

Conflict in Distributed Hypothesis Testing with Quantized Prior Probabilities*

Joong Bum Rhim[†], Lav R. Varshney[‡], and Vivek K Goyal[†]

[†]Research Laboratory of Electronics, Massachusetts Institute of Technology

[‡]IBM Thomas J. Watson Research Center

Abstract

The effect of quantization of prior probabilities in a collection of distributed Bayesian binary hypothesis testing problems over which the priors themselves vary is studied, with focus on conflicting agents. Conflict arises from differences in Bayes costs, even when all agents desire correct decisions and agree on the meaning of correct. In a setting with fusion of local binary decisions by majority rule, Nash equilibrium local decision strategies are found. Assuming that agents follow Nash equilibrium decision strategies, designing quantizers for prior probabilities becomes a strategic form game; we discuss its Nash equilibria. We also propose two different constrained quantizer design games, find Nash equilibrium quantizer designs, and compare performance. The system has deadweight loss: equilibrium decisions are not Pareto optimal.

1 Introduction

Traditional analyses of Bayesian distributed detection and data fusion systems assume that agents have noisy sensors and that their communication is constrained [1,2]. This paper focuses on an additional limitation that agents—especially human ones—face and on the resultant effects on group performance: in forming local decisions, an agent cannot use a decision rule that is perfectly specialized to the particular object under study; instead, the agent must assign the object to a category and use a decision rule designed for that category. This categorization is justified both by information-processing limitations of the agent [3] and by limitations in rates of learning from training data [4, Sec. 5.3].

Agents are also traditionally assumed to share a single set of Bayes costs [1,2]. However, there are situations—especially in human decision-making groups such as juries and business units—where each agent in the group has a different preference for Type I and Type II errors. Since agents derive differing utility from the group decision, conflicts of interest arise and lead to game-theoretic considerations. A companion paper studies the setting where Bayes costs are shared by agents [5].

With this inspiration, we study Bayesian distributed binary hypothesis testing under an observation model specified by likelihood functions $f_{Y_i|H}(y|h_0)$ and $f_{Y_i|H}(y|h_1)$, where i indexes the agents. To formalize having an ensemble of possible objects under study, we consider the prior probability distribution over hypotheses $\{h_0, h_1\}$ to itself

*This material based upon work supported by National Science Foundation Grant No. 0729069.

be random with a known distribution. Since $\mathbb{P}(H = h_1) = 1 - \mathbb{P}(H = h_0)$, the pmf of H is described by a single scalar $p_0 = \mathbb{P}(H = h_0)$. We model p_0 as a realization of random variable P_0 with a known pdf f_{P_0} whose support is on the 1-dimensional probability simplex $[0, 1]$. Categorization of the ensemble of objects is equated with quantization of P_0 to one of K levels. Our focus is on the design of quantizers for P_0 that minimize Bayes risk; agents have their own perceptions of Bayes risk computed from their local Bayes costs and the probabilities of error of the global decision, which is made through majority vote.

Related work. Varshney and Varshney recently initiated the precise study of quantization of prior probabilities in Bayesian hypothesis testing and its implications on human decision-making, focusing on decision-making by a single agent [6]. This paper extends to decision making by multiple agents with conflicts of interest. The introduction of multiple agents highlights the role of strategic voting and introduces novel strategic quantizer design methods. The conflicts of interest arise from differing Bayes costs. Such differences are studied as *preference heterogeneity* in political science and economics [7] but seem to be absent in the engineering literature on distributed detection [1].

Most previous work on the effect of quantization in Bayesian distributed detection is focused on quantization of observations or on communication topology and rates among agents and to the fusion center [1, 8]. Here we do not consider quantization of observations, besides implicitly in forming local decisions.

Game theory originated in studies of decision-making with conflicts of interest [9], but the formulation of quantizer design games seems to be unique to this work.

Paper organization and preview of main results. Sec. 2 formalizes the setting and defines notation; in particular, it defines the Bayes risk objective functions of interest and the functions that take their place when prior probabilities are quantized. Sec. 3 discusses Nash equilibrium decision making for a group of agents and demonstrates that equilibrium strategies are not Pareto optimal. Sec. 4 formulates the game for design of prior probability quantizers and characterizes Nash equilibrium quantizer designs. Restriction to common quantizer partitions leads to two other quantizer design games with differing computable Nash equilibrium strategies. By measuring system performance, deadweight loss from conflicts of interest is quantified. The presence of deadweight loss demonstrates that group decision-making is best when agents share a common goal.

2 Problem Statement

Consider a group of agents deciding between $H = h_0$ and $H = h_1$ when $\mathbb{P}(H = h_0 | P_0 = p_0) = p_0$ is given. The scenario of interest is shown in Fig. 1 for the case of three decision making agents. For each i , Agent i (marked D_i) observes Y_i satisfying likelihood function $f_{Y_i|H}$ and sends a local decision $\hat{H}_i \in \{h_0, h_1\}$ to a fusion center. The observations are assumed to be conditionally independent given H . The fusion center determines $\hat{H} \in \{h_0, h_1\}$ by majority rule.

