PROBABILITY REWEIGHTING IN SOCIAL LEARNING: OPTIMALITY AND SUBOPTIMALITY

Daewon Seo, Ravi Kiran Raman, Lav R. Varshney

University of Illinois at Urbana-Champaign

ABSTRACT

This work explores sequential Bayesian binary hypothesis testing in the social learning setup under expertise diversity. We consider a two-agent (say *advisor-learner*) sequential binary hypothesis test where the learner infers the hypothesis based on the decision of the advisor, a prior private signal, and individual belief. In addition, the agents have varying expertise, in terms of the noise variance in the private signal.

Under such a setting, we first investigate the behavior of optimal agent beliefs and observe that the nature of optimal agents could be inverted depending on expertise levels. We also discuss suboptimality of the Prelec reweighting function under diverse expertise. Next, we consider an advisor selection problem wherein the belief of the learner is fixed and the advisor is to be chosen for a given prior. We characterize the decision region for choosing such an advisor and argue that a learner with beliefs varying from the true prior often ends up selecting a suboptimal advisor.

Index Terms— social learning, sequential binary hypothesis test, cumulative prospect theory

1. INTRODUCTION

Team decision making typically involves individual decisions that are influenced by the 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 [1–3]. Individual agents are selfish and aim to minimize their perceived Bayes risk, according to beliefs as reinforced by earlier decisions. In particular, a team of two agents can be treated as an *advisor* followed by a *learner*.

Social learning, also referred to as observational learning, has been widely studied and we provide a non-exhaustive listing of some of the relevant works. Aspects of conformism and "herding" were studied in [4–6]. The concept of herding is further highlighted to be a consequence of boundedly informative private signals in [7]. Further convergence properties of actions taken under social learning were explored under imperfect information in [8].

Rhim and Goyal [3] studied a sequential binary hypothesis test in the social learning framework, termed *social teaching*, and characterized optimal beliefs of agents that minimize the Bayes risk of the last-acting agent. In two and three agent contexts, they showed, counterintuitively, that it is optimal for agents to use beliefs that do not match the true prior. Specifically, the optimal advisor in the social learning context is one who is *open-minded*, i.e., overweights the belief for small prior, and underweights when it is large. On the other hand, the corresponding optimal learner is one who is *closed-minded* and behaves in the opposite way to the advisor.

Human actions are typically affected by individual perceptions of the underlying context. Cumulative prospect theory [9] seeks to provide a psychological understanding of human behaviors under risk. It introduces the notion of probability reweighting functions to explain irrational human behaviors. Among reweighting functions, the Prelec reweighting function [10] satisfies a majority of the axiomatic behavior of the prospect theory. Interestingly, the Prelec function spans a class of open- and closed-minded beliefs and hence one might expect it to emerge as the information-theoretically optimal choice under social learning. However, we will discuss that it does not capture all behavioral patterns for the optimal beliefs of agents.

In particular, we consider observation models with different noise variances, which translates to varying agent *expertise*. The expertise of the advisor and learner affect the nature of optimal beliefs of the agents. Specifically, when the learner has more expertise than the advisor, the Prelec function does not capture the behavior of optimal beliefs. We also identify interesting properties of the optimal beliefs.

We are ultimately interested in the Bayes risk of the learner, and thus it is important that the learner uses the correct set of advisors for the task. To this end, we also consider team selection for such sequential hypothesis testing, and characterize the criterion for advisor selection.

2. PROBLEM DESCRIPTION

Consider a two-agent sequential decision making problem. The underlying hypothesis, $H \in \{0, 1\}$, is a binary signal with prior $\mathbb{P}[H = 0] = p_0$, and $\mathbb{P}[H = 1] = 1 - p_0$. Quantities of the first agent (advisor) are denoted by subscript 1, and

those of the second agent (learner) by subscript 2. Each agent perceives p_0 differently as $q_n, n = 1, 2$, called belief.

Each agent receives the private signal $Y_n = H + Z_n$, where Z_n is an independent additive Gaussian noise with zero mean and variance σ_n^2. Let $\mathcal{N}(y; \mu, \sigma^2)$ be the Gaussian probability density with mean μ and variance σ^2 at y. Then the received signal probability densities for $H = h$ are

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

In addition, the learner acquires the decision of the advisor $\hat{H}_1$, and makes a decision $\hat{H}_2$ based on $(\hat{H}_1, Y_2)$.

