

Social Teaching: Being Informative vs. Being Right in Sequential Decision Making

Joong Bum Rhim

Research Laboratory of Electronics
Massachusetts Institute of Technology
Email: jbrhim@mit.edu

Vivek K Goyal

Research Laboratory of Electronics
Massachusetts Institute of Technology
Email: v.goyal@ieee.org

Abstract—We consider sequential Bayesian binary hypothesis testing where each individual agent makes a binary decision motivated only by minimization of her own perception of the Bayes risk. The information available to each agent is an initial belief, a private signal, and decisions of all earlier-acting agents; it follows that each agent should apply a standard Bayesian update of her belief as in social learning. The effect of the set of initial beliefs on the decision-making performance of the last agent is studied. In general, the optimal initial beliefs are not equal to the actual prior probability. When the private signals are described by Gaussian likelihoods, they also are not haphazard, but rather follow a systematic pattern: The earlier-acting agents should act as if the prior probability is larger than it is in reality when the true prior probability is small, and vice versa. We interpret this as being *open-minded* toward the unlikely hypothesis. Such *open-mindedness* increases but does not maximize the mutual information between the true hypothesis and a decision.

I. INTRODUCTION

Consider decision-making agents facing the task of choosing between two alternatives. Each agent has a private signal, which is not visible to the other agents. The agents sequentially make their individual decisions in some predetermined order. Every agent is selfish in the sense that she does whatever she thinks right, which we formalize as the minimization of her perception of expected Bayes cost. This is the standard behavior of idealized economic agents.

An agent should naturally use all available information, and we assume that decisions are made public. Each decision contains some information about the right (or better) choice, so the following agents can learn from the action and reflect it in their own actions. For example, when you want to choose between two alternatives when buying a new phone, the choices made by your colleagues can affect your judgment. This is also a standard principle in all literature on sequential decision making or social learning.

The novelty of this paper is to not assume that the agents know true prior probabilities. This assumption is motivated, for example, by jurors who judge a defendant's prior probability of committing a crime by how intelligent, how attractive, how friendly, and how threatening the defendant looks [1]. Thus, even if they are rationally applying Bayesian reasoning—as humans have been argued to do [2]—they are doing so based on *initial beliefs* rather than prior probabilities. We study the

effect of the perceived, inaccurate beliefs on the decisions of the later-acting agents.

Being influenced by earlier-acting agents has been termed *social learning* [3]. It has generally been studied in settings where each agent has no motivation beyond making a correct choice for herself. In this paper, we study the effect of agents' actions on subsequent agents and find that making correct decisions is generally not equivalent to providing information to other agents that maximally benefits them in their decision-making tasks. Accounting for the effect on other agents could be termed *social teaching*. This term does not necessarily mean that the earlier-acting agents want to be beneficial to the later-acting agents. Also, in any scenario with social learning, the earlier-acting agents can be seen as *advisers* to the later-acting agents. As will be detailed later, we find that a good adviser should be *open-minded* in the sense of being more receptive to the *a priori* less likely alternative than she would have been if she were only interested in being right rather than also interested in being informative. We interpret the phenomenon of open-mindedness as increasing mutual information between advisers' private signals and their decisions.

The framework of sequential decision making with social learning was independently introduced in [4] and [5]. These works focused primarily on *herding*, which is for all agents beyond some index to take the same action. Subsequently, [6] showed that learning is incomplete—beliefs are not eventually focused on the true state—if private signals are boundedly informative, but agents will asymptotically settle on the optimal action otherwise. Recently, [7] extends the result to general network topologies where each agent can observe decisions made by its neighbors instead of all previous agents.

In another related line of work, [8] studies the effect of social learning in a quickest detection problem, in which agents keep updating their beliefs based on previous decisions and detect the time at which an underlying state changes. It has a similar framework to [9], which studied update of private information in a finite memory.

This paper differentiates itself from the literature in that it considers unbounded private signals and does not focus on herding behavior. In addition, we focus largely on the effect of prior probabilities in decision making. We do not assume that an agent knows a correct prior probability for the decision at hand.