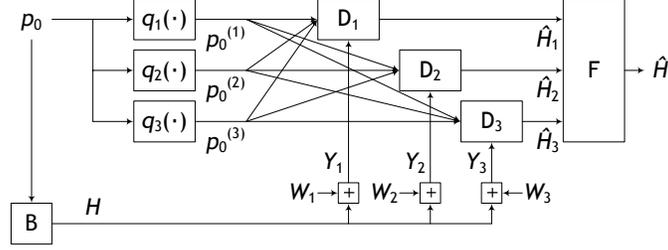


Fig. 1: A schematic diagram depicting the problem information pattern. The environment B generates a Bernoulli signal H . Its prior probability p_0 is quantized by three separate quantizers; the results are used by local agents D_i . Each agent also has access to H corrupted by i.i.d. noise W_i and all agents' Bayes costs $c_{10}^{(i)}$ and $c_{01}^{(i)}$. The fusion center F determines \hat{H} from the local decisions \hat{H}_i .

In the sequel we restrict attention to three agents and the additive white Gaussian noise (AWGN) measurement model depicted in Fig. 1, with noise variables $\{W_i\}_{i=1}^3$ i.i.d. $\sim \mathcal{N}(0, \sigma^2)$. Generalizations are certainly possible.

Agent i has (local) Type I and Type II error probabilities given by

$$P_{e,i}^I = \mathbb{P}(\hat{H}_i = h_1 | H = h_0) \quad \text{and} \quad P_{e,i}^{II} = \mathbb{P}(\hat{H}_i = h_0 | H = h_1).$$

Since global errors occur exactly when the majority of agents make local errors, global error probabilities can be expressed in terms of the local error probabilities:

$$P_E^I = \mathbb{P}(\hat{H} = h_1 | H = h_0) = P_{e,1}^I P_{e,2}^I + P_{e,2}^I P_{e,3}^I + P_{e,3}^I P_{e,1}^I - 2P_{e,1}^I P_{e,2}^I P_{e,3}^I, \quad (1)$$

$$P_E^{II} = \mathbb{P}(\hat{H} = h_0 | H = h_1) = P_{e,1}^{II} P_{e,2}^{II} + P_{e,2}^{II} P_{e,3}^{II} + P_{e,3}^{II} P_{e,1}^{II} - 2P_{e,1}^{II} P_{e,2}^{II} P_{e,3}^{II}. \quad (2)$$

All error probabilities depend on p_0 and on the decision rules used by the agents.

When p_0 is known, the goal of Agent i is to minimize the expected value of the i th *Bayes risk*

$$\tilde{R}_i = c_{10}^{(i)} p_0 P_E^I + c_{01}^{(i)} (1 - p_0) P_E^{II}, \quad (3)$$

where $c_{10}^{(i)}$ and $c_{01}^{(i)}$ are the positive *Bayes costs* for Agent i ; correct decisions incur no cost. Through the definitions of P_E^I and P_E^{II} , it is clear that \tilde{R}_i is the conditional mean of the Bayes cost given $P_0 = p_0$.

As depicted in Fig. 1, Agent i quantizes p_0 to $p_0^{(i)} = q_i(p_0)$ due to information-processing limitations. Thus, Agent i makes decisions to minimize i th *perceived Bayes risk*

$$\bar{R}_i = c_{10}^{(i)} p_0^{(i)} P_E^I + c_{01}^{(i)} (1 - p_0^{(i)}) P_E^{II}. \quad (4)$$

The decision threshold, and consequently $P_{e,i}^I$ and $P_{e,i}^{II}$, of each agent is determined based on $\{\bar{R}_i, i = 1, 2, 3\}$. However, the true Bayes risk is \tilde{R}_i , where P_E^I and P_E^{II} in (3) are affected by the perceived Bayes risks. We define mean Bayes risk (MBR) as a fidelity criterion for the quantizer q_i of $f_{P_0}(p_0)$:

$$\mathbb{E}[\tilde{R}_i] = \int_0^1 (c_{10}^{(i)} p_0 P_E^I(q_i(p_0)) + c_{01}^{(i)} (1 - p_0) P_E^{II}(q_i(p_0))) f_{P_0}(p_0) dp_0; \quad (5)$$

see also [6]. The MBR of an agent differs from that of other agents, but depends on the quantizers of all agents. Thus designing quantizers q_i is a game. Later we will find Nash equilibrium quantizer designs, but first we discuss a game-theoretic formulation of decision-making itself.