The advisor and the learner both are selfish and aim to minimize individual Bayes risk with costs $c(\hat{H}, H)$. We restrict to the case where the cost is 0 for correct decisions and let $c_{01} = c(0, 1), c_{10} = c(1, 0)$. Then, the Bayes risk for the nth agent is given by

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

Thus each agent performs the likelihood ratio test, but the learner assumes that the advisor has the same beliefs as her, as she is unaware of the advisor's belief q_1. Therefore in this social learning scenario, the decision made by the advisor reinforces the decision of the learner by appropriately strengthening the posterior probability of the underlying hypothesis. The advisor follows the likelihood ratio test,

$$\mathcal{L}(y_1) \triangleq \frac{f(y_1|1)}{f(y_1|0)} \underset{\hat{H}_1=0}{\overset{\hat{H}_1=1}{\gtrless}} \frac{c_{10}q_1}{c_{01}(1 - q_1)},$$

and from [3], the learner decides according to

$$\mathcal{L}(y_2) \triangleq \frac{f(y_2|1)}{f(y_2|0)} \underset{\hat{H}_2=0}{\overset{\hat{H}_2=1}{\gtrless}} \frac{c_{10}q_2}{c_{01}(1 - q_2)} \frac{p_{\hat{H}_1|H}(\hat{h}_1|0)_{[2]}}{p_{\hat{H}_1|H}(\hat{h}_1|1)_{[2]}},$$

where the subscript [2] of probabilities indicates probability distribution 'seen by the learner', i.e., the probability computed as if the advisor also has q_2 as belief.

We formally introduce the Prelec reweighting function.

Definition 1 ([10]). *For $\alpha, \beta > 0$, the Prelec reweighting function $w : [0, 1] \mapsto [0, 1]$ is $w(p) = \exp(-\beta(-\log p)^\alpha)$.*

A more generic form, termed composite Prelec weighting function has been defined in [11]. Notice that 1) $w(p)$ is strictly increasing; 2) has a unique fixed point $w(p) = p$ 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 $\alpha > 1$.

3. DIVERSE EXPERTISE LEVELS

Consider the two-agent team with observational noise variances σ_1^2, σ_2^2. Note that smaller noise variance implies the

agent is more likely to infer correctly and so has more *expertise*.

Recall the decision threshold for the Gaussian binary hypothesis test with prior p, and variance σ^2 is given by

$$\lambda(p, \sigma^2) \triangleq \frac{1}{2} + \sigma^2 \log\left(\frac{c_{10}p}{c_{01}(1 - p)}\right).$$

Then, because the advisor thinks the prior is q_1, the decision threshold for the advisor is given by $\lambda_1 = \lambda(q_1, \sigma_1^2)$. But, the learner presumes that the advisor decides according to the threshold $\lambda_{1,[2]} = \lambda(q_2, \sigma_2^2)$.

Let $P_{e,1}^{\mathrm{I}}, P_{e,1}^{\mathrm{II}}$, and $P_{e,1,[2]}^{\mathrm{I}}, P_{e,1,[2]}^{\mathrm{II}}$, be the true Type-I and Type-II error probabilities of the advisor and those as perceived by the learner, respectively. Further, let $P_{e,2}^{\mathrm{I}_h}, P_{e,2}^{\mathrm{II}_h}$ be the error probabilities of the learner upon observing $\hat{H}_1 = h$.

Let the learner's posterior upon observing $\hat{H}_1 = h$ be q_2^h, and let the corresponding decision threshold be $\lambda_2^h = \lambda(q_2^h, \sigma_2^2)$. The posterior probabilities satisfy

$$\frac{q_2^0}{1 - q_2^0} = \frac{q_2}{1 - q_2} \frac{1 - P_{e,1,[2]}^{\mathrm{I}}}{P_{e,1,[2]}^{\mathrm{II}}}, \quad \frac{q_2^1}{1 - q_2^1} = \frac{q_2}{1 - q_2} \frac{P_{e,1,[2]}^{\mathrm{I}}}{1 - P_{e,1,[2]}^{\mathrm{II}}}.$$

The optimal beliefs of advisor and learner q_1^*, q_2^* that minimize R_2, are obtained by solving $\frac{dR_2}{dq_1} = \frac{dR_2}{dq_2} = 0$. From [3], the optimal belief of the advisor satisfies

$$\frac{q_1^*}{1 - q_1^*} = \frac{p_0}{1 - p_0} \frac{P_{e,2}^{\mathrm{I}_1} - P_{e,2}^{\mathrm{I}_0}}{P_{e,2}^{\mathrm{II}_0} - P_{e,2}^{\mathrm{II}_1}}. \quad (2)$$

From (2), we observe some properties of q_1^*, q_2^*.

