On optimal quantization rules for some sequential decision problems

XuanLong Nguyen
Department of EECS
University of California, Berkeley
xuanlong@cs.berkeley.edu

Martin J. Wainwright
Department of Statistics and Department of EECS
University of California, Berkeley
wainwrig@stat.berkeley.edu

Michael I. Jordan
Department of Statistics and Department of EECS
University of California, Berkeley
jordan@stat.berkeley.edu

June 14, 2006

Technical Report 708
Department of Statistics
University of California, Berkeley

Abstract

We consider the problem of sequential decentralized detection, a problem that entails several interdependent choices: the choice of a stopping rule (specifying the sample size), a global decision function (a choice between two competing hypotheses), and a set of quantization rules (the local decisions on the basis of which the global decision is made). In this paper we resolve an open problem concerning whether optimal local decision functions for the Bayesian formulation of sequential decentralized detection can be found within the class of stationary rules. We develop an asymptotic approximation to the optimal cost of stationary quantization rules and show how this approximation yields a negative answer to the stationarity question. We also consider the class of blockwise stationary quantizers and show that asymptotically optimal quantizers are likelihood-based threshold rules.[1]

Keywords: sequential detection, decentralized detection, quantizer design, decision-making under constraints

1 Introduction

In this paper, we consider the problem of sequential decentralized detection (see, e.g., [13, 14, 7]). Detection is a classical discrimination or hypothesis-testing problem, in which observations $\{X_1, X_2, \ldots\}$ are assumed to be drawn i.i.d. from the conditional distribution $\mathbb{P}(\cdot \mid H)$ and the goal is to infer the value of the random variable H, which takes values in $\{0, 1\}$. Placing this problem in a communication-theoretic context, a

[1]Part of this work was presented at the International Symposium on Information Theory, July 2006, Seattle, WA.

decentralized detection problem is a hypothesis-testing problem in which the decision-maker is not given access to the raw data points X_n, but instead must infer H based only on a set of *quantization rules* or *local decision functions*, $\{U_n = \phi_n(X_n)\}$. Finally, placing the problem in a real-time context, the *sequential decentralized detection problem* involves a data sequence, $\{X_1, X_2, \ldots\}$, and a corresponding sequence of summary statistics, $\{U_1, U_2, \ldots\}$, determined by a sequence of local decision rules $\{\phi_1, \phi_2, \ldots\}$. The goal is to design both the local decision functions and to specify a global decision rule so as to predict H in a manner that optimally trades off accuracy and delay. In short, the sequential decentralized detection problem is the communication-constrained extension of classical formulation of sequential centralized decision-making problems [see, e.g, 10, 6] to the decentralized setting.

In setting up a general framework of sequential decentralized problems, Veeravalli et al. [15] defined five problems ("Case A" through "Case E"), distinguished from one another by the amount of information available to the local sensors. In particular, in Case E, the local sensors are provided with memory and with feedback from the global decision-maker (also known as the *fusion center*), so that each sensor has available to it the current data, X_n, as well as all of the summary statistics from all of the other local sensors. In other words, each local sensor has the same snapshot of past state as the fusion center; this is an instance of a so-called "quasi-classical information structure" [5] for which dynamic programming (DP) characterizations of the optimal decision functions are available. Veeravalli et al. [15] exploit this fact to show that the decentralized case has much in common with the centralized case, in particular that likelihood ratio tests are optimal local decision functions at the sensors and that a variant of a sequential probability ratio test is optimal for the decision-maker.

Unfortunately, however, part of the spirit of the decentralized detection is arguably lost in Case E, which requires full feedback. In particular, in applications such as power-constrained sensor networks, we generally do not wish to assume that there are high-bandwidth feedback channels from the decision-maker to the sensors, nor do we wish to assume that the sensors have unbounded memory. Most suited to this perspective—and the focus of this paper—is Case A, in which the local decisions are of the simplified form $\phi_n(X_n)$; i.e., neither local memory nor feedback are assumed to be available.

Noting that Case A is not amenable to dynamic programming and is presumably intractable, Veeravalli et al. [15] suggest restricting the analysis to the class of *stationary* local decision functions; i.e., local decision functions ϕ_n that are independent of n. They conjecture that stationary decision functions may actually be optimal in the setting of Case A (given the intuitive symmetry and high degree of independence of the problem in this case), even though it is not possible to verify this optimality via DP arguments. The truth or falsity of this conjecture has remained open since it was first posed by Veeravalli et al. [15, 14].[2]

In this paper, we resolve this question by showing that stationary decision functions are, in fact, *not* optimal for decentralized problems of type A. More specifically, we do so by providing a simple lemma characterizing the asymptotically optimal cost in the Bayesian formulation (i.e., when the cost per sample goes to 0). This lemma has analogs in the Neyman-Pearson formulation (see, e.g., [11, 6]) and proves to be very useful in the study of optimal decision functions in the Bayesian formulation. It allows us to construct counterexamples to the stationarity conjecture, both in an exact and an asymptotic setting. In the asymptotic setting, we show that in general there are always a range of prior probabilities for which stationary strategies are suboptimal. We note in passing that an intuition for the source of this suboptimality is easily provided—it is due to the asymmetry of the Kullback-Leibler (KL) divergence.

[2]In a related formulation of the problem, i.e., the Neyman-Pearson formulation, Mei [7] surmised that the pair of optimal quantizers for minimizing the expected stopping times with respect to two hypotheses (which are proven to be stationary) need not be the same. In some sense, this open question appears to be the equivalent of Veervalli's conjecture in the Baysian formulation, which we address in this paper.

It is well known that optimal quantizers *when unrestricted* are necessarily likelihood-based threshold rules [13]. Our counterexamples and analysis imply that the thresholds need not be stationary (i.e., the threshold may differ from sample to sample). In the remainder of the paper, we addresses a partial converse to this result: specifically, if we restrict ourselves to stationary (or blockwise stationary) quantizer designs, then there exists an optimal design that is a threshold rule based on the likelihood ratio. We prove this result by establishing a quasiconcavity result for the asymptotically optimal cost function.

The remainder of this paper is organized as follows. We begin in Section 2 with background on the Bayesian formulation of sequential detection problems, and Wald's approximation. Section 3 provides a simple asymptotic approximation of the optimal cost that underlies our main analysis in Section 4. In Section 5, we establish the existence of optimal decision rules which are likelihood-based threshold rules, when restricted to be blockwise stationary. We conclude in Section 6.

2 Background

This section provides necessary background on the Bayesian formulation of sequential detection problems, a dynamic programming chacterization and Wald's approximation of the optimal cost.

2.1 Sequential detection and dynamic programming

Let $\mathbb{P}_0$ and $\mathbb{P}_1$ represent the class-conditional distributions of X, when conditioned on $\{H = 0\}$ and $\{H = 1\}$ respectively. Focusing the Bayesian formulation of the sequential detection problem [10, 14], we let $\pi^1 = \mathbb{P}(H = 1)$ and $\pi^0 = \mathbb{P}(H = 0)$ denote the prior probabilities of the two hypotheses. A sequential decision rule consists of a *stopping time* N (defined with respect to the sigma field $\sigma(X_1, \ldots, X_N)$), and a decision function γ (measurable with respect to $\sigma(X_1, \ldots, X_N)$). The cost function is a weighted sum of the sample size N and the probability of incorrect decision

$$R(N, \gamma) := \mathbb{E}\{cN + \mathbb{I}[\gamma(X_1, \ldots, X_N) \neq H]\}, \tag{1}$$

where $c > 0$ is the incremental cost of each sample. The overall goal is to choose the pair (N, γ) so as to minimize the expected loss (1).

Assume that $\mathbb{P}_0$ and $\mathbb{P}_1$ are absolutely continuous with respect to one another. For convenience, in the sequel we shall frequently use $f^0(x)$ and $f^1(x)$ to denote the respective density functions with respect to some dominating measure (e.g., Lesbegue for continous variables, or counting measure for discrete-valued variables).

It is well known that the optimal solution of the sequential decision problem can be characterized recursively using dynamic programming (DP) arguments [1, 17, 10]. Furthermore, as we develop in detail in Appendix A, the DP characterization holds even when $X_1, X_2, \ldots$ are independent but not identically distributed conditional on H. Although very useful for classical (centralized) sequential detection, the DP approach is not always straightforward to apply to *decentralized* versions of sequential detection. We will find, however, that the DP approach is useful for evaluating the accuracy of *approximations* to the optimal cost.