Fig. 1. A sequential decision making model with N agents (Alexis, Britta, ..., Norah). Agent n can observe $n - 1$ decisions made by the preceding agents.

Bayes-optimal agents need to know the prior probability in order to perform the likelihood ratio test [10]. Hence, intuitively, agents with biased beliefs should degrade the decision making and yield higher Bayes risk. Contrary to intuition, however, it turns out that biased beliefs may improve the decision of the final agent in sequential decision making.

Section II describes our sequential decision making model. Section III explains how agents interpret precedents, update their beliefs, and make decisions. It is proven in Section IV that the true prior probability is not the optimal initial belief for $N = 2$. Section V considers variations of our decision-making model and points out a possible relation to innate human behaviors. Section VI concludes the paper.

II. PROBLEM DESCRIPTION

Consider the sequential decision making model depicted in Fig. 1. There is an object in a binary state $H \in \{0, 1\}$ with probabilities $\mathbb{P}(\{H = 0\}) = p_0$ and $\mathbb{P}(\{H = 1\}) = 1 - p_0$. There are N agents that sequentially detect the state in a predetermined order. Agents do not know the true prior probability p_0 . Instead, Agent n perceives it as q_n . We call this a *belief* in order to distinguish it from the prior probability.

Agent n observes public signals—decisions of preceding agents, $\{\widehat{H}_1, \dots, \widehat{H}_{n-1}\}$ —as well as a private signal—a noisy observation about H given by $Y_n = H + W_n$. The noise variables $\{W_n\}_{n=1}^N$ are independent and identically distributed (iid) Gaussian random variables, also independent of H .

Our performance analysis focuses on the last agent (Norah) and her decision \widehat{H}_N . Upon observing her private signal Y_N and the $N - 1$ precedent decisions $\widehat{H}_1, \dots, \widehat{H}_{N-1}$, she determines her decision rule. The relative importance of correct decisions and errors can be abstracted as a cost function $c(\widehat{H}, H)$. For simplicity, we assume the correct decisions yield zero cost and use the shorthand notations $c_{10} = c(1, 0)$ for false alarm or Type I error (choosing $\widehat{H} = 1$ when $H = 0$) and $c_{01} = c(0, 1)$ for missed detection or Type II error (choosing $\widehat{H} = 0$ when $H = 1$). In addition, we assume that agents have the same costs; they are a team in the sense of Radner [11]. Then the Bayes risk is given by

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

The computation of (1) depends on the previous decisions $\widehat{H}_1, \dots, \widehat{H}_{N-1}$. Therefore, the computation of the expected

cost can be expanded as

$$R_N = \sum_{\widehat{h}_1, \dots, \widehat{h}_{N-1}} \left(c_{10}p_0p_{\widehat{H}_N, \widehat{H}_{N-1}, \dots, \widehat{H}_1|H}(1, \widehat{h}_{N-1}, \dots, \widehat{h}_1|0) + c_{01}(1 - p_0)p_{\widehat{H}_N, \widehat{H}_{N-1}, \dots, \widehat{H}_1|H}(0, \widehat{h}_{N-1}, \dots, \widehat{h}_1|1) \right). \quad (2)$$

In our model, the action of Agent n cannot be specifically to minimize the Bayes risk, because the agent does not know p_0 . Instead, the agent minimizes her *perceived Bayes risk*, which is the Bayes risk with prior probability p_0 replaced by initial belief q_n . The decisions $\{\widehat{H}_1, \dots, \widehat{H}_{n-1}\}$ of the earlier-acting agents reveal information about H and thus should be incorporated by Agent n . Since she does not know the initial beliefs $\{q_1, \dots, q_{n-1}\}$ of the earlier-acting agents, she aggregates information under the assumption of $q_1 = q_2 = \dots = q_n$.

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 .

III. SOCIAL LEARNING

Public signals convey some information and enrich agents' observations combining with their private signals. We discuss the utilization of the public signals as well as the private signals for Bayesian hypothesis testing. To be clear, the method derived in this section is the optimal decision-making strategy whether the agents know true prior probabilities or their initial beliefs are biased (inaccurate).