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

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 $p_0 \in \left\{0, \frac{c_{01}}{c_{01}+c_{10}}, 1\right\}$.*

Thm. 2 highlights the fact that if the learner believes the null hypothesis is more likely, then the ideal advisor underweights the prior, and vice versa. Additionally, for p_0 near zero (near one) the optimal advisor overweights (underweights) the prior. Proof is omitted due to space limitation.

In particular, let us consider two cases separately. First, let the advisor have more expertise, i.e., $\sigma_1^2 < \sigma_2^2$. Then the curves for optimal beliefs and the corresponding Bayes risk are as shown in Fig. 1. The behavior here is similar to the case with equal expertise [3], indicating that the additional expertise of the advisor does not alter the overall behaviors of beliefs, as the learner is unaware of this improved expertise.

On the other hand, when the learner has more expertise, i.e., $\sigma_1^2 > \sigma_2^2$, we notice that the nature of curves changes, as shown in Fig. 2. The behavior of the ideal agents indicates that when the advisor has significantly less expertise than the learner, the learner stays open-minded.
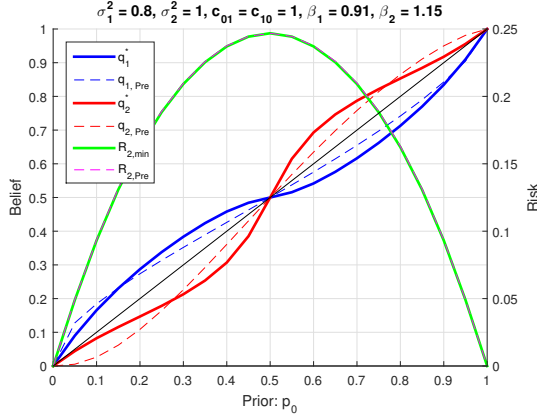
Fig. 1: Optimal beliefs as compared to Prelec-weighted beliefs when the advisor has more expertise.
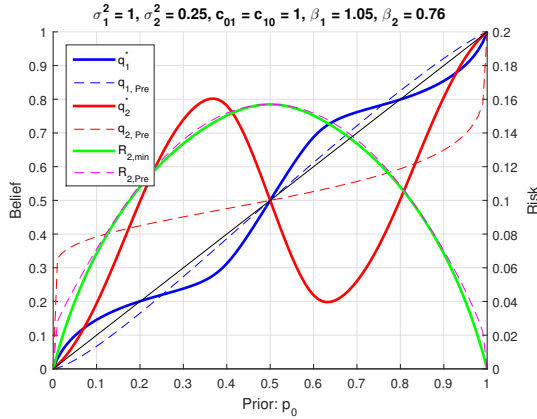
Fig. 2: Optimal beliefs as compared to Prelec-weighted beliefs when the learner has more expertise.

To discuss the Prelec function in social learning, we approximate the optimal belief q_n^* by the Prelec function. We first restrict to the Prelec family whose fixed point is identical with $p^* = \frac{c_{10}}{c_{01}+c_{10}}$ and then find best parameters (α_n, β_n) in the minimax absolute error sense, i.e., pick (α_n, β_n) such that

$$\min_{\alpha, \beta} \max_{p_0} |q_n^*(p_0) - \exp(-\beta(-\log p_0)^\alpha)|,$$

with fixed point at p^*. Such Prelec functions $(q_{1,\text{Pre}}, q_{2,\text{Pre}})$ result in corresponding Bayes risk $R_{2,\text{Pre}}$.

The Prelec quantities are shown in dotted lines in Figs. 1 and 2. When the advisor has more expertise as in Fig. 1, while the fit belief is not perfect, the Bayes risk does not increase by much. Compared to the trivial beliefs $p_0 = q_1 = q_2$ and its maximal loss $\max_{p_0}(R_{2,\text{tri}} - R_{2,\text{min}}) \approx 0.0039$ (not shown in Fig. 1), the best Prelec curves improve the loss ≈ 0.0009. The optimal Prelec-weighted agents also mimic the overall behavior of the ideal agents.