2.2 Wald's approximation

When all $X_1, X_2, \ldots$ are i.i.d. conditioned on H, there is an alternative to the DP approach for approximating cost of the optimal sequential test, originally due to Wald (cf. [11]), that we describe here. It can be shown that the optimal stopping rule for the cost function (1) takes the form

$$N = \inf\big\{n \geq 1 \;\big|\; L_n(X_1, \ldots, X_n) := \sum_{i=1}^{n} \log \frac{f^1(X_i)}{f^0(X_i)} \notin (a, b)\big\}, \tag{2}$$

for some real numbers $a < b$. Given this stopping rule, the optimal decision function has the form

$$\gamma(L_N) = \begin{cases} 1 & \text{if } L_N \geq b, \\ 0 & \text{if } L_N \leq a. \end{cases} \tag{3}$$

We now develop an alternative expression for the optimal cost of the decision rule (3). Consider the two types of error:

$$\begin{aligned} \alpha &= \mathbb{P}_0(\gamma(L_N) \neq H) = \mathbb{P}_0(L_N \geq b) \\ \beta &= \mathbb{P}_1(\gamma(L_N) \neq H) = \mathbb{P}_1(L_N \leq a). \end{aligned}$$

Now define $\mu^1 = \mathbb{E}_1[\log \frac{f^1(X_1)}{f^0(X_1)}] = D(f^1||f^0)$ and $\mu^0 = -\mathbb{E}_0[\log \frac{f^1(X_1)}{f^0(X_1)}] = D(f^0||f^1)$. With this notation, the cost $J(a, b)$ of the decision rule based on envelopes a and b can be written as

$$\begin{aligned} J(a, b) &= \mathbb{E}\big\{cN + \mathbb{I}[\gamma(X_1, \ldots, X_N) \neq H]\big\} \\ &= \pi^1 \mathbb{E}_1(cN + \mathbb{I}[L_N \leq a]) + \pi^0 \mathbb{E}_0(cN + \mathbb{I}[L_N \geq b]) \qquad (4) \\ &= c\pi^1 \frac{\mathbb{E}_1 L_N}{\mu^1} + c\pi^0 \frac{\mathbb{E}_0 L_N}{-\mu^0} + \pi^0 \alpha + \pi^1 \beta, \qquad (5) \end{aligned}$$

where the second line follows from Wald's equation [16]. With a slight abuse of notation, we shall also use $D(\alpha, \beta)$ to denote a function in $[0, 1]^2 \to \mathbb{R}$ such that:

$$D(\alpha, \beta) := \alpha \log \frac{\alpha}{\beta} + (1 - \alpha) \log \frac{1 - \alpha}{1 - \beta}.$$

Let us now develop a series of approximations of the cost, following Wald (cf. [11]). To begin, the errors α and β are related to a and b by the classical inequalties $\alpha \leq (1 - \beta)/e^b$ and $\beta \leq e^a(1 - \alpha)$. In general, these inequalities need not hold with equality because the likelihood ratio L_N at the optimal stopping time N might overshoot either a or b (instead of attaining precisely the value a or b at the stopping time). Wald's approximation is based on ignoring this overshoot and replacing the inequalities by equalities to solve for corresponding values of α and β:

$$\alpha \approx \frac{1 - e^a}{e^b - e^a} \text{ and } \beta \approx \frac{e^{-b} - 1}{e^{-b} - e^{-a}}. \tag{6}$$

We can also express a and b in terms of α and β:

$$a \approx a(\alpha, \beta) := \log \frac{\beta}{1 - \alpha} \text{ and } b \approx b(\alpha, \beta) := \log \frac{1 - \beta}{\alpha}. \tag{7}$$

The mapping (6) and (7) between (a, b) and (α, β) yields the following approximation

$$\frac{\mathbb{E}_0[L_N]}{-\mu^0} \approx \frac{(1 - \alpha)a + \alpha b}{-\mu^0} \tag{8a}$$

$$\frac{\mathbb{E}_1[L_N]}{\mu^1} \approx \frac{(1 - \beta)b + \beta a}{\mu^1} \tag{8b}$$

Plugging this approximation into (4) and using (7), we obtain *Wald's approximation* for the cost of a sequential test with error α and β:

$$G(\alpha, \beta) := c\pi^0 \frac{D(\alpha, 1-\beta)}{\mu^0} + c\pi^1 \frac{D(1-\beta, \alpha)}{\mu^1} + \pi^0\alpha + \pi^1\beta, \tag{9}$$

In the following section we exploit this approximation in our analysis of quantizer design.

3 Characterization of optimal quantizers

Turning now to the decentralized setting, the primary challenge lies in the design of the quantization rules ϕ_n. A quantization rule ϕ_n is a function that maps $\mathcal{X}$ to the discrete space $\mathcal{U} = \{0, \ldots, K-1\}$ for some natural number K. Any fixed set of quantization rules ϕ_n yields a sequence of compressed data $U_n = \phi_n(X_n)$, to which the classical theory can be applied. We are thus interested in choosing quantization rules $\phi_1, \phi_2, \ldots$ so that the error, resulting from applying the optimal sequential test to the sequence of sufficient statistics $U_1, U_2, \ldots$, is minimized over some space of quantization rules. In the decentralized setting, we use

$$f^i_{\phi_n}(u) \quad := \quad \mathbb{P}_i(\phi_n(X_n) = u), \qquad \text{for} \quad i = 0, 1,$$

to denote the distributions of the compressed data, conditioned on the hypothesis.

We say that a quantizer design is *stationary* if the rule ϕ_n is independent of n; in this case, we simplify the notation to f^1_ϕ and f^0_ϕ. In addition, we define the KL divergences $\mu^1_\phi := D(f^1_\phi || f^0_\phi)$ and $\mu^0_\phi := D(f^0_\phi || f^1_\phi)$. Moreover, let G_ϕ denote the analogue of the function G in equation (9), defined using $\mu^i_\phi, i = 0, 1$. In this section, we describe how—by exploiting Wald's approximation for sequential problems—it is possible to provide an asymptotic characterization of the optimal cost of any stationary quantization rule.

3.1 Approximate quantizer design

Given a fixed stationary quantizer ϕ, Wald's approximation (9) suggests the following strategy for approximating the cost of sequential detection strategy. For a given set of errors α and β, first assign the values of thresholds $a = a(\alpha, \beta)$ and $b = b(\alpha, \beta)$ using approximation (7). Then use the quantity $G_\phi(\alpha, \beta)$ as an approximation to the true cost $J_\phi(a, b)$. This approximation essentially ignores the overshoot of the likelihood ratio L_N. It is possible to analyze this overshoot to obtain a finer approximation (cf. [11, 6, 9]). For the purpose of quantizer design, however, as we shall see, the approximation error incurred from ignoring the overshoot is at most $O(c)$, whereas the choice of quantizer ϕ generally results in a change of the order $\Theta(c \log c^{-1})$.

A key assumption in our analysis is that

$$\sup_{\phi \in \Phi} \sup_{u \in \mathcal{U}} \log(f^1_\phi(u)/f^0_\phi(u)) \leq M \tag{10}$$

for some constant M over a class Φ of quantizer functions $\phi : \mathcal{X} \to \mathcal{U}$. Note that this assumption holds in many cases of interest. For instance, when X_n takes its values in a finite set of discrete elements $\mathcal{X}$, and both $\mathbb{P}_1$ and $\mathbb{P}_0$ place positive probability mass on all values in $\mathcal{X}$, the assumption clearly holds. The assumption also holds when $\mathcal{X}$ is a continuous domain and $\sup_{x \in \mathcal{X}} |\log(f^1(x)/f^0(x))| < +\infty$.

The following proposition guarantees that the approximation G_ϕ is asymptotically exact up to an additive error of size $O(c)$, and provides a basis for characterizing the optimal cost:

Proposition 1. (a) *The error in the approximation* (9) *is bounded as*

$$|J_\phi(a,b) - G_\phi(\alpha,\beta)| \le c\,M\left(\frac{\pi^0}{\mu_\phi^0} + \frac{\pi^1}{\mu_\phi^1}\right). \tag{11}$$

(b) *Define the optimal cost* $J_\phi^* = \inf_{a,b} J_\phi(a,b)$. *Then as* $c \to 0$, *we have*

$$J_\phi^* = \left(\frac{\pi^0}{\mu_\phi^0} + \frac{\pi^1}{\mu_\phi^1}\right) c\log\frac{1}{c} + O(c). \tag{12}$$

Proof: (a) We begin by bounding the error in the approximation (7). By definition of the stopping time N, we have either (i) $b \le L_N \le b + M$ or (ii) $a - M \le L_N \le a$. Consider all realizations $u_1, \ldots, u_n$ for which condition (i) holds; for any such sequence, we have

$$\begin{aligned} e^b P_0(u_1,\ldots,u_n) &\le P_1(u_1,\ldots,u_n) \\ e^{(b+M)} P_0(u_1,\ldots,u_n) &\ge P_1(u_1,\ldots,u_n). \end{aligned}$$

Taking a sum over all such realizations, using the definition of α and β, and performing some algebra yields the inequality $e^b\alpha \le 1-\beta \le e^{b+M}\alpha$, or equivalently $b \le b(\alpha,\beta) = \log\frac{1-\beta}{\alpha} \le b+M$. Similar reasoning for case (ii) yields $a - M \le a(\alpha,\beta) = \log\frac{\beta}{1-\alpha} \le a$. Now, note that

$$\mathbb{E}_0 L_N = \alpha\mathbb{E}_0[L_N | L_N \ge b] + (1-\alpha)\mathbb{E}_0[L_N | L_N \le a].$$

Conditioning on the event $L_N \in [b, b+M]$, we have $|L_N - b(\alpha,\beta)| \le M$. Similarly, conditioning on the event $L_N \in [a-M, a]$, we have $|L_N - b(\alpha,\beta)| \le M$. This yields $|\mathbb{E}_0 L_N - (-D(\alpha, 1-\beta))| \le M$. Similar reasoning yields $|\mathbb{E}_1 L_N - D(1-\beta,\alpha)| \le M$.