Agent n observes a private signal Y_n and public signals $\widehat{H}_1, \dots, \widehat{H}_{n-1}$. She, who believes the prior probability is q_n , uses the following likelihood ratio test to minimize her perceived Bayes risk:

$$\frac{f_{Y_n, \widehat{H}_1, \dots, \widehat{H}_{n-1}|H}(y_n, \widehat{h}_1, \dots, \widehat{h}_{n-1}|1)}{f_{Y_n, \widehat{H}_1, \dots, \widehat{H}_{n-1}|H}(y_n, \widehat{h}_1, \dots, \widehat{h}_{n-1}|0)} \stackrel{\widehat{H}_n(y_n)=1}{\underset{\widehat{H}_n(y_n)=0}{\gtrless}} \frac{c_{10}q_n}{c_{01}(1 - q_n)}. \quad (3)$$

Since her private signal is independent of the public signals conditioned on H , the likelihood ratio test (3) can be written in terms of the likelihood function:

$$\frac{f_{Y_n|H}(y_n|1)}{f_{Y_n|H}(y_n|0)} \stackrel{\widehat{H}_n(y_n)=1}{\underset{\widehat{H}_n(y_n)=0}{\gtrless}} \frac{c_{10}q_n}{c_{01}(1 - q_n)} \times \frac{p_{\widehat{H}_1, \dots, \widehat{H}_{n-1}|H}(\widehat{h}_1, \dots, \widehat{h}_{n-1}|0)_n}{p_{\widehat{H}_1, \dots, \widehat{H}_{n-1}|H}(\widehat{h}_1, \dots, \widehat{h}_{n-1}|1)_n}, \quad (4)$$

where the subscript “ n ” in the last terms indicates that the probabilities of the public signals are computed based on q_n even though, in fact, they are chosen based on q_1, \dots, q_{n-1} . Because Agent n does not know beliefs of other agents, she treats $\widehat{H}_1, \dots, \widehat{H}_{n-1}$ as if the preceding agents have the same beliefs as hers.

An agent whose belief q' satisfies

$$\frac{q'}{1-q'} = \frac{q_n}{1-q_n} \frac{p_{\widehat{H}_1, \dots, \widehat{H}_{n-1}|H}(\widehat{h}_1, \dots, \widehat{h}_{n-1}|0)_n}{p_{\widehat{H}_1, \dots, \widehat{H}_{n-1}|H}(\widehat{h}_1, \dots, \widehat{h}_{n-1}|1)_n} \quad (5)$$

would perform the same likelihood ratio test in a single-agent binary hypothesis testing problem. Thus, the public signals play a role of updating the belief of Agent n from q_n to q' :

$$q' = p_{H|\widehat{H}_1, \dots, \widehat{H}_{n-1}}(0|\widehat{h}_1, \dots, \widehat{h}_{n-1})_n. \quad (6)$$

Agent n need not update her belief at one time. The belief update (6) consists of a cascade of fragmentary updates considering one public signal at each time. For example, after Agent 1 (Alexis) makes a decision, Agent n observes it and updates her belief to q_n^A :

$$\frac{q_n^A}{1-q_n^A} = \frac{q_n}{1-q_n} \frac{p_{\widehat{H}_1|H}(\widehat{h}_1|0)_n}{p_{\widehat{H}_1|H}(\widehat{h}_1|1)_n}. \quad (7)$$

Similarly, after Agent 2 (Britta) makes a decision, Agent n updates her belief from q_n^A to q_n^{AB} :

$$\frac{q_n^{AB}}{1-q_n^{AB}} = \frac{q_n^A}{1-q_n^A} \frac{p_{\widehat{H}_2|\widehat{H}_1,H}(\widehat{h}_2|\widehat{h}_1,0)_n}{p_{\widehat{H}_2|\widehat{H}_1,H}(\widehat{h}_2|\widehat{h}_1,1)_n}. \quad (8)$$

Agent n updates her belief whenever an agent makes a decision and will have q' in (6) after Agent $n-1$ makes her decision publicly known. The update function is described in a recurrence form and explained in detail in [12].