3 Equilibrium Decision Making

When agents share one cost function, they can collaborate to make decisions that are optimal for all. Heterogeneous preferences, however, restrict collaboration. For example, consider two agents that incur greater cost for Type II errors than for Type I errors and the other agent that incurs greater cost for Type I errors than for Type II. Then, the first two agents may tend to exaggerate $\widehat{H}_i = h_1$ claims so as to decrease probability of global Type II error. Hence the global decision would be more likely to be h_1 irrespective of the last agent's \widehat{H}_i .

Game theory provides useful methods to analyze agents' decision-making strategies under competition [9]. It is straightforward to describe the decision-making problem in strategic form with players \mathcal{I} , strategies $(S_i)_{i \in \mathcal{I}}$, and payoffs $(u_i)_{i \in \mathcal{I}}$ [9]. This is written as **Game I**.¹

Theorem 1. *Dominant strategies do not exist in Game I.*

Proof. It is sufficient to consider dominant strategies of Agent 1 due to the symmetry among agents. By definition, s_1^* is dominant if for all $(s_1, s_2, s_3) \in S_1 \times S_2 \times S_3$,

$$u_1(s_1^*, s_2, s_3) \geq u_1(s_1, s_2, s_3). \quad (6)$$

By defining $f_1(s_2, s_3) = P_{e,2}^I + P_{e,3}^I - 2P_{e,2}^I P_{e,3}^I$ and $f_2(s_2, s_3) = P_{e,2}^{II} + P_{e,3}^{II} - 2P_{e,2}^{II} P_{e,3}^{II}$, (6) is equivalent to

$$\begin{aligned} & c_{10}^{(1)} p_0^{(1)} P_{e,1}^I(s_1^*) f_1(s_2, s_3) + c_{01}^{(1)} (1 - p_0^{(1)}) P_{e,1}^{II}(s_1^*) f_2(s_2, s_3) \\ & \leq c_{10}^{(1)} p_0^{(1)} P_{e,1}^I(s_1) f_1(s_2, s_3) + c_{01}^{(1)} (1 - p_0^{(1)}) P_{e,1}^{II}(s_1) f_2(s_2, s_3) \triangleq h(s_1), \end{aligned} \quad (7)$$

where $P_{e,1}^I(s_1)$ and $P_{e,1}^{II}(s_1)$ denote error probabilities when Agent 1's decision threshold is s_1 . If s_1^* is dominant, then $h(s_1)$ should have a global minimum point at $s_1 = s_1^*$ irrespective of s_2 and s_3 .

Since the measurement noise W_1 has a continuous pdf, the decision rule yields $P_{e,1}^I$ and $P_{e,2}^{II}$ such that $P_{e,1}^{II}(s_1)$ is monotonically decreasing and strictly convex in $P_{e,1}^I(s_1)$. Then the location of the minimal extreme of $h(s_1)$ depends on $f_1(s_2, s_3)$ and $f_2(s_2, s_3)$. This implies that for any $c_{10}^{(1)}$, $c_{01}^{(1)}$, and $p_0^{(1)}$, no s_1^* exists such that $h(s_1)$ is minimum at $s_1 = s_1^*$ for all $(s_2, s_3) \in S_2 \times S_3$.

Likewise Agents 2 and 3 do not have dominant strategies either. \square

In addition to the lack of existence of dominant strategies, the only dominated strategies for any agent are $s_i = \infty$ and $s_i = -\infty$. Therefore, Game I is not solvable by iterative dominance. Each agent's decision rule depends on other agents' rules. Thus, we consider Nash equilibrium strategies, introducing a modified game: **Game $\hat{\mathbf{I}}$** .

¹All strategic form game formulations are presented together at the end of the paper.

Lemma 2 ([9]). *An infinite game $(\mathcal{I}, (S_i)_{i \in \mathcal{I}}, (u_i)_{i \in \mathcal{I}})$ has a pure strategy Nash equilibrium if its strategy spaces S_i are nonempty compact convex sets; its payoff functions $u_i(s_i, s_{-i})$ are continuous in s_{-i} ; and $u_i(s_i, s_{-i})$ are quasi-concave in s_i , where s_{-i} denotes a strategy profile of every player except i .*

Lemma 3. *Game I is equivalent to Game \hat{I} .*

Proof. The functions $P_{e,i}^I : \mathbb{R} \mapsto [0, 1]$ are bijective since Gaussian densities are continuous and greater than zero. Thus, there exist inverse functions $(P_{e,i}^I)^{-1}$ that are also bijective; t_i are uniquely determined by s_i and vice versa. Since choosing either t_i or s_i does not change the players' payoff, Game I and Game \hat{I} are equivalent. \square

Theorem 4. *A pure Nash equilibrium always exists in Game I.*

Proof. Consider Game \hat{I} , where strategy sets are compact and convex.