(b) By part (a), it suffices to establish the asymptotic behavior (12) for the quantity $\widetilde{J}_\phi(a,b) = \inf_{\alpha,\beta} G_\phi(\alpha,\beta)$, where the infimum is taken among pairs of realizable error probabilities (α,β). Moreover, we only need to consider the asymptotic regime $\alpha+\beta \to 0$, since the error probabilities α and β vanish as $c \to 0$. It is simple to see that $D(1-\beta,\alpha) = \log(1/\alpha) + o(1)$, and $D(1-\alpha,\beta) = \log(1/\beta) + o(1)$. Hence, $\inf_{\alpha,\beta} G_\phi(\alpha,\beta)$ can be expressed as

$$\inf_{\alpha,\beta}\left\{\pi^0\alpha + \pi^1\beta + c\pi^0\frac{\log(1/\beta)}{\mu_\phi^0} + c\pi^1\frac{\log(1/\alpha)}{\mu_\phi^1}\right\} + o(c). \tag{13}$$

This infimum, taken over all positive (α,β), is achieved at $\alpha^* = \frac{c\pi^1}{\mu_\phi^1\pi^0}$ and $\beta^* = \frac{c\pi^0}{\mu_\phi^0\pi^1}$. We now show that these error probabilities can be approximately realized (for c small) by using a sufficiently large threshold $b > 0$ and small threshold $a < 0$ while incuring an approximation cost of order $O(c)$. Indeed, let us choose thresholds a' and b' such that $e^{-(b'+M)}/2 \le \alpha^* \le e^{-b'}$, and $e^{a'-M}/2 \le \beta^* \le e^{a'}$. Let α' and β' be the corresponding errors associated with these two thresholds. As proved above, we also have $\alpha' \in (e^{-(b'+M)}/2, e^{-b'})$ and $\beta' \in (e^{a'-M}/2, e^{a'})$. Clearly, $|\alpha^* - \alpha'| \le e^{-b'}(1-e^{-M}) = O(\alpha^*) = O(c)$. Similarly, $|\beta^* - \beta'| = O(c)$. By the mean value theorem,

$$|\log(1/\alpha^*) - \log(1/\alpha')| \le |\alpha^* - \alpha'|e^{b'+M} \le e^M(1-e^{-M}) = O(1).$$

Similarly, $\log(1/\beta^*) - \log(1/\beta') = O(1)$. Hence, the approximation of (α^*,β^*) by the realizable (α',β') incurs a cost at most $O(c)$.

c	0.01	0.009	0.008	0.007	0.006	0.005	0.004	0.003	0.002	0.001
J^*	0.0320	0.0297	0.0269	0.0241	0.0212	0.0178	0.0145	0.0112	0.0078	0.0042
$\tilde{J}$	0.0302	0.0277	0.0250	0.0224	0.0196	0.0168	0.0139	0.0108	0.0076	0.0041
$\tilde{J}/J^*$	0.9447	0.9313	0.9313	0.9288	0.9268	0.9409	0.9550	0.9682	0.9772	0.9737

Table 1. Comparison of the (exact) optimal cost J^* computed by the dynamic programming method, and an approximation of the optimal cost denoted by $\tilde{J}$ using Eq. (14a) as c decreases.

Plugging the quantities α^*, β^* into equation (13) yields

$$\inf_{\alpha,\beta} G_\phi(\alpha,\beta) = \left(\frac{\pi^0}{\mu_\phi^0} + \frac{\pi^1}{\mu_\phi^1}\right) c \log\frac{1}{c} + \left(\frac{\pi^0}{\mu_\phi^0} + \frac{\pi^1}{\mu_\phi^1} + \frac{\pi^0}{\mu_\phi^0}\log\frac{\mu_\phi^0\pi^1}{\pi^0} + \frac{\pi^1}{\mu_\phi^1}\log\frac{\mu_\phi^1\pi^0}{\pi^1}\right) c + O(c) \quad (14a)$$

$$= \left(\frac{\pi^0}{\mu_\phi^0} + \frac{\pi^1}{\mu_\phi^1}\right) c \log\frac{1}{c} + O(c) \quad (14b)$$

as claimed.

□

3.2 A simple illustration

Consider two simple hypotheses $\mathbb{P}_0$ and $\mathbb{P}_1$ where X takes its values from the set $\{1, 2\}$. In particular, $[\mathbb{P}_0(1)\ \mathbb{P}_0(2)] = [0.8\ 0.2]$ and $[\mathbb{P}_1(1)\ \mathbb{P}_1(2) = [0.01\ 0.99]$. The priors are $\pi^0 = \pi^1 = 0.5$. Table 1 shows that the approximation given by (14a) can provide a reasonable approximation to the cost for the optimal sequential test when c is sufficiently small.

4 Suboptimality of stationary designs

We now consider the structure of optimal quantizers. It was shown by Tsitsiklis [13] that optimal quantizers ϕ_n take the form of threshold rules based on the likelihood ratio $f^1(X_n)/f^0(X_n)$. Veeravalli et al. [15, 14] asked whether these rules can be taken to be stationary, a problem that has remained open. In this section, we resolve this question with a negative answer. First, we provide a counterexample in which the optimal quantizer is not stationary. Next, we show that using a stationary quantizer can be suboptimal even in an asymptotic sense (i.e., as $c \to 0$).

4.1 Illustrative counterexample

We begin with a simple but concrete demonstration of the suboptimality of stationary designs. Consider a problem in which $X \in \mathcal{X} = \{1, 2, 3\}$ and the conditional distributions take the form

$$f^0(x) = \left[\tfrac{8}{10}\quad \tfrac{1999}{10000}\quad \tfrac{1}{10000}\right] \text{ and } f^1(x) = \left[\tfrac{1}{3}\quad \tfrac{1}{3}\quad \tfrac{1}{3}\right].$$

Suppose that the prior probabilities are $\pi^1 = \frac{8}{100}$ and $\pi^0 = \frac{92}{100}$, and that the cost for each sample is $c = \frac{1}{100}$.

If we restrict to binary quantizers (i.e., $\mathcal{U} = \{0, 1\}$), then there are only three possible quantizers:

1. Design A: $\phi_A(X_n) = 0 \iff X_n = 1$. As a result, the corresponding distribution for U_n is specified by $f^0_{\phi_A}(u_n) = [\frac{4}{5}\ \frac{1}{5}]$ and $f^1_{\phi_A}(u) = [\frac{1}{3}\ \frac{2}{3}]$.

Method	$J_A(0.08)$	$J_B(0.08)$	$J_C(0.08)$	$J_*(0.08)$
Cost	0.0567	0.0532	0.0800	0.0528

Table 2. Numerically computed costs for the three stationary designs J_A, J_B and J_C, for the mixed design J_*.

2. Design B: $\phi_B(X_n) = 0 \iff X_n \in \{1, 2\}$. The corresponding distribution for U_n is given by $f^0_{\phi_B}(u) = [\frac{9999}{10000}\ \frac{1}{10000}]$ and $f^1_{\phi_B}(u) = [\frac{2}{3}\ \frac{1}{3}]$.
3. Design C: $\phi_C(X_n) = 0 \iff X_n \in \{1, 3\}$. The corresponding distribution for U_n is specified by $f^0_{\phi_C} \sim [\frac{8001}{10000}\ \frac{1999}{10000}]$ and $f^1_{\phi_C}(u) = [\frac{2}{3}\ \frac{1}{3}]$.

Now consider the three stationary strategies, each of which uses only one fixed design, A, B or C. For any given stationary quantization rule ϕ, we have a classical centralized sequential problem, for which the optimal cost (achieved by a SPRT) can be computed using a dynamic-programming procedure [17, 1]. Accordingly, for each stationary strategy, we compute the optimal cost function J for 10^6 points on the p-axis by performing 300 updates of Bellman's equation (24). In all cases, the difference in cost between the 299th and 300th updates is less than 10^{-6}. Let J_A, J_B and J_C denote the optimal cost function for sequential tests using all A's, all B's, and all C's, respectively. When evaluated at $\pi^1 = 0.08$, these computations yield the numerical values shown in Table 2.

Finally, we consider a non-stationary rule obtained by applying design A for only the first sample, and applying design B for the remaining samples. Again using Bellman's equation, we find that the cost for this design is

$$J_* = \min\{g(0.08), c + J_B(P(H=1|u_1=0))P(u_1=0) + J_B(P(H=1|u_1=1))P(u_1=1)\} = 0.052767,$$

which is better than any of the stationary strategies.

Note the slim improvement (0.0004) of J_* over the best stationary rule J_B. This is due in part to the choice of a small per-sample cost $c = 0.01$; for larger c we do not obtain a counterexample when using the distributions given above. More significantly, however, our non-stationary rule differs from design rule B by only the first sample. This suggests that one can achieve better cost by alternating between using design A and design B on the odd and even samples, respectively. We prove this formally in the asymptotic setting in the next section.

4.2 Asymptotic suboptimality of stationary designs

We now prove that there is always a range of prior probabilities for which stationary quantizer designs are suboptimal. Our result stems from the following observation: the form of the approximation (9) dictates that in order to achieve a small cost we need to choose a quantizer ϕ for which the KL divergences $\mu^0_\phi := D(f^0_\phi||f^1_\phi)$ and $\mu^1_\phi := D(f^1_\phi||f^0_\phi)$ are both as large as possible. Due to the asymmetry of the KL divergence, however, these maxima are not necessarily achieved by a single quantizer ϕ. This suggests that it should be possible to construct a non-stationary quantizer with better cost than a stationary design.

Proposition 2. *Let ϕ_1 and ϕ_2 be any two quantizers. If the following inequalities hold*

$$\mu^0_{\phi_1} < \mu^0_{\phi_2} \text{ and } \mu^1_{\phi_1} > \mu^1_{\phi_2} \tag{15}$$

then there exists a non-empty interval $(A, B) \subseteq (0, +\infty)$ such that as $c \to 0$,

$$\begin{aligned} J^*_{\phi_1} \le J^*_{\phi_1,\phi_2} \le J^*_{\phi_2} \quad &\text{if} \quad \frac{\pi^0}{\pi^1} \le A \\ J^*_{\phi_1,\phi_2} < \min\{J^*_{\phi_1}, J^*_{\phi_2}\} - \Theta(c \log c^{-1}) \quad &\text{if} \quad \frac{\pi^0}{\pi^1} \in (A, B) \\ J^*_{\phi_1} \ge J^*_{\phi_1,\phi_2} \ge J^*_{\phi_2} \quad &\text{if} \quad \frac{\pi^0}{\pi^1} \ge B, \end{aligned}$$

where $J^*_{\phi_i}$ denotes the optimal cost of the stationary design based on the quantizer ϕ_i, and $J^*_{\phi_1,\phi_2}$ denotes the optimal cost of a sequential test that alternates between using ϕ_1 and ϕ_2 on odd and even samples respectively.

