agent's behavior with belief q_2^* is as shown in Fig. 9(c). We note that the preceding agent chosen by the last agent differs from the optimal choice. Further, it is also evident that this choice consequently results in an increased Bayes risk. Such behavior in team selection highlights the significance of context and thus a social planner for identifying the right team.

VII. HUMAN-AI COLLABORATION SYSTEMS

In this section, we use mathematical results from previous sections to study the engineering design problem of constructing human-AI collaborative systems. To do so, we make the following assumptions from the behavioral sciences: Human agents perform Bayesian decision-making [42]–[45] and their perceptions follow the Prelec reweighting function [38]. In addition, agents experience varying observational noise which is additive and Gaussian (as it is a common model in human signal perception [46], [47]). As usual in sequential social learning setup, all agents make selfish decisions [6], [10].



(a) When the preceding agent (b) When the later-acting agent has smaller noise has smaller noise

Fig. 10. Optimal beliefs as compared to Prelec-weighted beliefs.

A. Approximation by Prelec Family

To design human-AI collaborative systems, we first determine whether optimal belief functions from previous sections are close to human behavior as modeled by cumulative prospect theory [35], [38].⁵

We approximate the optimal belief curves q_n^* by the Prelec function and study the resulting increase in the Bayes risk. We restrict to the Prelec family whose fixed point is identical to $p^* = \frac{c_{10}}{c_{01} + c_{10}}$, and then find best parameters (α_n, β_n) in the minimax absolute error sense, i.e.,

$$(\alpha_n, \beta_n) = \arg \min_{\alpha, \beta: w(p^*; \alpha, \beta) = p^*} \|q_n^*(\cdot) - w(\cdot; \alpha, \beta)\|_\infty.$$

Let the Prelec function approximations be $(q_{1, \text{Pre}}, q_{2, \text{Pre}})$.

The Prelec approximations for the two-agent case are shown in dotted curves in Fig. 10. When the preceding agent has smaller noise as in Fig. 10(a), the Prelec function approximates the

⁵Bounded rationality models have been categorized into two main classes—costly bounded rationality and truly bounded rationality [48]. Costly rationality considers the emergence of boundedly irrational behavior as optimization under some costs of decision-making such as computation and communication. On the other hand, truly bounded rationality is not based on an optimization framework. Though not the focus of the present paper, one might wonder whether people are (approximately) naturally optimal for social learning. That is, do cumulative prospect-theoretic models emerge from a costly rationality framework for social learning, since the optimal belief curves result from limitations in computation (selfish decision-making) and communication (public signal quantization).

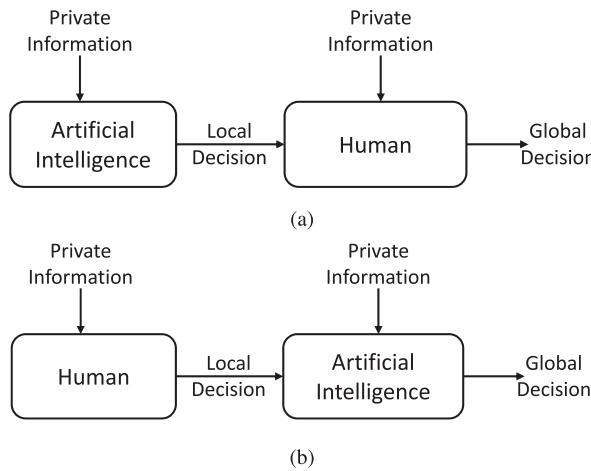


Fig. 11. Models of AI-human collaboration, where a machine provides input for human judgement or vice versa.

optimal beliefs well and the Bayes risk does not increase by much. To evaluate the loss from the approximation, consider the set of correct beliefs $q_1 = q_2 = p_0$, that result in a Bayes risk of $R_{2,\text{corr}}$. The maximal loss in terms of Bayes risk from using the correct beliefs is $\max_{p_0} (R_{2,\text{corr}} - R_{2,\text{min}}) \approx 0.0039$. On the other hand, the maximal loss from the best Prelec approximation is ≈ 0.0009 . This indicates that the natural cognitive biases of humans (i.e., Prelec reweighting) are effective for social learning when the preceding agent has smaller noise.