Due to the symmetry among players, it is sufficient to show concavity and continuity of the payoff function of Agent 1. From (1),

$$P_E^I = t_1 t_2 + t_2 t_3 + t_3 t_1 - 2t_1 t_2 t_3 = (t_2 + t_3 - 2t_2 t_3)t_1 + t_2 t_3,$$

which is linear in t_1 . Similarly, P_E^{II} in (2) is linear in τ_1 :

$$P_E^{II} = (\tau_2 + \tau_3 - 2\tau_2 \tau_3)\tau_1 + \tau_2 \tau_3,$$

where τ_i denotes the Type II error probability of Agent i and $0 \leq \tau_i \leq 1$. Note that $\tau_2 + \tau_3 - 2\tau_2 \tau_3 = \tau_2(1 - \tau_3) + (1 - \tau_2)\tau_3 \geq 0$. Since τ_1 is convex in t_1 by the characteristic of probabilities of errors [6], P_E^{II} is convex in t_1 . Hence, the payoff function of Agent 1, which is the sum of P_E^I and P_E^{II} multiplied by negative coefficients, is concave in t_1 .

Furthermore, $\hat{u}_1(t_1, t_2, t_3)$ is continuous in (t_2, t_3) as well as t_1 because τ_i is continuous in t_i .

Likewise, for other players, $\hat{u}_i(t_1, t_2, t_3)$ is concave in t_i and continuous in (t_{-i}, t_i) .

Thus by Lemma 2, Game \hat{I} has a pure Nash equilibrium (t_1^*, t_2^*, t_3^*) for any $c_{10}^{(i)}$, $c_{01}^{(i)}$, and $p_0^{(i)}$. By Lemma 3, Game I has a pure Nash equilibrium for any $c_{10}^{(i)}$, $c_{01}^{(i)}$, and $p_0^{(i)}$. \square

Due to existence, agents can always choose a Nash equilibrium strategy. A Nash equilibrium can be found by solving

$$\left. \frac{\partial \hat{u}_i(t_1, t_2, t_3)}{\partial t_i} \right|_{(t_1, t_2, t_3) = (t_1^*, t_2^*, t_3^*)} = 0, \quad i = 1, 2, 3, \quad (8)$$

and computing (s_1^*, s_2^*, s_3^*) that leads to (t_1^*, t_2^*, t_3^*) or directly by solving

$$\left. \frac{\partial u_i(s_1, s_2, s_3)}{\partial s_i} \right|_{(s_1, s_2, s_3) = (s_1^*, s_2^*, s_3^*)} = 0, \quad i = 1, 2, 3. \quad (9)$$

The reason that the latter is possible comes from the fact that

$$\frac{\partial \hat{u}_i(t_1, t_2, t_3)}{\partial t_i} = \frac{ds_i}{dt_i} \frac{\partial u_i(s_1, s_2, s_3)}{\partial s_i}, \quad i = 1, 2, 3, \quad (10)$$

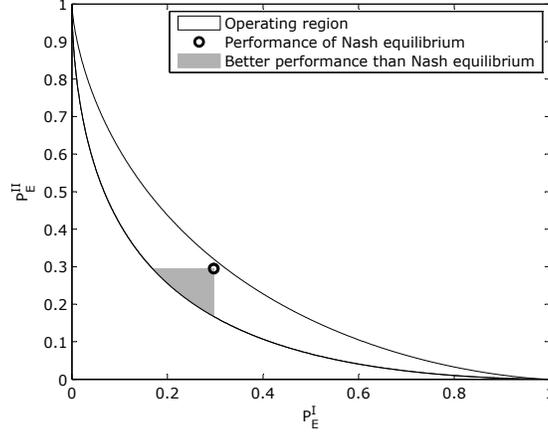


Fig. 2: The performance of the three agents whose Bayes costs are $(c_{10}^{(1)}, c_{01}^{(1)}) = (1, 4)$, $(c_{10}^{(2)}, c_{01}^{(2)}) = (4, 1)$, and $(c_{10}^{(3)}, c_{01}^{(3)}) = (4, 4)$ for $h_0 = 0$, $h_1 = 1$, $p_0^{(i)} = 0.5$ and $\sigma = 1$.

where $ds_i/dt_i < 0$. Note that (9) always has a solution s_i unless for both $j \neq i$, $s_j = \infty$ or $s_j = -\infty$. In that case, the global decision will always be h_0 or h_1 , respectively, regardless of the decision made by Agent i , which means Agent i does not play any role in the decision making.

The particular Nash equilibrium computed from (9) is not Pareto optimal, due to conflict. In Fig. 2, the operating point of the Nash equilibrium is located in the interior of the operating region. The agents can improve their performance by changing their decision rules such that any point to the lower left region is achieved. Since agents have different Bayes costs, however, they will not agree on how to optimize their decision rules and thus incur deadweight loss.

