

BAYESIAN SEQUENTIAL CHANGE DIAGNOSIS

SAVAS DAYANIK, CHRISTIAN GOULDING, AND H. VINCENT POOR

ABSTRACT. Sequential change diagnosis is the joint problem of detection and identification of a sudden and unobservable change in the distribution of a random sequence. In this problem, the common probability law of a sequence of i.i.d. random variables suddenly changes at some disorder time to one of finitely many alternatives. This disorder time marks the start of a new regime, whose fingerprint is the new law of observations. Both the disorder time and the identity of the new regime are unknown and unobservable. The objective is to detect the regime-change as soon as possible, and, at the same time, to determine its identity as accurately as possible. Prompt and correct diagnosis is crucial for quick execution of the most appropriate measures in response to the new regime, as in fault detection and isolation in industrial processes, and target detection and identification in national defense. The problem is formulated in a Bayesian framework. An optimal sequential decision strategy is found, and an accurate numerical scheme is described for its implementation. Geometrical properties of the optimal strategy are illustrated via numerical examples. The traditional problems of Bayesian change-detection and Bayesian sequential multi-hypothesis testing are solved as special cases. In addition, a solution is obtained for the problem of detection and identification of component failure(s) in a system with suspended animation.

1. INTRODUCTION

Sequential change diagnosis is the joint problem of detection and identification of a sudden change in the distribution of a random sequence. In this problem, one observes a sequence of i.i.d. random variables X_1, X_2, \dots , taking values in some measurable space (E, \mathcal{E}) . The common probability distribution of the X 's is initially some known probability measure \mathbb{P}_0 on (E, \mathcal{E}) , and, in the terminology of statistical process control, the system is said to be “in control.” Then, at some *unknown* and *unobservable* disorder time θ , the common probability distribution changes suddenly to another probability measure \mathbb{P}_μ for some *unknown* and *unobservable* index $\mu \in \mathcal{M} \triangleq \{1, \dots, M\}$, and the system goes “out of control.” The objective is to detect the change as quickly as possible, and, at the same time, to identify the new probability distribution as accurately as possible, so that the most suitable actions can be taken with the least delay.

Decision strategies for this problem have a wide array of applications, such as fault detection and isolation in industrial processes, target detection and identification in national defense, pattern recognition and machine learning, radar and sonar signal processing, seismology, speech and image processing, biomedical signal processing, finance, and insurance.

For example, suppose we perform a quality test on each item produced from a manufacturing process consisting of several complex processing components (labeled $1, 2, \dots, M$). As long as each processing component is operating properly, we can expect the distribution of our quality test statistic to be stationary. Now, if there occurs a sudden fault in one of the processing components, this can change the distribution of our quality test statistic depending on the processing component which caused the fault. It may be costly to continue manufacture of the items at a substandard quality level, so we must decide when to (temporarily) shut down the manufacturing process and repair the fault. However, it may also be expensive to dissect each and every processing component in order to identify the source of the failure and to fix it. So, not only do we want to detect quickly when a fault happens, but, at the same time we want also to identify accurately which processing component is the cause. The time and the cause of the fault will be distributed independently according to a geometric and a finite distribution, respectively, if each component fails independently according to some geometric distributions, which is a reasonable assumption for highly reliable components; see Section 5.5. As another example, an insurance company may monitor reported claims not only to detect a change in its risk exposure, but also to assess the nature of the change so that it can adjust its premium schedule or re-balance appropriately its portfolio of reserves to hedge against a different distribution of loss scenarios.

Sequential change diagnosis can be viewed as the fusion of two fundamental areas of sequential analysis: change detection and multi-hypothesis testing. In traditional change detection problems, $M = 1$ and there is only one change distribution, \mathbb{P}_1 ; therefore, the focus is exclusively on detecting the change time, whereas in traditional sequential multi-hypothesis testing problems, there is no change time to consider. Instead, every observation has common distribution \mathbb{P}_μ for some unknown μ , and the focus is exclusively on the inference of μ . Both change detection and sequential multi-hypothesis testing have been studied extensively. For recent reviews of these areas, we refer the reader to Basseville and Nikiforov [3], Dragalin, Tartakovsky and Veeravalli [8, 9], and Lai [14], and the references therein.

However, the sequential change diagnosis problem involves key trade-off decisions not taken into account by separately applying techniques for change detection and sequential multi-hypothesis testing. While raising an alarm as soon as the change occurs is advantageous for the change detection task, it is undesirable for the isolation task because the longer one waits to raise the alarm, the more observations one has to use for inferring the change distribution. Moreover, the unknown change time complicates the isolation task, and, as a result, adaptation of existing sequential multi-hypothesis testing algorithms is problematic.

The theory of sequential change diagnosis has not been broadly developed. Nikiforov [16] provides the first results for this problem, showing asymptotic optimality for a certain non-Bayesian approach, and Lai [13] generalizes these results through the development of information-theoretic bounds and the application of likelihood methods. In this paper, we follow a Bayesian approach to reveal a new sequential decision strategy for this problem,

which incorporates a priori knowledge regarding the distributions of the change time θ and of the change index μ . We prove that this strategy is optimal and we describe an accurate numerical scheme for its implementation.

In Section 2 we formulate precisely the problem in a Bayesian framework, and in Section 3 we show that it can be reduced to an optimal stopping of a Markov process whose state space is the standard probability simplex. In addition, we establish a simple recursive formula that captures the dynamics of the process and yields a sufficient statistic fit for online tracking.