On the other hand, when the last agent has smaller noise as in Fig. 10(b), the Prelec approximation does not accurately mimic the optimal behavior of agents. Recall that the Prelec function is always increasing and has only one crossing with unit slope line in $(0,1)$. Therefore, the Prelec function fails to account for all the variations in the optimal belief. Moreover, while the additional loss of Bayes risk by the Prelec fitting is ≈ 0.0187 , the loss from using the correct beliefs, $p_0 = q_1 = q_2$, is ≈ 0.0060 . This indicates that even though the Prelec weighting functions serve as good approximations with predecessors having less noisy observations, they do not model the optimal behavior in the case of predecessors having noisier observations. These results suggest that human agents following cumulative prospect theory models [35] yield small Bayes risk when predecessors have smaller noise.

B. Human-AI Teams

The previous subsection informs the design of AI-human collaboration structures [34]. In many human-AI joint teams, a human agent makes the final decision based on the advice of an AI component as depicted in Fig. 11(a), but the opposite structure of Fig. 11(b) is also possible. It is thus important to identify the best team configuration [48]. Indeed, D. Kahneman recently stated that “You can combine humans and machines, provided the machine has the last word” [49].

Our results indicate that an AI assistant with smaller noise could be an effective predecessor to the human decision-maker. In particular, an open-minded AI predecessor and a closed-minded human final decision-maker with appropriate Prelec

reweighted beliefs work well together, as in Fig. 6. However, an AI component with greater noise might not be a good predecessor to the human last agent who does not have beliefs that mimic the optimal behavior in Fig. 7 and so perhaps counterintuitively, the architecture of Fig. 11(b) should be adopted, with the AI agent having larger noise making the global decision.

Additionally, these results along with those of Sec. VI provide some insight into human-AI teams when the human agent picks an AI predecessor, given a choice among different agents. In particular, consider the AI-human team where the human, who has a Prelec-weighted belief, chooses one of two possible AI predecessors—one that has the optimal belief q_1^* and the other that is aware of the true prior p_0 . In case the human agent has larger noise, and a closed-minded Prelec belief as in Fig. 9(a), she unfortunately picks the AI predecessor with $q_1 = p_0$ and the team becomes suboptimal. However, if the human agent has smaller noise, and an open-minded Prelec belief, she picks the optimal AI component $q_1 = q_1^*$ and therefore can make the optimal decision as in Fig. 9(b). Thus it is evident that optimal team organization is feasible when the human has smaller noise and the appropriate open-minded belief.

VIII. CONCLUSION

We discussed the sequential social learning problem with individual biased beliefs. Unlike previous works on herding, we focused on the Bayes risk of the last-acting agent. We first derived the optimal belief update rule for general likelihoods and evaluated for Gaussian likelihoods. Counterintuitively, optimal beliefs that yield minimum Bayes risk are in general different from the true prior. Under equal noise levels, we observed that optimal preceding agents have open-minded beliefs, that is, overweight small priors and underweight large priors, while the optimal last agent has closed-minded belief. However, the trend may change depending on varying noise levels such that especially when the last agent has much smaller noise, optimal belief of the last agent is inverted as she becomes open-minded.

We also showed that the Prelec reweighting function from cumulative prospect theory approximates the behavior of the optimal beliefs under specific levels of noise, however, when the last agent has much smaller noise, it fails to capture all the behavioral traits of the optimal beliefs.

Finally, we considered the ability of agents to organize themselves into optimal teams and showed that in the absence of a social planner, the last agent can get paired with the wrong predecessor when the individual belief deviates significantly from the underlying prior value. The setup arises from the consideration of AI and it tells us without knowing the true prior, our human-machine team construction could be misorganized.