In the sequel, we assume that agents adopt Nash equilibrium decision rules.

4 Quantization of Prior Probabilities

Quantizer design can also be described in strategic form: **Game II**.

There are $2K - 1$ degrees of freedom in a strategy when agents use K -point quantizers: K for representation points $\{a_k^{(i)}\}_{k=1}^K$ and $K - 1$ for cell boundaries $\{b_k^{(i)}\}_{k=1}^{K-1}$. There is high computational complexity for $K \geq 2$ because each variable depends on some others. Consider examples of 2-point quantizers in Fig. 3. In the first example in Fig. 3a, conditional MBRs of Agent 1 within \mathcal{R}_1 and \mathcal{R}_2 are given by

$$\mathbb{E}[\tilde{R}_1]_{\mathcal{R}_1} = \int_0^{b_1^{(1)}} \left(c_{10}^{(1)} p_0 \bar{P}_E^I(a_1^{(1)}, a_1^{(2)}, a_1^{(3)}) + c_{01}^{(1)} (1 - p_0) \bar{P}_E^{II}(a_1^{(1)}, a_1^{(2)}, a_1^{(3)}) \right) f_{P_0}(p_0) dp_0, \quad (11)$$

$$\begin{aligned} \mathbb{E}[\tilde{R}_1]_{\mathcal{R}_2} &= \int_{b_1^{(1)}}^{b_1^{(2)}} \left(c_{10}^{(1)} p_0 \bar{P}_E^I(a_2^{(1)}, a_1^{(2)}, a_1^{(3)}) + c_{01}^{(1)} (1 - p_0) \bar{P}_E^{II}(a_2^{(1)}, a_1^{(2)}, a_1^{(3)}) \right) f_{P_0}(p_0) dp_0 \\ &+ \int_{b_1^{(2)}}^{b_1^{(3)}} \left(c_{10}^{(1)} p_0 \bar{P}_E^I(a_2^{(1)}, a_2^{(2)}, a_1^{(3)}) + c_{01}^{(1)} (1 - p_0) \bar{P}_E^{II}(a_2^{(1)}, a_2^{(2)}, a_1^{(3)}) \right) f_{P_0}(p_0) dp_0 \\ &+ \int_{b_1^{(3)}}^1 \left(c_{10}^{(1)} p_0 \bar{P}_E^I(a_2^{(1)}, a_2^{(2)}, a_2^{(3)}) + c_{01}^{(1)} (1 - p_0) \bar{P}_E^{II}(a_2^{(1)}, a_2^{(2)}, a_2^{(3)}) \right) f_{P_0}(p_0) dp_0, \end{aligned} \quad (12)$$

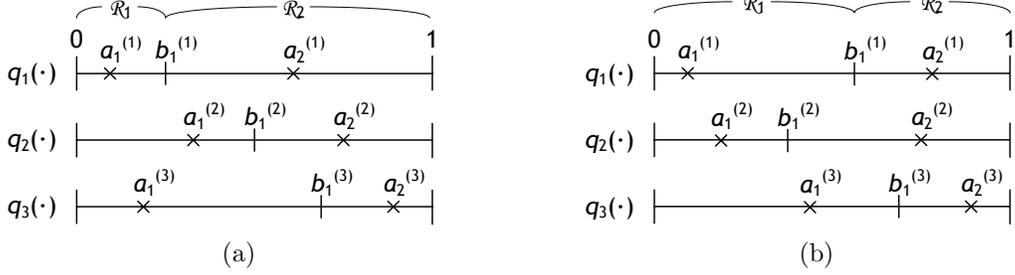


Fig. 3: Two examples of possible quantizers that agents use.

where $\bar{P}_E^I(x, y, z)$ and $\bar{P}_E^{II}(x, y, z)$ denote error probabilities when Agents 1, 2, and 3 use quantized prior probabilities x , y , and z . The dependency between $a_1^{(1)}$ and $a_2^{(1)}$ via $a_1^{(2)}$ and $a_1^{(3)}$ is revealed in (11) and (12). The dependency prevents agents from determining each variable separately.

Furthermore, the dependency varies with quantizer partition structure. In Fig. 3b, $a_1^{(1)}$ affects $a_2^{(2)}$ via $a_1^{(2)}$, $a_2^{(2)}$, and $a_1^{(3)}$. To find the best set of quantizers, we need to find the optimal set of quantizers for each of $\frac{(3(K-1))!}{(K-1)!(K-1)!(K-1)!}$ possible partition structures and compare them to choose the best one.

The dependency among representation points of different partitions occurs because quantizers used by agents have different partitions. If the quantizers have the same partitions, then each partition will be independent of other partitions and we can consider only dependency among agents.

Suppose all quantizers have the same partition, with fixed endpoints

$$0 = b_0 < b_1 < b_2 < \dots < b_{K-1} < b_K = 1.$$

The problem of designing quantizers reduces to the problem of choosing representation points, **Game III**.