Proof: According to Proposition 1, we have

$$J^*_{\phi_i} = \left(\frac{\pi^0}{\mu^0_{\phi_i}} + \frac{\pi^1}{\mu^1_{\phi_i}}\right) c \log c^{-1} + O(c), \quad i = 0, 1 \tag{16}$$

Now consider the sequential test that applies quantizers ϕ_1 and ϕ_2 alternately to odd and even samples. Furthermore, let this test consider two samples at a time. Let $f^0_{\phi_1\phi_2}$ and $f^1_{\phi_1\phi_2}$ denote the induced conditional probability distributions, jointly on the odd-even pairs of quantized variables. From the additivity of the KL divergence and assumption (15), there holds:

$$D(f^0_{\phi_1\phi_2} || f^1_{\phi_1\phi_2}) = \mu^0_{\phi_1} + \mu^0_{\phi_2} > 2\mu^0_{\phi_1} \tag{17a}$$

$$D(f^1_{\phi_1\phi_2} || f^0_{\phi_1\phi_2}) = \mu^1_{\phi_1} + \mu^1_{\phi_2} < 2\mu^1_{\phi_1}. \tag{17b}$$

Clearly, the cost of the proposed sequential test is an upper bound for $J^*_{\phi_1,\phi_2}$. Furthermore, the gap between this upper bound and the true optimal cost is no more than $O(c)$. Hence, as in the proof of Proposition 1, as $c \to 0$, the optimal cost $J^*_{\phi_1,\phi_2}$ can be written as

$$\left(\frac{2\pi^0}{\mu^0_{\phi_1} + \mu^0_{\phi_2}} + \frac{2\pi^1}{\mu^1_{\phi_1} + \mu^1_{\phi_2}}\right) c \log c^{-1} + O(c). \tag{18}$$

From equations (16) and (18), simple calculations yield the claim with

$$A = \frac{\mu^0_{\phi_1}(\mu^1_{\phi_1} - \mu^1_{\phi_2})(\mu^0_{\phi_1} + \mu^0_{\phi_2})}{\mu^1_{\phi_1}(\mu^1_{\phi_1} + \mu^1_{\phi_2})(\mu^0_{\phi_2} - \mu^0_{\phi_1})} < B = \frac{\mu^0_{\phi_2}(\mu^1_{\phi_1} - \mu^1_{\phi_2})(\mu^0_{\phi_1} + \mu^0_{\phi_2})}{\mu^1_{\phi_2}(\mu^1_{\phi_1} + \mu^1_{\phi_2})(\mu^0_{\phi_2} - \mu^0_{\phi_1})}.$$

□

Remarks: Let us return to the example provided in the previous subsection. Note that the two quantizers ϕ_A and ϕ_B satisfy assumption (15), since $D(f^0_{\phi_B}||f^1_{\phi_B}) = 0.4045 < D(f^0_{\phi_A}||f^1_{\phi_A}) = 0.45$ and $D(f^1_{\phi_B}||f^0_{\phi_B}) = 2.4337 > D(f^1_{\phi_A}||f^0_{\phi_A}) = 0.5108$. As a result, there exist priors for which a sequential test using stationary quantizer design (either ϕ_A, ϕ_B or ϕ_C for all samples) is not optimal.

5 On asymptotically optimal blockwise stationary designs

Despite the possible loss in optimality, it is useful to consider some form of stationarity in order to reduce computational complexity of the optimization and decision process. In this section, we consider the class of *blockwise stationary* designs, meaning that there exists some natural number T such that $\phi_{T+1} = \phi_1, \phi_{T+2} = \phi_2$, and so on. For each T, let C_T denote the class of all blockwise stationary designs with period T. We suppose throughout the analysis that each decision rule ϕ_n $(n = 1, \ldots, T)$ satisfies assumption (10). Thus, as T increases, we have a hierarchy of increasingly rich quantizer classes that will be seen to yield progressively better approximations to the optimal solution.

For a fixed prior (π^0, π^1) and $T > 0$, let $(\phi_1, \ldots, \phi_T)$ denote a quantizer design in C_T. The cost J^*_ϕ of asymptotically optimal sequential test using this quantizer design is

$$\left(\frac{T\pi^0}{\mu^0_{\phi_1} + \ldots + \mu^0_{\phi_T}} + \frac{T\pi^1}{\mu^1_{\phi_1} + \ldots + \mu^1_{\phi_T}}\right) c \log c^{-1} + O(c). \tag{19}$$

This formula reveals an important property of asymptotically optimal quantization—-namely, the asymptotically optimal quantizer design is the one that minimizes the multiplicative constant associated with the $c \log c^{-1}$ term. With a harmless abuse of notation, from now on we shall use J^*_ϕ to denote this constant, which becomes a function of ϕ. More precisely, J^*_ϕ is a function of the following vector of probabilities induced by the quantizer:

$$(f^0_\phi(1), \ldots, f^0_\phi(K-1), f^1_\phi(1), \ldots, f^1_\phi(K-1)).$$

We are interested in the properties of a quantization rule ϕ that minimizes J^*_ϕ.

We begin with a simple result on the structure of asymptotically optimal quantizer designs:

Proposition 3. *For a fixed prior (π^0, π^1) and $T > 0$, let $(\phi^*_1, \ldots, \phi^*_T)$ be the optimal quantizer design among those in C_T.*

(a) For any $1 \le i, j \le T$, there holds $(\mu^0_{\phi^*_i} - \mu^0_{\phi^*_j})(\mu^1_{\phi^*_i} - \mu^1_{\phi^*_j}) \le 0$.

(b) If there is a pair (i, j) such that $(\mu^0_{\phi^*_i} - \mu^0_{\phi^*_j})(\mu^1_{\phi^*_i} - \mu^1_{\phi^*_j}) < 0$ then there exists a prior π^0 so that $(\phi^*_1, \ldots, \phi^*_T)$ is not optimal in C_T.

Proof. To establish claim (a), suppose that for some pair (i, j), there holds $(\mu^0_{\phi^*_i} - \mu^0_{\phi^*_j})(\mu^1_{\phi^*_i} - \mu^1_{\phi^*_j}) > 0$. Without loss of generality, assume that $\mu^0_{\phi^*_i} > \mu^0_{\phi^*_j}$ and $\mu^1_{\phi^*_i} > \mu^1_{\phi^*_j}$). The asymptotic cost (19) can then be improved by replacing quantizer design ϕ^*_j at the periodic index j by the design ϕ^*_i. The proof of part (b) is similar to that of Proposition 2. □

It is well known [13] that optimal quantizers—*when unrestricted*—can be expressed as threshold rules based on the log likelihood ratio (LLR). Our counterexamples in the previous sections imply that the thresholds need not be stationary (i.e., the threshold may differ from sample to sample). In the remainder of this section, we addresses a partial converse to this issue: specifically, if we restrict ourselves to stationary (or blockwise stationary) quantizer designs, then there exists an optimal design consisting of LLR-based threshold rules.

In the analysis to follow, we assume that $T = 1$ so as to simplify the exposition. Our main result, stated below as Theorem 8, provides a characterization of the optimal quantizer ϕ^*_1, denoted more simply by ϕ^*. Due to the symmetry in the roles of individual quantizer functions, ϕ_n, for $n = 1, \ldots, T$, it is straightforward to show that results for $T = 1$ can be generalized to blockwise stationary quantizers $\{\phi_n,\ n = 1, \ldots, T\}$ for the case $T > 1$.

Definition 4. *The quantizer design function $\phi : \mathcal{X} \to \mathcal{U}$ is said to be a* likelihood ratio threshold rule *if there are thresholds $d_0 = -\infty < d_1 < \ldots < d_K = +\infty$, and a permutation $(u_1, \ldots, u_K)$ of $(0, 1, \ldots, K-1)$ such that for $l = 1, \ldots, K$, with $\mathbb{P}_0$-probability 1, we have:*

$$\phi(X) = u_l \text{ if } d_{l-1} \leq f^1(X)/f^0(X) \leq d_l,$$

When $f^1(X)/f^0(X) = d_{l-1}$, set $\phi(X) = u_{l-1}$ or $\phi(X) = u_l$ with $\mathbb{P}_0$-probability 1.[3]

Previous work on the extremal properties of likelihood ratio based quantizers guarantees that the Kullback-Leibler divergence is maximized by a LLR-based quantizer [12]. In our case, however, recall that the optimal cost function J^*_ϕ takes the form

$$J^*_\phi = \frac{\pi^0}{\mu^0_\phi} + \frac{\pi^1}{\mu^1_\phi}.$$

In particular, this function depends on the pair of KL divergences, μ^0_ϕ and μ^1_ϕ, which are related to one another in a nontrivial manner. Hence, establishing asymptotic optimality of LLR-based rules for this cost function does not follow from existing results, but rather requires further understanding of the interplay between these two KL divergences.

The following lemma concerns certain "unnormalized" variants of the Kullback-Leibler (KL) divergence. Given vectors $a = (a_0, a_1)$ and $b = (b_0, b_1)$, we define functions $\tilde{D}^0$ and $\tilde{D}^1$ mapping from $\mathbb{R}^4_+$ to the real line as follows:

$$\tilde{D}^0(a, b) \quad := \quad a_0 \log \frac{a_0}{a_1} + b_0 \log \frac{b_0}{b_1} \tag{20a}$$

$$\tilde{D}^1(a, b) \quad := \quad a_1 \log \frac{a_1}{a_0} + b_1 \log \frac{b_1}{b_0}. \tag{20b}$$

These functions are related to the standard (normalized) KL divergence via the relations $\tilde{D}^0(a, 1-a) \equiv D(a_0, a_1)$, and $\tilde{D}^1(a, 1-a) \equiv D(a_1, a_0)$.