On the other hand, when the learner has more expertise as in Fig. 2, the Prelec agents do not accurately mimic the

optimal 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 loss of Bayes risk by the Prelec fitting is ≈ 0.0187, the loss of trivial reweighting $p_0 = q_1 = q_2$ is ≈ 0.0060. This indicates that even though the Prelec weighting functions serve as good approximations with expert advisors, they do not model the optimal behavior in the case of poor advisors.

In addition, q_1^* has multiple crossings with p_0, i.e., $q_2^* = c_{01}/(c_{01} + c_{10})$. As expected, the ideal advisor is open-minded for near zero and one prior probabilities. However, when the hypotheses have Bayes risk around its peak, the ideal advisor chooses to favor the likely hypothesis. That is, around p^*, the learner stays open-minded as the decisions of the advisor are less accurate. To further understand the nature of such an advisor, we characterize the crossings of the curve with the prior. The complementary cumulative distribution function of the standard Gaussian is denoted by $Q(x)$.

Theorem 3. *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*

$$e^x = \frac{1 - \beta Q(\alpha + x)}{1 - \beta Q(\alpha - x)}, \quad (3)$$

where

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

We note that $p^* = \frac{c_{01}}{c_{01}+c_{10}}$ is always a solution to (3). We are particularly interested in when it has multiple solutions.

Corollary 4. *If*

$$2\beta \frac{\mathcal{N}(\alpha; 0, 1)}{1 - \beta Q(\alpha)} > 1, \quad (4)$$

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

Cor. 4 provides sufficient conditions on the expertise of agents under which there exists multiple crossings of the curves $q_1^*(p_0)$ and p_0. This is important as the crossings indicate a change in the perceived bias by the advisor. We omits proofs for Thm. 3 and Cor. 4 due to space limitation.

4. TEAM CONSTRUCTION CRITERION

Having studied the mathematical conditions for optimal reweighting of prior probabilities, 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 learner to identify the optimal advisors when a set of advisors is given. The proof of the next theorem is omitted due to space limitation.
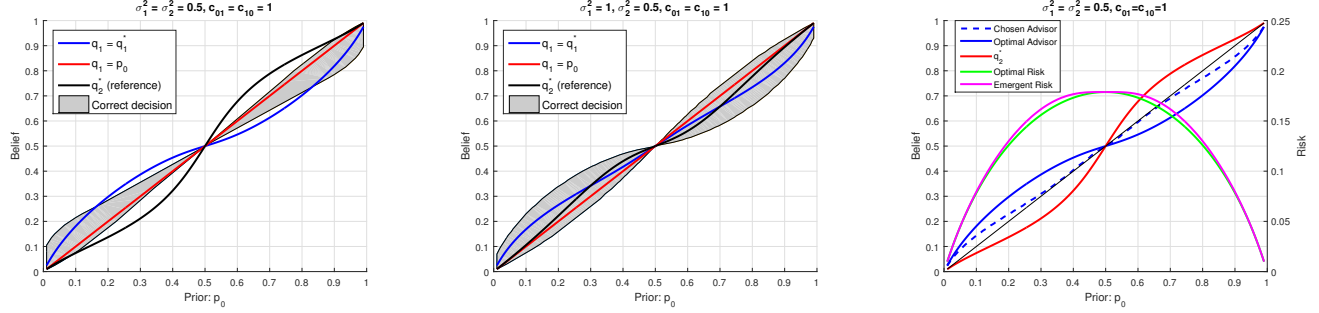
(a) Advisor selection under equal expertise. The optimal learner fails to recognize the optimal advisor and makes mistakes often.

(b) Advisor selection under expertise diversity. The optimal learner selects the optimal advisor in this case.

(c) Context-unaware advisor selection. The learner chooses advisor using (6) without p_0 and the Bayes risk increases as a result.

Fig. 3: Context unaware team selection.

Theorem 5. *Consider two advisors with $q_1 < q_{1'}$. Let $\lambda_1, \lambda_{1'}$ be the decision thresholds of the respective advisors. Then, the advisor with belief q_1 is the optimal choice if and only if*

$$\frac{\mathbb{P}_1\left[Y_1 \in [\lambda_1, \lambda_{1'}], Y_2 \in [\lambda_2^1, \lambda_2^0]\right]}{\mathbb{P}_0\left[Y_1 \in [\lambda_1, \lambda_{1'}], Y_2 \in [\lambda_2^1, \lambda_2^0]\right]} \geq \frac{c_{10}p_0}{c_{01}(1-p_0)}. \quad (5)$$

In other words, by rewriting (5) as a likelihood form,

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

where $\hat{H}_{2'}$ is the decision made by the learner following the decision of the advisor with belief $q_{1'}$.