Defining $\epsilon_k^I = \int_{\mathcal{R}_k} p_0 f_{P_0}(p_0) dp_0$ and $\epsilon_k^{II} = \int_{\mathcal{R}_k} (1 - p_0) f_{P_0}(p_0) dp_0$, a Nash equilibrium $(a_k^{(1)*}, a_k^{(2)*}, a_k^{(3)*})$ in Game III is given by

$$\frac{\partial v_{ik}(a_k^{(1)}, a_k^{(2)}, a_k^{(3)})}{\partial a_k^{(i)}} = -c_{10}^{(i)} \epsilon_k^I \frac{\partial P_E^I}{\partial a_k^{(i)}} - c_{01}^{(i)} \epsilon_k^{II} \frac{\partial P_E^{II}}{\partial a_k^{(i)}} = 0, \quad i = 1, 2, 3. \quad (13)$$

It is difficult to solve (13) because P_E^I and P_E^{II} are indirectly dependent on $a_k^{(i)}$. They are connected by the decision rule (s_1^*, s_2^*, s_3^*) , which is a Nash equilibrium in Game I for $p_0^{(i)} = a_k^{(i)}$, and that Nash equilibrium does not have a closed form.

There is an alternate way to find an equilibrium. For all objects that are categorized into \mathcal{R}_k , agents recognize their prior probabilities as $a_k^{(i)}$ and apply the same decision rule to them. Thus, the agents can directly choose their decision rules to apply to the objects, skipping the process of choosing representation points. This is written as **Game IV**.

Theorem 5. *There exists a Nash equilibrium in Game IV. As a result of Game IV, agents will have quantizers with the same representation points*

$$a_k^{(i)} = \mathbb{E}[P_0 | P_0 \in \mathcal{R}_k]. \quad (14)$$

Proof. We use the following payoff function instead of that in the definition of Game IV because scaling does not affect the choice of strategies:

$$\frac{\hat{v}_{ik}(\lambda_k^{(1)}, \lambda_k^{(2)}, \lambda_k^{(3)})}{\epsilon_k^I + \epsilon_k^{II}} = -c_{10}^{(i)} \left(\frac{\epsilon_k^I}{\epsilon_k^I + \epsilon_k^{II}} \right) P_E^I - c_{01}^{(i)} \left(1 - \frac{\epsilon_k^I}{\epsilon_k^I + \epsilon_k^{II}} \right) P_E^{II}. \quad (15)$$

Then Game I and Game IV are the same if $p_0^{(i)} = \epsilon_k^I / (\epsilon_k^I + \epsilon_k^{II})$, for $i = 1, 2, 3$. Therefore, choosing decision rules directly is equivalent to quantizing $p_0 \in \mathcal{R}_k$ to $\epsilon_k^I / (\epsilon_k^I + \epsilon_k^{II})$. Note that by definition, $\epsilon_k^I / (\epsilon_k^I + \epsilon_k^{II})$ is the average of P_0 conditioned on $P_0 \in \mathcal{R}_k$, which gives (14).

Game IV becomes Game I with (14). Therefore, Game IV has a Nash equilibrium by Theorem 4. \square

Working directly with decision thresholds rather than with representation points, a Nash equilibrium $(\lambda_k^{(1)*}, \lambda_k^{(2)*}, \lambda_k^{(3)*})$ is found by solving

$$\frac{\partial \hat{v}_{ik}(\lambda_k^{(1)}, \lambda_k^{(2)}, \lambda_k^{(3)})}{\partial \lambda_k^{(i)}} = 0, \quad i = 1, 2, 3. \quad (16)$$

Games III and IV look similar but they have different Nash equilibria. Game III hides the process of determining decision thresholds (s_1^*, s_2^*, s_3^*) by Game I for $(p_0^{(1)}, p_0^{(2)}, p_0^{(3)}) = (a_k^{(1)}, a_k^{(2)}, a_k^{(3)})$. Since not only s_1^* but also s_2^* and s_3^* depend on $a_k^{(1)}$,

$$\frac{\partial v_{ik}(a_k^{(1)}, a_k^{(2)}, a_k^{(3)})}{\partial a_k^{(i)}} = \sum_{j=1}^3 \frac{\partial \lambda_k^{(j)}}{\partial a_k^{(i)}} \frac{\partial \hat{v}_{ik}(\lambda_k^{(1)}, \lambda_k^{(2)}, \lambda_k^{(3)})}{\partial \lambda_k^{(j)}}, \quad i = 1, 2, 3.$$

Therefore, a Nash equilibrium $(a_k^{(1)*}, a_k^{(2)*}, a_k^{(3)*})$ of Game III does not necessarily lead to decision thresholds $(\lambda_k^{(1)*}, \lambda_k^{(2)*}, \lambda_k^{(3)*})$ that satisfy (16). Conversely, a Nash equilibrium $(\lambda_k^{(1)*}, \lambda_k^{(2)*}, \lambda_k^{(3)*})$ of Game IV is not the decision rules for $(p_0^{(1)}, p_0^{(2)}, p_0^{(3)}) = (a_k^{(1)*}, a_k^{(2)*}, a_k^{(3)*})$ that satisfy (13). The result in Fig. 4 shows that Games III and IV end up with different sets of quantizers. Note that the Bayes risks of agents who use quantized prior probabilities can be lower than those of agents who use true prior probabilities because decision rules of the latter agents are not optimal, as we discussed in Section 3.