Lemma 5. *For any positive scalars $a_1, b_1, c_1, a_0, b_0, c_0$ such that $\frac{a_1}{a_0} < \frac{b_1}{b_0} < \frac{c_1}{c_0}$,* at least one *of the following two conditions must hold:*

$$\tilde{D}^0(a, b+c) > \tilde{D}^0(b, c+a) \quad \textit{and} \quad \tilde{D}^1(a, b+c) > \tilde{D}^0(b, c+a), \ \textit{or} \tag{21a}$$

$$\tilde{D}^0(c, a+b) > \tilde{D}^0(b, c+a) \quad \textit{and} \quad \tilde{D}^1(c, a+b) > \tilde{D}^0(b, c+a). \tag{21b}$$

The above lemma implies that under certain conditions on the ordering of the probability ratios, one can increase *both* KL divergences by re-quantizing. This insight is used in the following lemma to establish that the optimal quantizer ϕ behaves almost like a likelihood ratio rule. To state the following result, recall that the *essential supremum* is the infimum of the set of all η such that $f(x) \leq \eta$ for $\mathbb{P}_0$-almost all x in the domain, for any measurable function f.

[3] This last requirement of the defi nition is termed the *canonical* likelihood ratio quantizer by Tsitsiklis [12]. Although one could consider performing additional randomization when there are ties, our later results (in particular, Lemma 7) establish that in this case, randomization will not further decrease the optimal cost J^*_ϕ.

Lemma 6. *If ϕ is an asymptotically optimal quantizer, then for any pairs of $(u_1, u_2) \in \mathcal{U}$, $u_1 \neq u_2$, there holds:*

$$\frac{f^1(u_1)}{f^0(u_1)} \notin \left(\text{ess} \inf_{x:\phi(x)=u_2} \frac{f^1(x)}{f^0(x)}, \text{ ess} \sup_{x:\phi(x)=u_2} \frac{f^1(x)}{f^0(x)} \right).$$

Note that a likelihood ratio rule guarantees something stronger: For $\mathbb{P}_0$-almost all x such that $\phi(x) = u_1$, $f^1(x)/f^0(x)$ takes a value either to the left or to the right, but not to both sides, of the interval specified above.

As we will show, the proof that there exists an optimal LLR-based rule turns out to reduce to the problem of showing that the optimal cost function J^*_ϕ is a *quasiconcave* function with respect the space of quantizers. (A function F is quasiconcave if and only if for any η, the level set $\{F(x) \geq \eta\}$ is a convex set). Since the minima of a quasiconcave function are generally extreme points of the function's domain [3], and the extreme points in the quantizer space are LLR-based rules [12], we deduce that there exists an optimal quantizer that is LLR-based. In Lemma 7, we present a proof of quasiconcavity for the case of binary quantizers. This result is sufficient to show that a LLR-based optimal quantizer exists.

Let $F : [0, 1]^2 \to R$ be given by

$$F(a_0, a_1) = \frac{c_0}{D(a_0, a_1) + d_0} + \frac{c_1}{D(a_1, a_0) + d_1}. \tag{22}$$

Lemma 7. *For any non-negative constants c_0, c_1, d_0 and d_1, the function F defined in equation* (22) *is quasiconcave.*

The following theorem is the main result of this section.

Theorem 8. *Restricting to the class of (blockwise) stationary and deterministic decision rules, there exists an asymptotically optimal quantizer ϕ that is a likelihood ratio rule.*

We present the full proof of this theorem in Appendix E, which follows from Lemma 6 and Lemma 7. We remark that our theorem is restricted to *deterministic* quantizer designs. We conjecture that J^*_ϕ is quasiconcave with respect to space of quantizers in general, which implies that Theorem 8 also holds in the (larger) space of randomized quantizer rules. As noted above, Lemma 7 establishes quasiconcavity of J^*_ϕ for binary quantizer rules, which implies that our conjecture certainly holds for the case of binary quantizers.

6 Discussion

The problem of decentralized sequential detection encompasses a wide range of problems involving different assumptions about the amount memory available at the local sensors and the nature of the feedback from the central decision-maker to the local sensors. A taxonomy has been provided by Veeravalli et al. [15]. Their analysis focused on Case E, the setting of a system with full feedback and memory restricted to past decisions. In this setting, the local sensors and the central decision-maker possess the same information state, and the problem can be attacked using dynamic programming and other tools of classical sequential analysis. This mathematical tractability is obtained, however, at a cost of realism. In many applications of the decentralized sequential detection it will not be feasible to feed back all of the decisions from all of the local sensors; indeed, in applications such as sensor networks there may be no feedback at all. Moreover, the local storage capacity may be very limited. In this paper we have focused on this more impoverished case, assuming that neither feedback nor local memory are available (Case A in the taxonomy of Veeravalli et al.).

We have provided an asymptotic characterization of the cost of the optimal sequential test in the setting of Case A. This characterization has allowed us to resolve the open question as to whether optimal quantizers are stationary. In particular, we have provided an explicit counterexample to the stationary conjecture. Moreover, we have shown that in the asymptotic setting (i.e., when the cost per sample goes to zero) we are guaranteed a range of prior probabilities for which stationary strategies are suboptimal. We have also presented a new result concerning the quasiconcavity of the optimal cost function. This result has allowed us to establish that asymptotically optimal quantizers are likelihood-based threshold rules when restricted to the class of blockwise stationary quantizers.

Acknowledgements

The research described in this paper was supported in part by Grant 0412995 from the National Science Foundation, by a grant from Microsoft Research, and by a Sloan Fellowship to Martin J. Wainwright.

A Dynamic-programming characterization

In this appendix, we describe how the optimal solution of the sequential decision problem can be characterized recursively using dynamic programming (DP) arguments [1, 17]. We assume that $X_1, X_2, \ldots$ are independent but not identically distributed conditioned on H. We use subscript n in $f_n^0(x)$ and $f_n^1(x)$ to denote the probability mass (or density) function conditioned on $H = 0$ and $H = 1$, respectively. It has been shown that the sufficient statistic for the DP analysis is the posterior probability $p_n = P(H = 1|X_1, \ldots, X_n)$, which can be updated as by:

$$p_0 = \pi^1; p_{n+1} = \frac{p_n f_{n+1}^1(X_{n+1})}{p_n f_{n+1}^1(X_{n+1}) + (1 - p_n) f_{n+1}^0(X_{n+1})}.$$

Finite horizon: First, let us restrict the stopping time N to a finite interval $[0, T]$ for some T. At each time step n, define $J_n^T(p_n)$ to be the minimum expected cost-to-go. At $n = T$, it is easily seen that

$$J_T^T(p_T) = g(p_T),$$

where $g(p) := \min\{p, 1-p\}$. In addition, the optimal decision function γ at time step T, which is a function of p_T, has the following form: $\gamma_T(p_T) = 1$ if $p \geq 1/2$ and 0 otherwise.

For $0 \leq n \leq T - 1$, a standard DP argument gives the following backward recursion:

$$J_n^T(p_n) = \min\{g(p_n), c + A_n^T(p_n)\},$$

where

$$A_n^T(p_n) = \mathbb{E}\{J_{n+1}^T(p_{n+1})|X_1, \ldots, X_n\} = \sum_{x_{n+1}} J_{n+1}^T(p_{n+1})(p_n f_{n+1}^1(x_{n+1}) + (1 - p_n) f_{n+1}^0(x_{n+1})).$$

The decision whether to stop depends on p_n: If $g(p_n) \leq c + A_n^T(p_n)$, there is no additional benefit of making one more observation, thus we stop. The final decision $\gamma(p_n)$ takes value 1 if $p_n \geq 1/2$ and 0 otherwise. The overall optimal cost function for the sequential test just described is J_0^T.

It is known that the functions J_n^T and A_n^T are concave and continuous in p that take value 0 when $p = 0$ and $p = 1$ [1]. Furthermore, the optimal region for which we decide $\hat{H} = 1$ is a convex set that contains

$p_n = 1$, and the optimal region for which we decide $\hat{H} = 0$ is a convex set that contains $p_n = 0$. Hence, we stop as soon as either $p_n \leq p_n^+$ or $p_n \geq p_n^-$ for some $0 < p_n^+ < p_n^-$. This corresponds to a likelihood ratio test: For some threshold $a_n < 0 < b_n$, let:

$$N = \inf\{n \geq 1 \,|L_n := \sum_{i=1}^{n} \log \frac{f_i^1(X_i)}{f_i^0(X_i)} \leq a_n \text{ or } L_n \geq b_n\}. \tag{23}$$

Set $\gamma(L_N) = 1$ if $L_n \geq b_n$ and 0 otherwise.

Infinite horizon: The original problem is solved by relaxing the restriction that the stopping time is bounded by a constant T. Letting $T \to \infty$, for each n, the optimal expected cost-to-go $J_n^T(p_n)$ decreases and tends to a limit denoted by $J(p_n) := \lim_{T\to\infty} J_n(p_n)$.

Note that since $X_1, X_2, \ldots$ are i.i.d. conditionally on a hypothesis H, the two functions $J_n^T(p)$ and $J_{n+1}^{T+1}(p)$ are equivalent. As a result, by lettting $T \to \infty$, $J_n(p)$ independent of n and can be denoted as $J(p)$. A similar time-shift argument also yields that the cost function $\lim_{T\to\infty} A_n^T(p)$ is independent of n. We denote this limit by $A(p)$. It is then easily seen that the optimal stopping time N is a likelihood ratio test where the thresholds a_n and b_n are independent of n. We use a to denote the former and b the latter. The functions $J(p)$ and $A(p)$ are related by the following Bellman equation [2]:

$$J(p) = \min\{g(p), c + A(p)\} \text{ for all } p \in [0, 1]. \tag{24}$$

The cost of the optimal sequential test of the problem is $J(\pi_1)$.

B Proof of Lemma 5

By renormalizing, we can assume w.l.o.g. that $a_1 + b_1 + c_1 = a_0 + b_0 + c_0 = 1$. Also w.l.o.g, assume that $b_1 \geq b_0$. Thus, $c_1 > c_0$ and $a_1 < a_0$. Replacing $c_1 = 1 - a_1 - b_1$ and $c_0 = 1 - a_0 - b_0$, the inequality $c_1/c_0 > b_1/b_0$ is equivalent to $a_1 < a_0 b_1/b_0 - (b_1 - b_0)/b_0$.