APPENDIX A PROOF OF THEOREM 5

Let us prove Thm. 5 starting with the premise that $q_1^* \geq p_0$. First, from (25), we have

$$q_1^* \geq p_0 \Leftrightarrow \frac{P_{e,2}^{I_1} - P_{e,2}^{I_0}}{P_{e,2}^{I_1} - P_{e,2}^{I_0}} \geq -1. \quad (33)$$

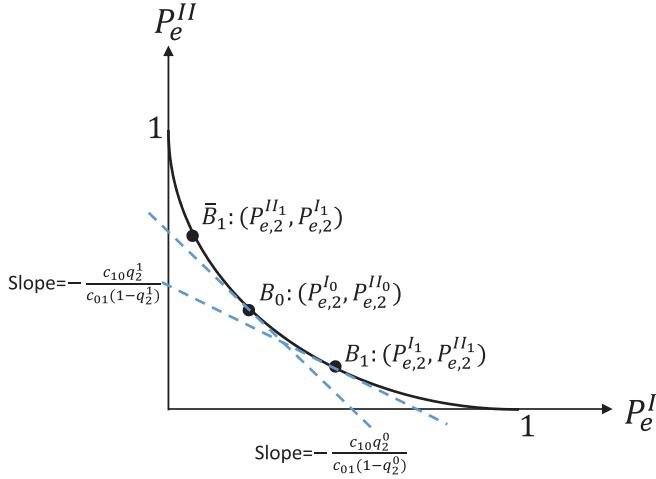


Fig. 12. The point B_0 always exists between points B_1 and \bar{B}_1 .

To study the ratio in (33), consider the Type I vs. Type II error curve for binary hypothesis testing under additive Gaussian noise.⁶ This is shown in Fig. 12, and as seen here is a convex function [50]. Note that on the curve, the Type I and Type II error probabilities, (P_e^I, P_e^{II}) , are the points on the curve that have tangents with slope matching $-(\frac{c_{10}q}{c_{01}(1-q)})$, where q is the corresponding probability, and σ^2 is the variance of the additive Gaussian noise.

First, from Thm. 2, we know that $q_2^0 \geq q_2^1$ which in turn implies that $\lambda_2^0 \geq \lambda_2^1$. This in turn indicates that

$$P_{e,2}^{I_0} = Q\left(\frac{\lambda_2^0}{\sigma_2}\right) \leq Q\left(\frac{\lambda_2^1}{\sigma_2}\right) = P_{e,2}^{I_1}.$$

Similarly, $P_{e,2}^{I_0} \geq P_{e,2}^{II_1}$, and thus, as shown in the figure, the point $B_0 = (P_{e,2}^{I_0}, P_{e,2}^{II_0})$ lies to the left of $B_1 = (P_{e,2}^{I_1}, P_{e,2}^{II_1})$.

Further, since B_1 lies on the curve, so does the point $\bar{B}_1 = (P_{e,2}^{II_1}, P_{e,2}^{I_1})$ as it caters to the error probabilities corresponding to the probability of the null hypothesis $\mathbb{P}[H = 0] = 1 - q_2^1$. Thus, the line $\overline{B_1 \bar{B}_1}$ has a slope of -1 .

Note that the condition (33) translates to the slope of the line $\overline{B_0 B_1}$ is greater than -1 . Observe that if \bar{B}_1 lies to the right of B_1 then it implies that the slope of $\overline{B_0 B_1}$ is less than -1 , violating (33). Similarly, if B_0 lies to the left of \bar{B}_1 , then again the (33) is violated.

On the other hand, if B_0 lies between \bar{B}_1 and B_1 , then we know that the slope of $\overline{B_0 B_1}$ is greater than that of $\overline{B_1 \bar{B}_1}$, therein satisfying (33). Thus, (33) is true if and only if the point B_0 lies between the two points B_1 and \bar{B}_1 .