5 Conclusions

This paper explores group decision-making problems under quantized prior probabilities. Agents who have different cost functions encounter conflicts of interest in decision making (Game I); the conflicts result in a Nash equilibrium that is not Pareto optimal. The agents also face conflicts of interest in quantizing prior probabilities. The conflicts lead to a quantizer design game (Game II) which is completely novel in the quantization literature. The complexity of the game is too high to compute an equilibrium due to dependence on the partition structure, however, we apply a reasonable simplification that the agents use identical categories. Then we have Games III

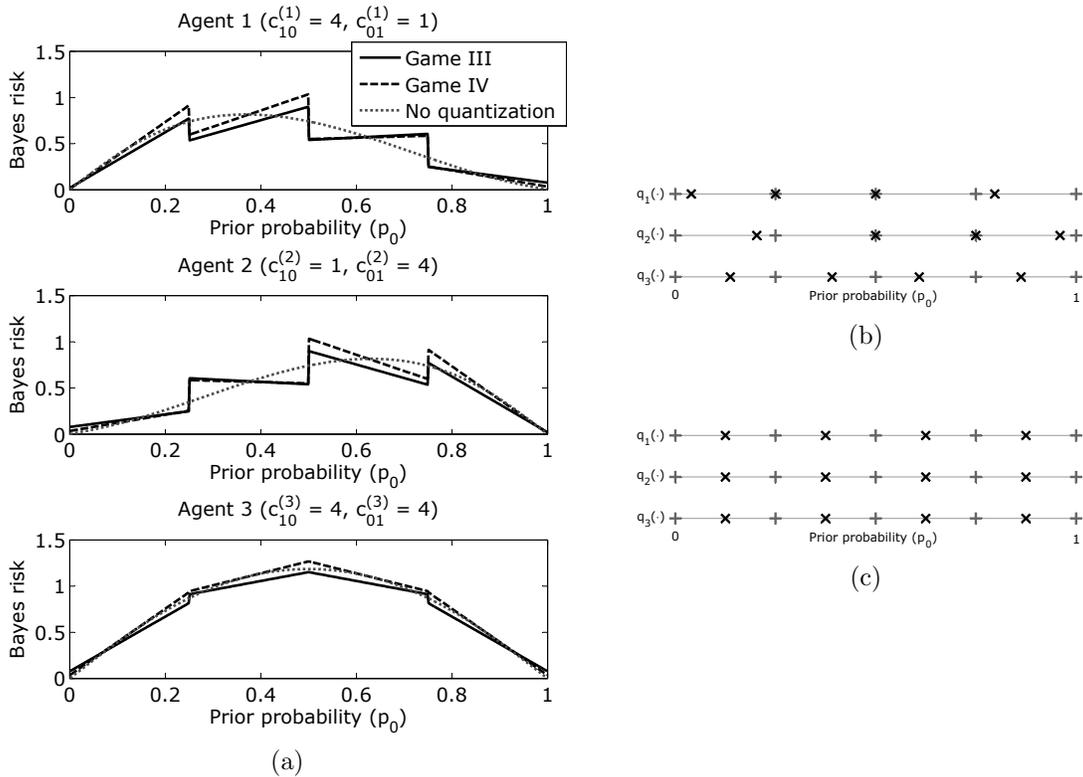


Fig. 4: A result of designing quantizers for uniformly distributed P_0 and given a set of endpoints $\{0, 0.25, 0.5, 0.75, 1\}$. (a) The resulting Bayes risks. (b) Quantizers obtained by Game III. (c) Quantizers obtained by Game IV. Endpoints and representation points of partitions are marked by + and \times , respectively.

and IV; both give equilibria in each category but are different in some sense. Game III designs representation points of quantizers, which affect decision making indirectly through the process of choosing decision rules (Game I). On the other hand, Game IV determines decision rules of agents for prior probabilities in each partition, which directly control the decision making. Interestingly, the Nash equilibrium of the latter is exactly the single-agent centroid condition. That is, each agent locally optimizes quantization after the coordination of partition regions is imposed.