We fix values of b, and consider varying $a \in A$, where A denotes the domain for (a_0, a_1) governed by the following equality and inequality constraints:

$$0 < a_1 < 1 - b_1 \tag{25a}$$
$$0 < a_0 < 1 - b_0 \tag{25b}$$
$$a_1 < a_0 \tag{25c}$$
$$a_1 < a_0 b_1/b_0 - (b_1 - b_0)/b_0. \tag{25d}$$

Note that the third constraint is redundant due to the other three constraints. In particular, constraint (25d) corresponds to a line passing through $((b_1 - b_0)/b_1, 0)$ and $(1 - b_0, 1 - b_1)$ in the (a_0, a_1) coordinates. As a result, A is the interior of the triangle defined by this line and two other lines given by $a_1 = 0$ and $a_0 = 1 - b_0$ (see Figure B).

It is straightforward to check that both $\tilde{D}^0(a, 1-a)$ and $\tilde{D}^1(a, 1-a)$ are convex functions with respect to (a_0, a_1). In addition, the derivatives with respect to a_1 are $\frac{a_1 - a_0}{a_1(1-a_1)} < 0$ and $\log \frac{a_1(1-a_0)}{a_0(1-a_1)} < 0$, respectively. Hence, both functions can be (strictly) bounded from below by increasing a_1 while keeping a_0 unchanged, i.e., by replacing a_1 by a_1' so that (a_0, a_1') lies on the line given by (25d), which is equivalent to the constraint $c_1/c_0 = b_1/b_0$. Let $c_1' = 1 - b_1 - a_1'$, then $c_1'/c_0 = b_1/b_0$.

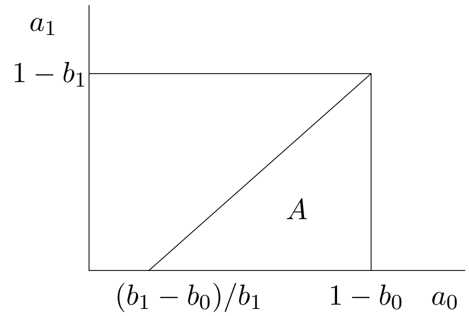

Figure 1: Illustration of the domain A.

We have

$$\begin{aligned}
\tilde{D}^0(a, b+c) &\overset{(a)}{>} a_1' \log \frac{a_1'}{a_0} + (b_1 + c_1') \log \frac{b_1 + c_1'}{b_0 + c_0} && \text{(26a)}\\
&\overset{(b)}{=} a_1' \log \frac{a_1'}{a_0} + c_1' \log \frac{c_1'}{c_0} + b_1 \log \frac{b_1}{b_0} && \text{(26b)}\\
&\overset{(c)}{\geq} (a_1' + c_1') \log \frac{a_1' + c_1'}{a_0 + c_0} + b_1 \log \frac{b_1}{b_0} && \text{(26c)}\\
&= \tilde{D}^0(a + c, b), && \text{(26d)}
\end{aligned}$$

where inequality (c) follows from an application of the log-sum inequality [4]. A similar conclusion holds for $\tilde{D}^1(a, b+c)$ as well.

C Proof of Lemma 6

Suppose the opposite is true, that there exist two sets S_1, S_2 with positive $\mathbb{P}_0$-measure such that $\phi(X) = u_2$ for any $X \in S_1 \cup S_2$, and

$$\frac{f^1(S_1)}{f^0(S_1)} < \frac{f^1(u_1)}{f^0(u_1)} < \frac{f^1(S_2)}{f^0(S_2)}. \tag{27}$$

By reassigning S_1 or S_2 to the quantile u_1, we are guaranteed to have a new quantizer ϕ' such that $\mu^0_{\phi'} > \mu^0_{\phi^*}$ and $\mu^1_{\phi'} > \mu^1_{\phi^*}$, thanks to Lemma 5. As a result, ϕ' has a smaller sequential cost $J^*_{\phi'}$, which is a contradiction.

D Proof of Lemma 7

The proof of this lemma is conceptually straightforward, but the algebra is involved. To simplify the notation, we replace a_0 by x, a_1 by y, the function $D(a_0, a_1)$ by $f(x, y)$, and the function $D(a_1, a_0)$ by $g(x, y)$. Finally, we assume that $d_0 = d_1 = 0$; the proof will reveal that this case is sufficient to establish the more general result with arbitrary non-negative scalars d_0 and d_1.

We have $f(x, y) = x \log(x/y) + (1 - x) \log(1 - x/1 - y)$ and $g(x, y) = y \log(y/x) + (1 - y) \log(1 - y/1 - x)$. Note that both f and g are convex functions and are non-negative in their domains, and moreover that we have $F(x, y) = c_0/f(x, y) + c_1/g(x, y)$. In order to establish the quasiconcavity of F, it suffices to show that for any (x, y) in the domain of F, whenever vector $h = [h_0\ h_1] \in \mathbb{R}^2$ such that $h^T \nabla F(x, y) = 0$, there holds

$$h^T \nabla^2 F(x, y)\ h \leq 0. \tag{28}$$

Here we adopt the standard notation of ∇F for the gradient vector of F, and $\nabla^2 F$ for its Hessian matrix. We also use F_x to denote the partial derivative with respect to variable x, F_{xy} to denote the partial derivative with respect to x and y, and so on.

We have $\nabla F = -\frac{c_0 \nabla f}{f^2} - \frac{c_1 \nabla g}{g^2}$. Thus, it suffices to prove relation (28) for vectors of the form

$$h = \left[\left(-\frac{c_0 f_y}{f^2} - \frac{c_1 g_y}{g^2} \right) \quad \left(\frac{c_0 f_x}{f^2} + \frac{c_1 g_x}{g^2} \right) \right]^T.$$

It is convenient to write $h = c_0 v_0 + c_1 v_1$, where $v_0 = [-f_y/f^2 \quad f_x/f^2]^T$ and $v_1 = [-g_y/g^2 \quad g_x/g^2]^T$.

The Hessian matrix $\nabla^2 F$ can be written as $\nabla^2 F = c_0 H_0 + c_0 H_1$, where

$$H_0 = -\frac{1}{f^3} \begin{bmatrix} f_{xx} f - 2 f_x^2 & f_{xy} f - 2 f_x f_y \\ f_{xy} f - 2 f_x f_y & f_{yy} f - 2 f_y^2 \end{bmatrix},$$

and

$$H_1 = -\frac{1}{g^3} \begin{bmatrix} g_{xx} g - 2 g_x^2 & g_{xy} g - 2 g_x g_y \\ g_{xy} g - 2 g_x g_y & g_{yy} g - 2 g_y^2 \end{bmatrix}.$$

Now observe that

$$h^T \nabla^2 F h \;=\; (c_0 v_0 + c_1 v_1)^T (c_0 H_0 + c_1 H_1)(c_0 v_0 + c_1 v_1),$$

which can be simplified to

$$h^T \nabla^2 F h \;=\; c_0^3 v_0^T H_0 v_0 \,+\, c_1^3 v_1^T H_1 v_1 \,+\, c_0^2 c_1 (2 v_0^T H_0 v_1 \,+\, v_0^T H_1 v_0) \,+\, c_0 c_1^2 (2 v_0^T H_1 v_1 \,+\, v_1^T H_0 v_1).$$

This function is a polynomial in c_0 and c_1, which are restricted to be non-negative scalars. Therefore, it suffices to prove that all the coefficients of this polynomial (with respect to c_0 and c_1) are non-positive. In particular, we shall show that

(i) $v_0^T H_0 v_0 \leq 0$, and

(ii) $2 v_0^T H_0 v_1 + v_0^T H_1 v_0 \leq 0$.

The non-positivity of the other two coefficients follows from entirely analogous arguments.

First, some straightforward algebra shows that inequality (i) is equivalent to the relation

$$f_{xx} f_y^2 + f_{yy} f_x^2 \geq 2 f_x f_y f_{xy}.$$

But note that f is a convex function, so $f_{xx} f_{yy} \geq f_{xy}^2$. Hence, we have

$$f_{xx} f_y^2 + f_{yy} f_x^2 \;\overset{(a)}{\geq}\; 2\sqrt{f_{xx} f_{yy}} f_x f_y \;\overset{(b)}{\geq}\; 2 f_x f_y f_{xy},$$

thereby proving (i). (In this argument, inequality (a) follows from the fact that $a^2 + b^2 \geq 2ab$, whereas inequality (b) follows from the convexity of f.)