From the convexity of the curve and comparing coordinates of B_0 and \bar{B}_1 , we have

$$q_1^* \geq p_0$$

$$\Leftrightarrow P_{e,2}^{I_0} \geq P_{e,2}^{II_1} \text{ and } P_{e,2}^{II_0} \leq P_{e,2}^{I_1}$$

⁶It is also called Receiver Operating Characteristic (ROC) curve [50], [51] when the curve is vertically inverted.

$$\begin{aligned} &\stackrel{(a)}{\Leftrightarrow} Q\left(\frac{\lambda_2^0}{\sigma_2}\right) \geq 1 - Q\left(\frac{\lambda_2^1 - 1}{\sigma_2}\right) \\ &\text{and } Q\left(\frac{\lambda_2^1}{\sigma_2}\right) \geq 1 - Q\left(\frac{\lambda_2^0 - 1}{\sigma_2}\right) \\ &\stackrel{(b)}{\Leftrightarrow} \lambda_2^0 + \lambda_2^1 \leq 1 \\ &\stackrel{(c)}{\Leftrightarrow} 2\lambda_{1,[2]} + \sigma_2^2 \log\left(\frac{P_{e,1,[2]}^I (1 - P_{e,1,[2]}^I)}{P_{e,1,[2]}^{II} (1 - P_{e,1,[2]}^{II})}\right) \leq 1, \quad (34) \end{aligned}$$

where (a) follows from the false alarm and missed detection probabilities in terms of the Q -function of the standard Gaussian random variable; (b) follows from the fact that the Q -function is monotonically decreasing and that $1 - Q(x) = Q(-x)$; and (c) follows from (22), (23), and $\lambda_{1,[2]} = \lambda_2(q_2)$.

From (28), we have

$$\lambda_{1,[2]} = \frac{1}{2} + \sigma_2^2 \log\left(\frac{c_{10}q_2^*}{c_{01}(1 - q_2^*)}\right).$$

Substituting in (34), we have

$$\begin{aligned} q_1^* \geq p_0 &\Leftrightarrow 2 \log\left(\frac{c_{10}q_2^*}{c_{01}(1 - q_2^*)}\right) \\ &\leq \log\left(\frac{P_{e,1,[2]}^{II} (1 - P_{e,1,[2]}^{II})}{P_{e,1,[2]}^I (1 - P_{e,1,[2]}^I)}\right). \end{aligned}$$

Letting $x := \log(\frac{c_{10}q_2^*}{c_{01}(1 - q_2^*)}) = \frac{1}{\sigma_2^2}(\lambda_2 - \frac{1}{2})$ and using $Q(\cdot)$ representation of error probabilities, we have

$$q_1^* \geq p_0 \Leftrightarrow 2x \leq \log\left(\frac{Q\left(\sigma_2x - \frac{1}{2\sigma_2}\right)Q\left(-\sigma_2x + \frac{1}{2\sigma_2}\right)}{Q\left(\sigma_2x + \frac{1}{2\sigma_2}\right)Q\left(-\sigma_2x - \frac{1}{2\sigma_2}\right)}\right). \quad (35)$$

From Cor. 1, we know that the function

$$\tilde{g}(x) = x + \log\left(\frac{Q\left(\sigma x + \frac{1}{2\sigma}\right)}{Q\left(\sigma x - \frac{1}{2\sigma}\right)}\right)$$

is an increasing function of x . Thus, reformulating (35) using $\tilde{g}(\cdot)$,

$$\begin{aligned} q_1^* \geq p_0 &\Leftrightarrow \tilde{g}(x) \leq \tilde{g}(-x) \\ &\Leftrightarrow x \leq 0 \Leftrightarrow q_2^* \leq \frac{c_{01}}{c_{01} + c_{10}}. \end{aligned}$$

The condition for equality follows from observing the condition for equality at all the inequalities, proving the first part of the result.

The second part follows directly from the first, taking into account the trivial cases of $p_0 \in \{0, 1\}$.

APPENDIX B PROOF OF THEOREM 6

We will consider the case of $c_{01} = c_{10} = 1$ for convenience. The proof extends directly by a simple scaling argument.