IV. AN OPEN-MINDED ADVISER

The agents we are considering have personal beliefs, which may be different from the true prior probability, p_0 . It is worth asking what beliefs are beneficial to the last agent. This discussion is not an optimization of the beliefs of the agents, and it does not change their intentions. It is a demonstration that biased beliefs of agents can happen to be beneficial. We discuss a two-agent case (Alexis and Britta) for simplicity.

The criterion is Britta's Bayes risk, R_2 , which is determined by Alexis's and Britta's error probabilities:

$$\begin{aligned} R_2 &= c_{10}p_{\widehat{H}_2,H}(1,0) + c_{01}p_{\widehat{H}_2,H}(0,1) \\ &= c_{10} [P_{e,2}^{I_0} (1 - P_{e,1}^I) + P_{e,2}^{I_1} P_{e,1}^I] p_0 \\ &\quad + c_{01} [P_{e,2}^{II_0} P_{e,1}^{II} + P_{e,2}^{II_1} (1 - P_{e,1}^{II})] (1 - p_0), \end{aligned} \quad (9)$$

where $P_{e,1}^I$ and $P_{e,1}^{II}$ denote probabilities of Alexis's Types I and II errors and $P_{e,2}^{I_h}$ and $P_{e,2}^{II_h}$ denote probabilities of Britta's Types I and II errors with h in superscripts specifying Alexis's decision $\widehat{H}_1 = h$. The Bayes risk (9) is a function of Alexis's and Britta's initial beliefs, q_1 and q_2 .

One might guess that $q_1 = q_2 = p_0$ minimizes (9). In this case, Alexis will make the best decision she can. Also, Britta will not misunderstand her decision because they have the same beliefs. Surprisingly, however, it turns out that it is not optimal. It is simple to verify that Alexis's optimal belief, q_1^* , is not equal to p_0 .

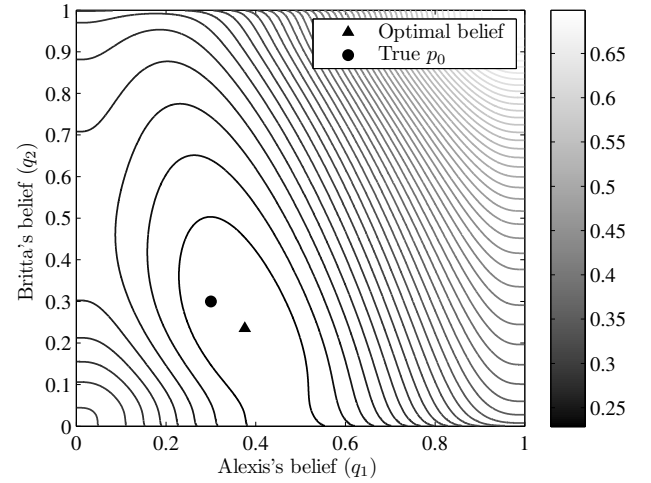


Fig. 2. Visualization of Britta's Bayes risk for various q_1 and q_2 for $N = 2$, $c_{10} = c_{01} = 1$, $p_0 = 0.3$, and additive Gaussian noise with zero mean and unit variance. Alexis's and Britta's optimal initial beliefs (▲) are different from the true prior probability (●).

Let us consider the first derivative of (9) with respect to q_1 :

$$\frac{dR_2}{dq_1} = c_{10}p_0(P_{e,2}^{I_1} - P_{e,2}^{I_0}) \frac{dP_{e,1}^I}{dq_1} + c_{01}(1-p_0)(P_{e,2}^{II_0} - P_{e,2}^{II_1}) \frac{dP_{e,1}^{II}}{dq_1}.$$

Let us denote Alexis's decision threshold by λ_1 : Alexis chooses $\widehat{H}_1 = 0$ if $y_1 < \lambda_1$ and $\widehat{H}_1 = 1$ otherwise. Using