References

- [1] P. K. Varshney, *Distributed Detection and Data Fusion*. Springer-Verlag, 1997.
- [2] D. Austen-Smith and J. S. Banks, "Information aggregation, rationality, and the Condorcet jury theorem," *Am. Polit. Sci. Rev.*, vol. 90, no. 1, pp. 34–45, Mar. 1996.
- [3] G. A. Miller, "The magical number seven, plus or minus two: Some limits on our capacity for processing information," *Psych. Rev.*, vol. 63, no. 2, pp. 81–97, Mar. 1956.
- [4] K. R. Varshney, "Frugal hypothesis testing and classification," Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA, Jun. 2010.

- [5] J. B. Rhim, L. R. Varshney, and V. K. Goyal, “Collaboration in distributed hypothesis testing with quantized prior probabilities,” in *IEEE Data Compression Conf.*, Mar. 2011, to appear.
- [6] K. R. Varshney and L. R. Varshney, “Quantization of prior probabilities for hypothesis testing,” *IEEE Trans. Signal Process.*, vol. 56, no. 10, pp. 4553–4562, Oct. 2008.
- [7] P. J. Coughlan, “In defense of unanimous jury verdicts: Mistrials, communication, and strategic voting,” *Am. Polit. Sci. Rev.*, vol. 94, no. 2, pp. 375–393, Jun. 2000.
- [8] T. S. Han and S.-I. Amari, “Statistical inference under multiterminal data compression,” *IEEE Trans. Inf. Theory*, vol. 44, no. 6, pp. 2300–2324, Oct. 1998.
- [9] D. Fudenberg and J. Tirole, *Game Theory*. Cambridge, MA: MIT Press, 1991.

Strategic Form Game Formulations

All games have players $\mathcal{I} = \{1, 2, 3\}$.

Game I: Determination of decision rules

- The strategy of Agent $i \in \mathcal{I}$ is decision threshold $s_i \in \mathbb{R}$.
- The payoff function of Agent $i \in \mathcal{I}$ is $u_i = -\bar{R}_i = -c_{10}^{(i)} p_0^{(i)} P_E^I - c_{01}^{(i)} (1 - p_0^{(i)}) P_E^{II}$, which is the negative perceived Bayes risk.

Game $\hat{\mathbf{I}}$: Determination of decision rules in terms of the probability of error

- The strategy of Agent $i \in \mathcal{I}$ is the Type I error probability $t_i \in [0, 1]$. (This is related to the decision threshold through $t_i = P_{e,i}^I(s_i)$.)
- The payoff function of Agent $i \in \mathcal{I}$ is $\hat{u}_i(t_1, t_2, t_3) = u_i(s_1, s_2, s_3) = -\bar{R}_i = -c_{10}^{(i)} p_0^{(i)} P_E^I - c_{01}^{(i)} (1 - p_0^{(i)}) P_E^{II}$, which is the negative perceived Bayes risk.

Game II: Determination of quantizers for prior probabilities

- The strategy of Agent $i \in \mathcal{I}$ is $s_i = (a_1^{(i)}, \dots, a_K^{(i)}, b_1^{(i)}, \dots, b_{K-1}^{(i)})$, which represents a quantizer with representation point $a_k^{(i)}$ for k th region $\mathcal{R}_k = (b_{k-1}^{(i)}, b_k^{(i)}) \subset [0, 1]$.
- The payoff function of Agent $i \in \mathcal{I}$ is $v_i = -\mathbb{E}[\tilde{R}_i] = -\int \tilde{R}_i f_{P_0}(p_0) dp_0$, which is the negative mean Bayes risk.

Game III: Determination of representation points for fixed category $\mathcal{R}_k = (b_{k-1}, b_k]$

- The strategy of Agent $i \in \mathcal{I}$ is representation point $a_k^{(i)} \in \mathcal{R}_k = (b_{k-1}, b_k]$.
- The payoff function of Agent $i \in \mathcal{I}$ is $v_{ik}(a_k^{(1)}, a_k^{(2)}, a_k^{(3)}) = -c_{10}^{(i)} \epsilon_k^I P_E^I - c_{01}^{(i)} \epsilon_k^{II} P_E^{II}$, which is the negative mean Bayes risk.

Game IV: Determination of decision rules for fixed category $\mathcal{R}_k = (b_{k-1}, b_k]$

- The strategy of Agent $i \in \mathcal{I}$ is $\lambda_k^{(i)}$, which is a decision threshold for $p_0 \in \mathcal{R}_k$.
 - The payoff function of Agent $i \in \mathcal{I}$ is $\hat{v}_{ik}(\lambda_k^{(1)}, \lambda_k^{(2)}, \lambda_k^{(3)}) = -c_{10}^{(i)} \epsilon_k^I P_E^I - c_{01}^{(i)} \epsilon_k^{II} P_E^{II}$, which is the negative mean Bayes risk.
-