The optimal belief of worker two satisfies $\frac{\partial R_2}{\partial q_2} = 0$. Thus, differentiating (24) with respect to q_2 and rearranging,

$$p_0 \left[(1 - P_{e,1}^I) f_{Y_2|H}(\lambda_2^0|0) \frac{\partial \lambda_2^0}{\partial q_2} + P_{e,1}^I f_{Y_2|H}(\lambda_2^1|0) \frac{\partial \lambda_2^1}{\partial q_2} \right] \\ = (1 - p_0) \left[P_{e,1}^{\text{II}} f_{Y_2|H}(\lambda_2^0|1) \frac{\partial \lambda_2^0}{\partial q_2} + (1 - P_{e,1}^{\text{II}}) f_{Y_2|H}(\lambda_2^1|1) \frac{\partial \lambda_2^1}{\partial q_2} \right].$$

Let $x = \log \left(\frac{p_0}{1-p_0} \right)$. For $q_2^* = 1/2$ and $q_1^* = p_0$, we have

$$\lambda_1 = \frac{1}{2} + \sigma_1^2 x \text{ and } \lambda_{1,[2]} = \frac{1}{2}.$$

It implies $P_{e,1,[2]}^I = P_{e,1,[2]}^{\text{II}} = Q(1/2\sigma_2)$. Then,

$$\mathcal{L}(\lambda_2^0) = \frac{f_{Y_2|H}(\lambda_2^0|1)}{f_{Y_2|H}(\lambda_2^0|0)} \\ = \frac{q_2}{1-q_2} \frac{(1 - P_{e,1,[2]}^I)}{P_{e,1,[2]}^{\text{II}}} = \frac{Q(-1/2\sigma_2)}{Q(1/2\sigma_2)} =: \frac{1}{c},$$

$$\mathcal{L}(\lambda_2^1) = \frac{f_{Y_2|H}(\lambda_2^1|1)}{f_{Y_2|H}(\lambda_2^1|0)} \\ = \frac{q_2}{1-q_2} \frac{P_{e,1,[2]}^I}{(1 - P_{e,1,[2]}^{\text{II}})} = \frac{Q(1/2\sigma_2)}{Q(-1/2\sigma_2)} = c.$$

Equivalently, this implies that

$$\lambda_2^0 = \frac{1}{2} + \sigma^2 \log \left(\frac{1}{c} \right), \quad \lambda_2^1 = \frac{1}{2} - \sigma^2 \log \left(\frac{1}{c} \right).$$

Thus, $\lambda_2^0 + \lambda_2^1 = 1$, and so,

$$f_{Y_2|H}(\lambda_2^1|1) = \frac{1}{\sqrt{2\pi\sigma_2}} \exp \left(-\frac{(\lambda_2^1 - 1)^2}{2\sigma_2^2} \right) \\ = \frac{1}{\sqrt{2\pi\sigma_2}} \exp \left(-\frac{(\lambda_2^0)^2}{2\sigma_2^2} \right) \\ = f_{Y_2|H}(\lambda_2^0|0).$$

Similarly, we also have

$$f_{Y_2|H}(\lambda_2^1|0) = f_{Y_2|H}(\lambda_2^0|1).$$

Further, from (22) and (23), we have

$$\frac{d\lambda_2^0}{dq_2} = \frac{d\lambda_2^0}{d\lambda_{1,[2]}} \frac{d\lambda_{1,[2]}}{dq_2} \\ = \left[\frac{\sigma_2^2 \phi \left(\frac{\lambda_{1,[2]}}{\sigma_2} \right)}{1 - P_{e,1,[2]}^I} - \frac{\sigma_2^2 \phi \left(\frac{\lambda_{1,[2]} - 1}{\sigma_2} \right)}{P_{e,1,[2]}^{\text{II}}} \right] \frac{d\lambda_{1,[2]}}{dq_2},$$

$$\frac{d\lambda_2^1}{dq_2} = \frac{d\lambda_2^1}{d\lambda_{1,[2]}} \frac{d\lambda_{1,[2]}}{dq_2} \\ = \left[\frac{\sigma_2^2 \phi \left(\frac{\lambda_{1,[2]}}{\sigma_2} \right)}{1 - P_{e,1,[2]}^I} + \frac{\sigma_2^2 \phi \left(\frac{\lambda_{1,[2]} - 1}{\sigma_2} \right)}{1 - P_{e,1,[2]}^{\text{II}}} \right] \frac{d\lambda_{1,[2]}}{dq_2}.$$