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

the first derivative equals to zero when

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

Note that λ_1 is obtained based on Alexis's belief and, thus, a solution to

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

Therefore Alexis's optimal belief q_1^* needs to satisfy

$$\frac{q_1^*}{1-q_1^*} = \frac{p_0(P_{e,2}^{I_1} - P_{e,2}^{I_0})}{(1-p_0)(P_{e,2}^{II_0} - P_{e,2}^{II_1})}, \quad (11)$$

where the value of $(P_{e,2}^{I_1} - P_{e,2}^{I_0}) / (P_{e,2}^{II_0} - P_{e,2}^{II_1})$ does not have to be 1. In fact, in additive Gaussian noise cases, it is not equal to one except for $p_0 = c_{01} / (c_{10} + c_{01})$. Therefore, the optimal value of q_1 is not p_0 in general.

For example, Fig. 2 depicts Britta's Bayes risk for various q_1 and q_2 when private signals, Y_1 and Y_2 , are distorted by iid additive Gaussian noise with zero mean and unit variance. Even though the prior probability is 0.3, Alexis's and Britta's optimal beliefs are different from 0.3. Alexis's optimal belief is larger than 0.3 and Britta's belief is smaller.

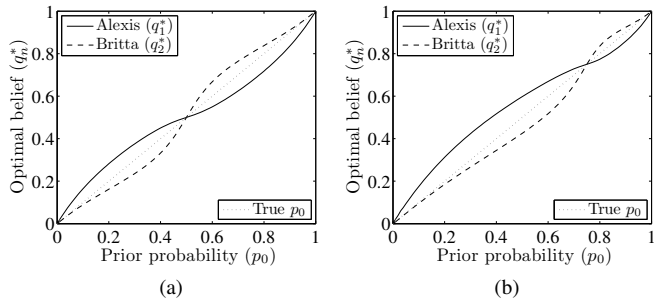


Fig. 3. The trend of the optimal initial beliefs for varying p_0 for $N = 2$ (Alexis and Britta). (a) $c_{10} = c_{01} = 1$. (b) $c_{10} = 1, c_{01} = 3$.

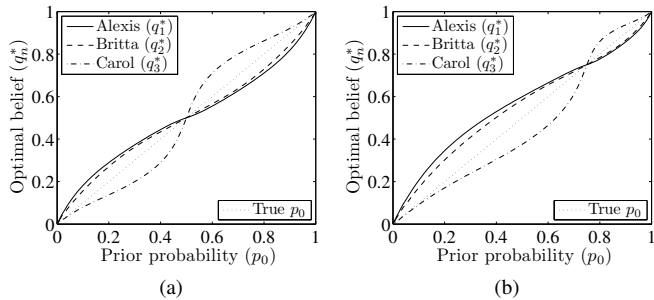


Fig. 4. The trend of the optimal initial beliefs for varying p_0 for $N = 3$ (Alexis, Britta, and Carol). (a) $c_{10} = c_{01} = 1$. (b) $c_{10} = 1, c_{01} = 3$.

In fact, these optimal beliefs follow systematic patterns as shown in Figs. 3 and 4. First, the non-terminal agents (Alexis for $N = 2$ and Alexis and Britta for $N = 3$) should have beliefs larger than p_0 when p_0 is small and beliefs smaller than p_0 when p_0 is large. We call this *open-mindedness* because it is to assign a higher initial belief to outcomes that are very unlikely. Second, the last agent (Britta for $N = 2$ and Carol for $N = 3$) should have a belief smaller than p_0 when p_0 is small and a belief larger than p_0 when p_0 is large. This is necessary to compensate for the biases of preceding agents. Last, there is a unique fixed point, except for $p_0 = 0$ or $p_0 = 1$, where all agents' optimal beliefs are the same as the true prior probability. It occurs at $p_0 = c_{01}/(c_{10} + c_{01})$. The first and the last patterns are proven for $N = 2$ in [12].