Regarding (ii), some further algebra reduces it to the inequality

$$G_1 + G_2 - G_3 \geq 0, \tag{29}$$

where

$$\begin{aligned} G_1 &= 2(f_y g_y f_{xx} + f_x g_x f_{yy} - (f_y g_x + f_x g_y) f_{xy}), \\ G_2 &= f_y^2 g_{xx} + f_x^2 g_{yy} - 2 f_x f_y g_{xy}, \\ G_3 &= \frac{2}{g}(f_y g_x - f_x g_y)^2. \end{aligned}$$

At this point in the proof, we need to exploit specific information about the functions f and g, which are defined in terms of KL divergences. To simplify notation, we let $u = x/y$ and $v = (1-x)/(1-y)$. Computing derivatives, we have

$$\begin{aligned} f_x(x,y) &= \log(x/y) - \log((1-x)/(1-y)) = \log(u/v), \\ f_y(x,y) &= (1-x)/(1-y) - x/y = v - u, \\ g_x(x,y) &= (1-y)/(1-x) - y/x = 1/v - 1/u, \\ g_y(x,y) &= \log(y/x) - \log((1-y)/(1-x)) = \log(v/u), \\ \nabla^2 f(x,y) &= \begin{bmatrix} \frac{1}{x(1-x)} & -\frac{1}{y(1-y)} \\ -\frac{1}{x(1-x)} & \frac{1-x}{(1-y)^2} + \frac{x}{y^2} \end{bmatrix}, \qquad \text{and} \qquad \nabla^2 g(x,y) = \begin{bmatrix} \frac{1-y}{(1-x)^2} + \frac{y}{x^2} & -\frac{1}{x(1-x)} \\ -\frac{1}{x(1-x)} & \frac{1}{y(1-y)} \end{bmatrix}. \end{aligned}$$

Noting that $f_x = -g_y;\ g_{xy} = -f_{xx};\ f_{xy} = -g_{yy}$, we see that equation (29) is equivalent to

$$2(f_x g_x f_{yy} + f_y g_x g_{yy}) - f_x^2 g_{yy} + f_y^2 g_{xx} \;\; \geq \;\; \frac{2}{g}(f_y g_x - f_x g_y)^2. \tag{30}$$

To simplify the algebra further, we shall make use of the inequality $(\log t^2)^2 \leq (t - 1/t)^2$, which is valid for any t. This implies that

$$f_y g_x = (v-u)(1/v - 1/u) \leq f_x g_y = -(\log(u/v))^2 = -f_x^2 = -g_y^2 \leq 0.$$

Thus, $-f_x^2 g_{yy} \geq f_y g_x g_{yy}$, and $\frac{2}{g}(f_y g_x - f_x g_y)^2 \leq \frac{2}{g} f_y g_x (f_y g_x - f_x g_y)$. As a result, (30) would follow if we can show that

$$2(f_x g_x f_{yy} + f_y g_x g_{yy}) + f_y g_x g_{yy} + f_y^2 g_{xx} \geq \frac{2}{g} f_y g_x (f_y g_x - f_x g_y).$$

For all $x \neq y$, we may divide both sides by $-f_y(x,y) g_x(x,y) > 0$. Consequently, it suffices to show that:

$$-2 f_x f_{yy}/f_y - f_y g_{xx}/g_x - 3 g_{yy} \geq \frac{2}{g}(f_x g_y - g_x f_y),$$

or, equivalently,

$$2\log(u/v)\left(\frac{v}{u-1} + \frac{u}{1-v}\right) + \left(\frac{u}{1-x} + \frac{v}{x}\right) - \frac{3}{y(1-y)} \geq \frac{2}{g}\left(\frac{(u-v)^2}{uv} - (\log\frac{u}{v})^2\right),$$

or, equivalently,

$$2\log(u/v)\frac{(u-v)(u+v-1)}{(u-1)(1-v)}+\frac{(u-v)^2(u+v-4uv)}{uv(u-1)(1-v)}\geq\frac{2}{g}\left(\frac{(u-v)^2}{uv}-(\log\frac{u}{v})^2\right). \qquad (31)$$

Due to the symmetry, it suffices to prove (31) for $x<y$. In particular, we shall use the following inequality for logarithm mean [8], which holds for $u\neq v$:

$$\frac{3}{2\sqrt{uv}+(u+v)/2}<\frac{\log u-\log v}{u-v}<\frac{1}{(uv(u+v)/2)^{1/3}}.$$

We shall replace $\frac{\log(u/v)}{u-v}$ in (31) by appropriate upper and lower bounds. In addition, we shall also bound $g(x,y)$ from below, using the following argument. When $x<y$, we have $u<1<v$, and

$$\begin{aligned} g(x,y) &= y\log\frac{y}{x}+(1-y)\log\frac{1-y}{1-x}>\frac{3y(y-x)}{2\sqrt{xy}+(x+y)/2}+\frac{(1-y)(x-y)}{[(1-x)(1-y)(1-(x+y)/2)]^{1/3}} \\ &= \frac{3(1-v)(1-u)}{(u-v)(2\sqrt{u}+\frac{u+1}{2})}+\frac{(u-1)(1-v)}{(u-v)(v(v+1)/2)^{1/3}}>0. \end{aligned}$$

Let us denote this lower bound by $q(u,v)$.

Having got rid of the logarithm terms, (31) will hold if we can prove the following:

$$\frac{6(u-v)^2(u+v-1)}{(2\sqrt{uv}+(u+v)/2)(u-1)(1-v)}+\frac{(u-v)^2(u+v-4uv)}{uv(u-1)(1-v)}\geq\frac{2}{q(u,v)}\left(\frac{(u-v)^2}{uv}-\frac{9(u-v)^2}{(2\sqrt{uv}+(u+v)/2)^2}\right),$$

or equivalently,

$$\begin{aligned} &\left(\frac{6(u+v-1)}{(2\sqrt{uv}+(u+v)/2)}+\frac{(u+v-4uv)}{uv}\right)\left(\frac{3}{(v-u)(2\sqrt{u}+\frac{u+1}{2})}-\frac{1}{(v-u)(v(v+1)/2)^{1/3}}\right) \\ &\qquad\geq 2\left(\frac{1}{uv}-\frac{9}{(2\sqrt{uv}+(u+v)/2)^2}\right), \qquad (32) \end{aligned}$$

which is equivalent to

$$\begin{aligned} &\frac{(u+v-2\sqrt{uv})((u+v)/2+3\sqrt{uv}+4uv)}{(2\sqrt{uv}+(u+v)/2)uv}\frac{3(v(v+1)/2)^{1/3}-(2\sqrt{u}+(u+1)/2)}{(v-u)(2\sqrt{u}+(u+1)/2)(v(v+1)/2)^{1/3}} \\ &\qquad\geq\frac{(u+v-2\sqrt{uv})((u+v)/2+5\sqrt{uv})}{uv(2\sqrt{uv}+(u+v)/2)^2} \qquad (33) \end{aligned}$$

and also equivalent to

$$\begin{aligned} &((u+v)/2+2\sqrt{uv})((u+v)/2+3\sqrt{uv}+4uv)[3(v(v+1)/2)^{1/3}-(2\sqrt{u}+(u+1)/2)] \\ &\qquad\geq(2\sqrt{u}+(u+1)/2)(v(v+1)/2)^{1/3}((u+v)/2+5\sqrt{uv})(v-u) \qquad (34) \end{aligned}$$

It can be checked by tedious but straightforward calculus that inequality (34) holds for any $u\leq 1\leq v$, and equality holds when $u=1=v$, i.e., $x=y$.

E Proof of Theorem 8

Suppose that ϕ is not a likelihood ratio rule. Then there exist positive $\mathbb{P}_0$-probability disjoint sets S_1, S_2, S_3 such that for any $X_1 \in S_1, X_2 \in S_2, X_3 \in S_3$,

$$\phi(X_1) = \phi(X_3) = u_1 \tag{35a}$$
$$\phi(X_2) = u_2 \neq u_1 \tag{35b}$$
$$\frac{f^1(X_1)}{f^0(X_1)} < \frac{f^1(X_2)}{f^0(X_2)} < \frac{f^1(X_3)}{f^0(X_3)}. \tag{35c}$$

Define the probability of the quantiles as:

$$f^0(u_1) := \mathbb{P}_0(\phi(X) = u_1), \quad \text{and} \quad f^0(u_2) := \mathbb{P}_0(\phi(X) = u_2),$$
$$f^1(u_1) := \mathbb{P}_1(\phi(X) = u_1), \quad \text{and} \quad f^1(u_2) := \mathbb{P}_1(\phi(X) = u_2).$$

Similarly, for the sets S_1, S_2 and S_3, we define

$$a_0 = f^0(S_1), \quad b_0 = f^0(S_2) \quad \text{and} \quad c_0 = f^0(S_3),$$
$$a_1 = f^1(S_1), \quad b_1 = f^1(S_2), \quad \text{and} \quad c_1 = f^1(S_3).$$

Finally, let p_0, p_1, q_0 and q_1 denote the probability measures of the “residuals”:

$$p_0 \;=\; f^0(u_2) - b_0, \qquad p_1 \;=\; f^1(u_2) - b_1,$$
$$q_0 \;=\; f^0(u_1) - a_0 - c_0, \qquad q_1 \;=\; f^1(u_1) - a_1 - c_1.$$

Note that we have $\frac{a_1}{a_0} < \frac{b_1}{b_0} < \frac{c_1}{c_0}$. In addition, the sets S_1 and S_3 were chosen so that $\frac{a_1}{a_0} \leq \frac{q_1}{q_0} \leq \frac{c_1}{c_0}$. From Lemma 6, there holds $\frac{p_1+b_1}{p_0+b_0} = \frac{f^1(u_2)}{f^0(u_2)} \notin \left(\frac{a_1}{a_0}, \frac{c_1}{c_0}\right)$. We may assume without loss of generality that $\frac{p_1+b_1}{p_0+b_0} \leq \frac{a_1}{a_0}$. Then, $\frac{p_1+b_1}{p_0+b_0} < \frac{b_1}{b_0}$, so $\frac{p_1}{p_0} < \frac{p_1+b_1}{p_0+b_0}$. Overall, we are guaranteed to have the ordering

$$\frac{p_1}{p_0} < \frac{p_1 + b_1}{p_0 + b_0} \leq \frac{a_1}{a_0} < \frac{b_1}{b_0} < \frac{c_1}{c_0}. \tag{36}$$

Our strategy will be to modify the quantizer ϕ only for those X for which $\phi(X)$ takes the values u_1 or u_2, such that the resulting quantizer is defined by a LLR-based threshold, and has a smaller (or equal) value of the corresponding cost J^*_ϕ. For simplicity in notation, we use $\mathcal{A}$ to denote the set with measures under $\mathbb{P}_0$ and $\mathbb{P}_1$ equal to a_0 and a_1; the sets $\mathcal{B}$, $\mathcal{C}$, $\mathcal{P}$ and $\mathcal{Q}$ are defined in an analogous manner. We begin by observing that we have either $\frac{a_1}{a_0} \leq \frac{q_1+a_1}{q_0+a_0} < \frac{b_1}{b_0}$ or $\frac{b_1}{b_0} < \frac{q_1+c_1}{q_0+c_0} \leq \frac{c_1}{c_0}$. Thus, in our subsequent manipulation of sets, we always bundle $\mathcal{Q}$ with either $\mathcal{A}$ or $\mathcal{C}$ accordingly without changing the ordering of the probability ratios. Without loss of generality, then, we may disregard the corresponding residual set corresponding to $\mathcal{Q}$ in the analysis to follow.