When, $\lambda_{1,[2]} = \frac{1}{2}$, $P_{e,1,[2]}^I = P_{e,1,[2]}^{\text{II}} = Q(\frac{1}{2\sigma_2})$, and $\phi(\frac{\lambda_{1,[2]}}{\sigma_2}) = \phi(\frac{\lambda_{1,[2]} - 1}{\sigma_2})$. Thus, $\frac{d\lambda_2^0}{dq_2} = \frac{d\lambda_2^1}{dq_2}$.

Using these, the values of prior for which $q_1^* = p_0$, $q_2^* = 1/2$ are given by

$$\frac{p_0}{1-p_0} = \frac{Q\left(\frac{-1}{2\sigma_2}\right)Q\left(\frac{-1}{2\sigma_1} - \sigma_1 x\right) + Q\left(\frac{1}{2\sigma_2}\right)Q\left(\frac{1}{2\sigma_1} + \sigma_1 x\right)}{Q\left(\frac{-1}{2\sigma_2}\right)Q\left(\frac{-1}{2\sigma_2} + \sigma_1 x\right) + Q\left(\frac{1}{2\sigma_2}\right)Q\left(\frac{1}{2\sigma_1} - \sigma_1 x\right)}. \quad (36)$$

Using the definitions of x, α, β in (36), and the fact that $Q(-y) = 1 - Q(y)$, the result follows.

APPENDIX C PROOF OF THEOREM 7

From (1), we note that the Bayes risk for social learning with beliefs (q_1, q_2) is

$$R_2(q_1, q_2) = c_{10} p_0 [P_{e,2}^{\text{I}_0} (1 - P_{e,1}^I) + P_{e,2}^{\text{I}_1} P_{e,1}^I] \\ + c_{01} (1 - p_0) [P_{e,2}^{\text{II}_0} P_{e,1}^{\text{II}} + P_{e,2}^{\text{II}_1} (1 - P_{e,1}^{\text{II}})].$$

Then, the difference in Bayes risk between the two choices of advisors is given by

$$\Delta R_2 = R_2(q_1, q_2) - R_2(q_1', q_2) \\ = c_{10} p_0 (P_{e,1}^I - P_{e,1'}^I) (P_{e,2}^{\text{I}_1} - P_{e,2}^{\text{I}_0}) \\ + c_{01} (1 - p_0) (P_{e,1}^{\text{II}} - P_{e,1'}^{\text{II}}) (P_{e,2}^{\text{II}_0} - P_{e,2}^{\text{II}_1}). \quad (37)$$

Since $q_1 < q_1'$, the decision thresholds satisfy $\lambda_1 < \lambda_{1'}$. Thus, from (37) and independence of Y_1, Y_2 given H , we see that $\Delta R_2 \leq 0$ if and only if (31) holds.