V. VARIATIONS AND INTERPRETATIONS

A. A Team of Agents Knowing Each Other's Beliefs

Agents in our decision-making model do not know the others' beliefs. However, even if we relax this constraint, the optimality of an open-minded adviser does not disappear. Let us consider a team of agents such that, for any n , Agent n knows beliefs of all preceding agents, or more specifically, their decision thresholds. Then Agent n interprets public signals $\widehat{H}_1, \dots, \widehat{H}_{n-1}$ correctly and her updated belief will depend on q_1, \dots, q_{n-1} as well as on q_n . We can drop the subscript " n " in the equations in Section III.

The fact that Alexis's optimal belief for Britta is not the same as the true prior probability can be proven similarly to the process in Section IV. In addition, the optimal beliefs of




Fig. 5. The trend of the optimal initial beliefs for varying p_0 for $c_{10} = c_{01} = 1$ when agents know what decision thresholds others use. (a) $N = 2$ (Alexis and Britta) (b) $N = 3$ (Alexis, Britta, and Carol).




Fig. 6. Mutual information $I(H_1; \widehat{H}_1)$ for Alexis's various choices of belief for $N = 2$: From top to bottom, maximum $I(H_1; \widehat{H}_1)$ (—), her optimal belief for Britta (---), and the belief equal to prior probability (⋯).

earlier-acting agents show the same kind of pattern, as shown in Fig. 5. A substantial change is the optimal belief of the last agent. In this model, it is optimal for the last agent to know p_0 because the interpretation of the public signals is already accurate and the last agent need not compensate for the biases of preceding agents.

This observation implies that the open-mindedness comes not from ignorance, and it supports our idea that it comes from transfer of information, which we will discuss in the next subsection.

B. The Way to Be Informative

Alexis's open-minded belief increases the mutual information $I(H; \widehat{H}_1)$. Fig. 6 compares the mutual information in the cases when her belief is equal to q_1^* in Fig. 3a and when it is equal to p_0 . For comparison, the maximum mutual information is depicted as well. Her optimal belief for Britta yields not maximum but higher mutual information than using the true prior p_0 .

All agents we have considered until now are selfish. However, we can think of earlier-acting agents that know p_0 but are selfless. Such agents will want to balance between being right for themselves and being informative for the later-acting agents. Therefore, we can come up with a family of utilities that linearly combine the Bayes risk and the mutual information:

$$\max_{q_n} -R_n + \alpha_n I(H; \widehat{H}_n), \quad (12)$$




Fig. 7. Comparison of Alexis's belief that maximizes her utility (12) and her belief that minimizes Britta's or Carol's Bayes risk for $c_{10} = c_{01} = 1$.

where R_n denotes the Bayes risk following the definition of (1) and α_n is a nonnegative parameter that represents the extent of selflessness of Agent n .

If Agent n modifies her belief to maximize her utility (12), then her behavior will be open-minded. For example, when Alexis has $\alpha_1 = 1$, her belief maximizing (12) is very close to her optimal belief for Britta or Carol. They are depicted in Fig. 7.

C. Connection to Natural Human Behavior

In previous sections, we have analyzed binary decision-making performance of agents with mathematically-idealized rational behavior, but a type of bounded rationality in which they use their initial beliefs in place of the prior probability p_0 . In this section, we describe a similarity between our prediction of the behavior of a good adviser and well-established economic models that have been designed to empirically match human behavior.

In behavioral economics, prospect theory [13] and cumulative prospect theory [14] have been introduced to describe observed systematic deviations from utility maximization and suboptimal behaviors of human decision makers. In addition to framing effects and a nonlinear value curve, experiments in cumulative prospect theory reveal that people use a nonlinear probability scale, which overweights small probabilities and underweights moderate and high probabilities.

The probability weighting function is evidence of human suboptimality. Upon observing the resemblance between the characteristic of a good adviser and the innate probability weighting function, however, we do not think that this behavior is just suboptimal in an arbitrary way. This resemblance may imply that humans sacrifice optimality to have the quality of a good adviser.

VI. CONCLUSION

We have discussed decision making sequentially performed in a predetermined order by a team of agents that make decisions based on individual initial beliefs that may be biased. Instead of investigating herding on a wrong action, we have assumed unbounded private signals and focused on the belief update of the agents. The Bayes-optimal updated belief turns out to be the probability of each hypothesis conditioned on the decisions made by previous agents.