Thus selecting an ideal advisor requires a social planner who is aware of the context p_0. Without this, the learner selects an advisor according to his personal belief q_2. That is, the learner verifies condition (5) by replacing p_0 by q_2. Such a choice of advisor might not always conform to the optimal choice when the belief of the learner deviates significantly from the prior. To illustrate, we consider the problem of choosing between two advisors with belief $q_1(p_0) = q_1^*(p_0)$ and $q_{1'}(p_0) = p_0$. Let $q(p_0, q_2)$ be the belief of the optimal advisor 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 expertise levels are equal, the region in which the learner picks the correct advisor is shown in Fig. 3a. We note that the correct region is relatively small and does not include q_2^*. In particular, the learner with optimal belief chooses the wrong advisor always, whereas a suboptimal learner with beliefs in the shaded region picks the correct one.

On the other hand, when the learner has more expertise than the advisor, the corresponding region is as shown in Fig. 3b. Here we note that the learner with optimal belief picks the correct advisor always.

Thus, we note that knowledge of the mathematically optimal beliefs does not guarantee selection of the right advisor.

Further, we also observe that the diversity of expertise levels may increase the feasibility of selecting the right advisor when the learner has optimal belief.

We also explore the optimal choice of advisor for the given optimal learner in the absence of knowledge of the prior probability. From (2), the belief of the optimal advisor, $\tilde{q}_1$ chosen by a learner, in the absence of context (prior probability p_0) satisfies

$$\frac{\tilde{q}_1}{1-\tilde{q}_1} = \frac{p_0}{1-p_0} \frac{P_{e,2}^{\mathrm{I}_1} - P_{e,2}^{\mathrm{I}_0}}{P_{e,2}^{\mathrm{II}_0} - P_{e,2}^{\mathrm{II}_1}}. \quad (6)$$

The learner's behavior with belief q_2^* is as shown in Fig. 3c. We note that the advisor chosen by the learner 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.

5. CONCLUSION

We considered the problem of sequential social learning under varying agent expertise and investigated the question of optimal probability reweighting in systems with two agents—advisor and learner.

Under specific levels of expertise, we showed that the Prelec reweighting function approximates the behavior of the optimal beliefs of the agents, however when the learner has much more expertise, the behavior of the optimal agents is inverted as the learner becomes open-minded about the problem. In this case, the Prelec reweighting function 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 learner can get paired with the wrong advisor when the individual belief deviates significantly from the underlying prior value.

6. REFERENCES

[1] Glenn Ellison and Drew Fudenberg, "Rules of thumb for social learning," *J. Polit. Econ.*, vol. 101, no. 4, pp. 612–643, Aug. 1993.

[2] Vikram Krishnamurthy and H. Vincent 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.

[3] Joong Bum Rhim and Vivek K Goyal, "Social teaching: Being informative vs. being right in sequential decision making," in *Proc. 2013 IEEE Int. Symp. Inf. Theory*, July 2013, pp. 2602–2606.

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

[5] Sushil Bikhchandani, David Hirshleifer, and Ivo Welch, "Learning from the behavior of others: Conformity, fads, and informational cascades," *J. Econ. Perspect.*, vol. 12, no. 3, pp. 151–170, 1998.

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

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

[8] Boğaçhan Çelen and Shachar Kariv, "Observational learning under imperfect information," *Games Econ. Behav.*, vol. 47, no. 1, pp. 72–86, Apr. 2004.

[9] Amos Tversky and Daniel Kahneman, "Advances in prospect theory: Cumulative representation of uncertainty," *J. Risk Uncertainty*, vol. 5, no. 4, pp. 297–323, Oct. 1992.

[10] Drazen Prelec, "The probability weighting function," *Econometrica*, vol. 66, no. 3, pp. 497–527, May 1998.

[11] Ali al-Nowaihi and Sanjit Dhami, "Probability weighting functions," *Wiley Encyclopedia of Operations Research and Management Science*, Feb. 2011.