In the remainder of the proof, we shall show that either one of the following two modifications of the quantizer ϕ will improve (decrease) the sequential cost J^*_ϕ:

(i) Assign $\mathcal{A}, \mathcal{B}$ and $\mathcal{C}$ to the same quantization level u_1, and leave $\mathcal{P}$ to the level u_2, or

(ii) Assign $\mathcal{P}$, $\mathcal{A}$ and $\mathcal{B}$ to the same level u_2, and leave c to the level u_1.

It is clear that this modified quantizer design respects the likelihood ratio rule for the quantization indices xu_1 and u_2. By repeated application of this modification for every such pair, we are guaranteed to arrive at a likelihood ratio quantizer that is optimal, thereby completing the proof.

Let a'_0, b'_0, c'_0, p'_0 be normalized versions of a_0, b_0, c_0, p_0, respectively (i.e., $a'_0 = a_0/(p_0 + a_0 + b_0 + c_0)$, and so on). Similarly, let a'_1, b'_1, c'_1, p'_1 be normalized versions of a_1, b_1, c_1, p_1, respectively. With this notation, we have the relations

$$
\begin{aligned}
\mu^0_\phi &= \sum_{u \neq u_1, u_2} f^0(u) \log \frac{f^0(u)}{f^1(u)} + (p_0 + b_0) \log \frac{p_0 + b_0}{p_1 + b_1} + (a_0 + c_0) \log \frac{a_0 + c_0}{a_1 + c_1} \\
&= A_0 + (f^0(u_1) + f^0(u_2)) \left((p'_0 + b'_0) \log \frac{p'_0 + b'_0}{p'_1 + b'_1} + (a'_0 + c'_0) \log \frac{a'_0 + c'_0}{a'_1 + c'_1} \right) \\
&= A_0 + (f^0(u_1) + f^0(u_2)) \tilde{D}^0(p' + b', a' + c'), \\
\mu^1_\phi &= \sum_{u \neq u_1, u_2} f^1(u) \log \frac{f^1(u)}{f^0(u)} + (p_1 + b_1) \log \frac{p_1 + b_1}{p_0 + b_0} + (a_1 + c_1) \log \frac{a_1 + c_1}{a_0 + c_0} \\
&= A_1 + (f^1(u_1) + f^1(u_2)) \tilde{D}^1(p' + b', a' + c'),
\end{aligned}
$$

where we define

$$
\begin{aligned}
A_0 &:= \sum_{u \neq u_1, u_2} f^0(u) \log \frac{f^0(u)}{f^1(u)} + (f^0(u_1) + f^0(u_2)) \log \frac{f^0(u_1) + f^0(u_2)}{f^1(u_1) + f^1(u_2)} \geq 0, \\
A_1 &:= \sum_{u \neq u_1, u_2} f^1(u) \log \frac{f^1(u)}{f^0(u)} + (f^1(u_1) + f^1(u_2)) \log \frac{f^1(u_1) + f^1(u_2)}{f^0(u_1) + f^0(u_2)} \geq 0
\end{aligned}
$$

due to the non-negativity of the KL divergences.

Note that from (36) we have

$$
\frac{p'_1}{p'_0} < \frac{p'_1 + b'_1}{p'_0 + b'_0} \leq \frac{a'_1}{a'_0} < \frac{b'_1}{b'_0} < \frac{c'_1}{c'_0},
$$

in addition to the normalization constraints that $p'_0 + a'_0 + b'_0 + c'_0 = p'_1 + a'_1 + b'_1 + c'_1 = 1$. It follows that $\frac{p'_1 + b'_1}{p'_0 + b'_0} < \frac{p'_1 + a'_1 + b'_1 + c'_1}{p'_0 + a'_0 + b'_0 + c'_0} = 1$.

Let us consider varying the values of a'_1, b'_1, while fixing all other variables and ensuring that all the above constraints hold. Then, $a'_1 + b'_1$ is constant, and both $\tilde{D}^0(p' + b', a' + c')$ and $\tilde{D}^1(p' + b', a' + c')$ increase as b_1 decreases and a_1 increases. In other words, if we define $a''_0 = a'_0$, $b''_0 = b'_0$ and a''_1 and b''_1 such that

$$
\frac{a''_1}{a'_0} = \frac{b''_1}{b'_0} = \frac{1 - p'_1 - c'_1}{1 - p'_0 - c'_0},
$$

then we have

$$
\tilde{D}^0(p' + b', a' + c') \leq \tilde{D}^0(p' + b'', a'' + c') \text{ and } \tilde{D}^1(p' + b', a' + c') \leq \tilde{D}^1(p' + b'', p'' + c'). \tag{37}
$$

Now note that vector (b''_0, b''_1) in $\mathbb{R}^2$ is a convex combination of $(0, 0)$ and $(a''_0 + b''_0, a''_1 + b''_1)$. It follows that $(p'_0 + b''_0, p'_1 + b''_1)$ is a convex combination of (p'_0, p'_1) and $(p'_0 + a''_0 + b''_0, p'_1 + a''_1 + b''_1) = (p'_0 + a'_0 + b'_0, p'_1 + a'_1 + b'_1)$.

By (37) and the quasiconcavity result in Lemma 7, we have:

$$
\begin{aligned}
J^*_\phi &= \frac{\pi^0}{\mu^0_\phi} + \frac{\pi^1}{\mu^1_\phi} \\
&= \frac{\pi^0}{A_0 + (f^0(u_1) + f^0(u_2))\tilde{D}^0(p' + b', a' + c')} + \frac{\pi^1}{A_1 + (f^1(u_1) + f^1(u_2))\tilde{D}^1(p' + b', a' + c')} \\
&\geq \frac{\pi^0}{A_0 + (f^0(u_1) + f^0(u_2))\tilde{D}^0(p' + b'', a'' + c')} + \frac{\pi^1}{A_1 + (f^1(u_1) + f^1(u_2))\tilde{D}^1(p' + b'', a'' + c')} \\
&= \frac{\pi^0}{A_0 + (f^0(u_1) + f^0(u_2))D(p'_0 + b''_0, p'_1 + b''_1)} + \frac{\pi^1}{A_1 + (f^1(u_1) + f^1(u_2))D(p'_1 + b''_1, p'_0 + b''_0)} \\
&\geq \min\Bigg\{ \frac{\pi^0}{A_0 + (f^0(u_1) + f^0(u_2))D(p'_0, p'_1)} + \frac{\pi^1}{A_1 + (f^1(u_1) + f^1(u_2))D(p'_1, p'_0)}, \\
&\qquad \frac{\pi^0}{A_0 + (f^0(u_1) + f^0(u_2))D(p'_0 + a'_0 + b'_0, p'_1 + a'_1 + b'_1)} + \\
&\qquad \frac{\pi^1}{A_1 + (f^1(u_1) + f^1(u_2))D(p'_1 + a'_1 + b'_1, p'_0 + a'_0 + b'_0)} \Bigg\}.
\end{aligned}
$$

But the two arguments of the minimum in the final equation are the costs of the two possible modifications of ϕ. Hence, the proof is complete.

References

[1] K. J. Arrow, D. Blackwell, and M. A. Girshick. Bayes and minimax solutions of sequential decision problems. *Econometrica*, 17(3/4):213–244, 1949.

[2] D. P. Bertsekas. *Dynamic Programming and Stochastic Control*, volume 1. Athena Scientific, Belmont, MA, 1995.

[3] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

[4] T. Cover and J. Thomas. *Elements of Information Theory*. Wiley, 1991.

[5] Y. C. Ho. Team decision problems and information structures. *Proceedings of the IEEE*, 68:644–654, 1980.

[6] T. L. Lai. Sequential analysis: Some classical problems and new challenges (with discussion). *Statist. Sinica*, 11:303–408, 2001.

[7] Y. Mei. *Asymptotically optimal methods for sequential change-point detection*. PhD thesis, California Institute of Technology, 2003.

[8] D. S. Mitrinović, J. E. Pecarić, and A. M. Fink. *Classical and New Inequalities in Analysis*. Kluwer Academic Publishers, 1993.

[9] H. V. Poor. *An Introduction to Signal Detection and Estimation*. Springer-Verlag, New York, NY, 1994.

[10] A. N. Shiryayev. *Optimal Stopping Rules*. Springer-Verlag, 1978.

[11] D. Siegmund. *Sequential Analysis*. Springer-Verlag, 1985.

[12] J. Tsitsiklis. Extremal properties of likelihood-ratio quantizers. *IEEE Trans. on Communication*, 41(4):550–558, 1993.

[13] J. N. Tsitsiklis. On threshold rules in decentralized detection. In *Proc. 25th IEEE Conf. Decision Control*, pages 232–236, 1986.

[14] V. V. Veeravalli. Sequential decision fusion: theory and applications. *Journal of the Franklin Institute*, 336:301–322, 1999.

[15] V. V. Veeravalli, T. Basar, and H. V. Poor. Decentralized sequential detection with a fusion center performing the sequential test. *IEEE Trans. Info. Theory*, 39(2):433–442, 1993.

[16] A. Wald. *Sequential Analysis*. John Wiley and Sons, Inc., New York, 1947.

[17] A. Wald and J. Wolfowitz. Optimum character of the sequential probability ratio test. *Annals of Statistics*, 19:326–339, 1948.