The possibly-wrong beliefs held by preceding agents affect the probabilities that following agents choose each hypothesis. Contrary to intuition, however, wrong beliefs are not always bad. In fact, the optimal beliefs of agents (those that lead to the minimum Bayes risk for the last agent) are not usually equal to the true prior probability. Specifically, in the case of observations corrupted by iid additive Gaussian noises, an agent biased toward $c_{01}/(c_{10} + c_{01})$ can be more beneficial to subsequent agents than is an agent with accurate belief.

The idea of an open-minded adviser is related to the amount of information conveyed in the public signals. A public signal is a quantized version of a private signal while also simultaneously reflecting an agent's belief. Britta would want a public signal that is more informative of Alexis's private signal and thus want her to be open-minded. We have seen that Alexis's open-minded belief increases the mutual information between the true hypothesis and the public signal. The open-mindedness is innate human behavior as revealed by economic studies.

While some conclusions of our study depend on having Gaussian likelihoods and may not hold for different types of additive noise, it is more generally true that the optimal initial beliefs are different from the true prior probability.

ACKNOWLEDGMENT

This material is based upon work supported by the National Science Foundation under Grant No. 1101147.

REFERENCES

- [1] M. J. Brown, E. Henriquez, and J. Groscup, "The effects of eyeglasses and race on juror decisions involving a violent crime," *American Journal of Forensic Psychology*, vol. 26, no. 2, pp. 25–43, Sep. 2008.
- [2] W. K. Viscusi, "Are individuals Bayesian decision makers?" *American Economics Review*, vol. 75, no. 2, pp. 381–385, May 1985.
- [3] G. Ellison and D. Fudenberg, "Rules of thumb for social learning," *Journal of Political Economy*, vol. 101, no. 4, pp. 612–643, Aug. 1993.
- [4] A. V. Banerjee, "A simple model of herd behavior," *The Quarterly Journal of Economics*, vol. 107, no. 3, pp. 797–817, Aug. 1992.
- [5] S. Bikhchandani, D. Hirshleifer, and I. Welch, "A theory of fads, fashion, custom, and cultural change as informational cascades," *Journal of Political Economy*, vol. 100, no. 5, pp. 992–1026, Oct. 1992.
- [6] L. Smith and P. Sørensen, "Pathological outcomes of observational learning," *Econometrica*, vol. 68, no. 2, pp. 371–398, Mar. 2000.
- [7] D. Acemoglu, M. A. Dahleh, I. Lobel, and A. Ozdaglar, "Bayesian learning in social networks," *Review of Economic Studies*, vol. 78, no. 4, pp. 1201–1236, Mar. 2011.
- [8] V. Krishnamurthy, "Quickest detection POMDPs with social learning: Interaction of local and global decision makers," *IEEE Transactions on Information Theory*, vol. 58, no. 8, pp. 5563–5587, Aug. 2012.
- [9] M. E. Hellman and T. M. Cover, "Learning with finite memory," *The Annals of Mathematical Statistics*, vol. 41, no. 3, pp. 765–782, Jun. 1970.
- [10] J. Neyman and E. S. Pearson, "On the problem of the most efficient tests of statistical hypotheses," *Philosophical Transactions of the Royal Society of London. Series A*, vol. 231, pp. 289–337, Jan. 1933.
- [11] R. Radner, "Team decision problems," *Ann. Math. Stat.*, vol. 33, no. 3, pp. 857–881, Sep. 1962.
- [12] J. B. Rhim and V. K. Goyal, "Social teaching: Being informative vs. being right in sequential decision making," arXiv:1212.6592v1 [cs.IT], Dec. 2012.
- [13] D. Kahneman and A. Tversky, "Prospect theory: An analysis of decision under risk," *Econometrica*, vol. 47, no. 2, pp. 263–291, Mar. 1979.
- [14] A. Tversky and D. Kahneman, "Advances in prospect theory: Cumulative representation of uncertainty," *Journal of Risk and Uncertainty*, vol. 5, no. 4, pp. 297–323, Oct. 1992.