REFERENCES

- [1] J. B. Rhim and V. K. Goyal, "Social teaching: Being informative vs. being right in sequential decision making," in *Proc. IEEE Int. Symp. Inf. Theory*, Jul. 2013, pp. 2602–2606.
- [2] D. Seo, R. K. Raman, and L. R. Varshney, "Probability reweighting in social learning: Optimality and suboptimality," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, Apr. 2018, pp. 6966–6970.
- [3] G. Ellison and D. Fudenberg, "Rules of thumb for social learning," *J. Political Economy*, vol. 101, no. 4, pp. 612–643, Aug. 1993.
- [4] 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.
- [5] A. V. Banerjee, "A simple model of herd behavior," *Quart. J. Econ.*, vol. 107, no. 3, pp. 797–817, Aug. 1992.
- [6] S. Bikhchandani, D. Hirshleifer, and I. Welch, "Learning from the behavior of others: Conformity, fads, and informational cascades," *J. Econ. Perspectives*, vol. 12, no. 3, pp. 151–170, 1998.
- [7] V. Bala and S. Goyal, "Conformism and diversity under social learning," *Econ. Theory*, vol. 17, no. 1, pp. 101–120, Jan. 2001.
- [8] L. Smith and P. Sørensen, "Pathological outcomes of observational learning," *Econometrica*, vol. 68, no. 2, pp. 371–398, Mar. 2000.
- [9] B. Çelen and S. Kariv, "Observational learning under imperfect information," *Games Econ. Behav.*, vol. 47, no. 1, pp. 72–86, Apr. 2004.
- [10] D. Gale and S. Kariv, "Bayesian learning in social networks," *Games Econ. Behav.*, vol. 45, no. 2, pp. 329–346, Nov. 2003.
- [11] 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.
- [12] J. B. Rhim, L. R. Varshney, and V. K. Goyal, "Quantization of prior probabilities for collaborative distributed hypothesis testing," *IEEE Trans. Signal Process.*, vol. 60, no. 9, pp. 4537–4550, Sep. 2012.

- [13] J. B. Rhim and V. K. Goyal, "Distributed hypothesis testing with social learning and symmetric fusion," *IEEE Trans. Signal Process.*, vol. 62, no. 23, pp. 6298–6308, Dec. 2014.
- [14] V. V. Veeravalli, T. Başar, and H. V. Poor, "Decentralized sequential detection with a fusion center performing the sequential test," *IEEE Trans. Inf. Theory*, vol. 39, no. 2, pp. 433–442, Mar. 1993.
- [15] 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.
- [16] T. Berger, Z. Zhang, and H. Viswanathan, "The CEO problem," *IEEE Trans. Inf. Theory*, vol. 42, no. 3, pp. 887–902, May 1996.
- [17] 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.
- [18] S. Kar and J. M. F. Moura, "Distributed consensus algorithms in sensor networks with imperfect communication: Link failures and channel noise," *IEEE Trans. Signal Process.*, vol. 57, no. 1, pp. 355–369, Jan. 2009.
- [19] S. Kar and J. M. F. Moura, "Distributed consensus algorithms in sensor networks: Quantized data and random link failures," *IEEE Trans. Signal Process.*, vol. 58, no. 3, pp. 1383–1400, Mar. 2010.
- [20] M. E. Hellman and T. M. Cover, "Learning with finite memory," *Ann. Math. Statist.*, vol. 41, no. 3, pp. 765–782, 1970.
- [21] 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, Sep./Oct. 1991.
- [22] J. D. Papastavrou and M. Athans, "Distributed detection by a large team of sensors in tandem," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 28, no. 3, pp. 639–653, Jul. 1992.
- [23] 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.
- [24] M. Alanyali, S. Venkatesh, O. Savas, and S. Aeron, "Distributed Bayesian hypothesis testing in sensor networks," in *Proc. Amer. Control Conf.*, Jun./Jul. 2004, vol. 6, pp. 5369–5374.
- [25] K. R. Rad and A. Tahbaz-Salehi, "Distributed parameter estimation in networks," in *Proc. 49th IEEE Conf. Decis. Control*, Dec. 2010, pp. 5050–5055.
- [26] 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.
- [27] 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.
- [28] A. Nedić, A. Olshevsky, and C. A. Uribe, "Nonasymptotic convergence rates for cooperative learning over time-varying directed graphs," in *Proc. Amer. Control Conf.*, Jul. 2015, pp. 5884–5889.
- [29] 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.
- [30] 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.
- [31] P. K. Varshney, *Distributed Detection and Data Fusion*. New York, NY, USA: Springer-Verlag, 1997.
- [32] R. K. Raman and S. Pattabiraman, "Selfish learning: Leveraging the greed in social learning," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, Apr. 2018, pp. 2261–2265.
- [33] R. Radner, "Team decision problems," *Ann. Math. Statist.*, vol. 33, no. 3, pp. 857–881, Sep. 1962.
- [34] A. McAfee and E. Brynjolfsson, *Machine, Platform, Crowd: Harnessing Our Digital Future*. New York, NY, USA: Norton, 2017.
- [35] A. Tversky and D. Kahneman, "Advances in prospect theory: Cumulative representation of uncertainty," *J. Risk Uncertainty*, vol. 5, no. 4, pp. 297–323, Oct. 1992.
- [36] V. S. S. Nadendla, S. Brahma, and P. K. Varshney, "Towards the design of prospect-theory based human decision rules for hypothesis testing," in *Proc. 54th Annu. Allerton Conf. Commun. Control Comput.*, Sep. 2016, pp. 766–773.
- [37] V. S. S. Nadendla, E. Akyol, C. Langbort, and T. Başar, "Strategic communication between prospect theoretic agents over a Gaussian test channel," in *Proc. Mil. Commun. Conf.*, Oct. 2017, pp. 109–114.
- [38] D. Prelec, "The probability weighting function," *Econometrica*, vol. 66, no. 3, pp. 497–527, May 1998.
- [39] A. al-Nowaihi and S. Dhami, "Probability weighting functions," in *Wiley Encyclopedia of Operations Research and Management Science*, J. J. Cochran, L. A. Cox, Jr., P. Keskinocak, J. P. Kharoufeh, J. C. Smith, Eds. New York, NY, USA: Wiley, Feb. 2011.
- [40] J. P. Mills, "Table of the ratio: Area to bounding ordinate, for any portion of normal curve," *Biometrika*, vol. 18, no. 3/4, pp. 395–400, Nov. 1926.
- [41] M. R. Sampford, "Some inequalities on Mill's ratio and related functions," *Ann. Math. Statist.*, vol. 24, no. 1, pp. 130–132, Mar. 1953.
- [42] J. A. Swets, W. P. Tanner, Jr., and T. G. Birdsall, "Decision processes in perception," *Psychological Rev.*, vol. 68, no. 5, pp. 301–340, Sep. 1961.
- [43] W. K. Viscusi, "Are individuals Bayesian decision makers?" *Amer. Econ. Rev.*, vol. 75, no. 2, pp. 381–385, May 1985.
- [44] G. L. Brase, L. Cosmides, and J. Tooby, "Individuation, counting, and statistical inference: The role of frequency and whole-object representations in judgment under uncertainty," *J. Exp. Psychol., Gen.*, vol. 127, no. 1, pp. 3–21, Mar. 1998.
- [45] M. Glanzer, A. Hilford, and L. T. Maloney, "Likelihood ratio decisions in memory: Three implied regularities," *Psychonomic Bull. Rev.*, vol. 16, no. 3, pp. 431–455, Jun. 2009.
- [46] *Perception as Bayesian Inference*. D. C. Knill and W. Richards, Eds. Cambridge, U.K.: Cambridge Univ. Press, 1996.
- [47] W. A. Yost, A. N. Popper, and R. R. Fay, *Human Psychophysics*. New York, NY, USA: Springer-Verlag, 1993.
- [48] M. H. Jarrahi, "Artificial intelligence and the future of work: Human-AI symbiosis in organizational decision making," *Bus. Horiz.*, vol. 61, no. 4, pp. 577–586, 2018.
- [49] MIT IDE, "Where humans meet machines: Intuition, expertise and learning," May 2018. [Online]. Available: <https://medium.com/mit-initiative-on-the-digital-economy/where-humans-meet-machines-intuition-expertise-and-learning-be639f00bade>
- [50] H. L. Van Trees, *Detection, Estimation, and Modulation Theory*. Hoboken, NJ, USA: Wiley, 1968.
- [51] H. V. Poor, *An Introduction to Signal Detection and Estimation*. Berlin, Germany: Springer, 1988.