In Section 4 we use optimal stopping theory to substantiate the optimality equation for the value function of the optimal stopping problem. Moreover, we prove that this value function is bounded, concave, and continuous on the standard probability simplex. Furthermore, we prove that the optimal decision strategy uses a finite number of observations on average and we establish some important characteristics of the associated optimal stopping/decision region. In particular, we show that the optimal stopping region of the state space for the problem consists of M non-empty, convex, closed, and bounded subsets. Also, we consider a truncated version of the problem that allows at most N observations from the sequence of random measurements. We establish an explicit bound (inversely proportional to N) for the approximation error associated with this truncated problem.

In Section 5 we show that the separate problems of change detection and sequential multi-hypothesis testing are solved as special cases of the overall joint solution. We illustrate some geometrical properties of the optimal method and demonstrate its implementation by numerical examples for the special cases $M = 2$ and $M = 3$. Specifically, we show instances in which the M convex subsets comprising the optimal stopping region are connected and instances in which they are not. Likewise, we show that the continuation region (i.e., the complement of the stopping region) need not be connected. We provide a solution to the problem of detection and identification of component failure(s) in a system with suspended animation. Finally, we outline in Section 6 how the change-diagnosis algorithm may be implemented with a computer in general. Proofs of most results are deferred to the Appendix.

2. PROBLEM STATEMENT

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space hosting random variables $\theta : \Omega \mapsto \{0, 1, \dots\}$ and $\mu : \Omega \mapsto \mathcal{M} \triangleq \{1, \dots, M\}$ and a process $X = (X_n)_{n \geq 1}$ taking values in some measurable space (E, \mathcal{E}) . Suppose that for every $t \geq 1$, $i \in \mathcal{M}$, $n \geq 1$, and $(E_k)_{k=1}^n \subseteq \mathcal{E}$

$$(2.1) \quad \mathbb{P}\{\theta = t, \mu = i, X_1 \in E_1, \dots, X_n \in E_n\} \\ = (1 - p_0)(1 - p)^{t-1} p \nu_i \prod_{1 \leq k \leq (t-1) \wedge n} \mathbb{P}_0(E_k) \prod_{t \vee 1 \leq \ell \leq n} \mathbb{P}_i(E_\ell)$$

for some given probability measures $\mathbb{P}_0, \mathbb{P}_1, \dots, \mathbb{P}_M$ on (E, \mathcal{E}) , known constants $p_0 \in [0, 1]$, $p \in (0, 1)$, and $\nu_i > 0, i \in \mathcal{M}$ such that $\nu_1 + \dots + \nu_M = 1$, where $x \wedge y \triangleq \min\{x, y\}$ and $x \vee y \triangleq \max\{x, y\}$. Namely, θ is independent of μ ; it has a zero-modified geometric

distribution with parameters p_0 and p in the terminology of Klugman, Panjer, and Willmot [12, Sec. 3.6], which reduces to the standard geometric distribution with success probability p when $p_0 = 0$. Moreover, ν_i is the probability that the change type μ is i for every $i = 1, \dots, M$.

Conditionally on θ and μ , the random variables X_n , $n \geq 1$ are independent; $X_1, \dots, X_{\theta-1}$ and $X_\theta, X_{\theta+1}, \dots$ are identically distributed with common distributions \mathbb{P}_0 and \mathbb{P}_μ , respectively. The probability measures $\mathbb{P}_0, \mathbb{P}_1, \dots, \mathbb{P}_M$ always admit densities with respect to some σ -finite measure m on (E, \mathcal{E}) ; for example, we can take $m = \mathbb{P}_0 + \mathbb{P}_1 \cdots + \mathbb{P}_M$. So, we fix m and denote the corresponding densities by f_0, f_1, \dots, f_M , respectively.

Suppose now that we observe sequentially the random variables X_n , $n \geq 1$. Their common probability density function f_0 changes at stage θ to some other probability density function f_μ , $\mu \in \mathcal{M}$. Our objective is to detect the change time θ as quickly as possible *and* isolate the change index μ as accurately as possible. More precisely, given costs associated with detection delay, false alarm, and false isolation of the change index, we seek a strategy that minimizes the expected total change detection *and* isolation cost.

In view of the fact that the observations arrive sequentially, we are interested in sequential diagnosis schemes. Specifically, let $\mathbb{F} = (\mathcal{F}_n)_{n \geq 0}$ denote the natural filtration of the observation process X , where

$$\mathcal{F}_0 = \{\emptyset, \Omega\} \quad \text{and} \quad \mathcal{F}_n = \sigma(X_1, \dots, X_n), \quad n \geq 1.$$

A *sequential decision strategy* $\delta = (\tau, d)$ is a pair consisting of a *stopping time* (or *stopping rule*) τ of the filtration \mathbb{F} and a *terminal decision rule* $d : \Omega \mapsto \mathcal{M}$ measurable with respect to the history $\mathcal{F}_\tau = \sigma(X_{n \wedge \tau}; n \geq 1)$ of observation process X through stage τ . Applying a sequential decision strategy $\delta = (\tau, d)$ consists of announcing at the end of stage τ that the common probability density function has changed from f_0 to f_d at or before stage τ . Let

$$\Delta \triangleq \{(\tau, d) \mid \tau \in \mathbb{F}, \text{ and } d \in \mathcal{F}_\tau \text{ is an } \mathcal{M}\text{-valued random variable}\}$$

denote the collection of all such sequential decision strategies (“ $\tau \in \mathbb{F}$ ” means that τ is a stopping time of filtration \mathbb{F}). Let us specify the possible losses associated with a sequential decision strategy $\delta = (\tau, d) \in \Delta$ as follows:

- (i) *Detection delay loss.* Let us denote by a fixed positive constant c the detection delay cost per period. Then the expected decision delay cost for δ is $\mathbb{E}[c(\tau - \theta)^+]$, possibly infinite, where $(x)^+ \triangleq \max\{x, 0\}$.
- (ii) *Terminal decision loss.* Here we identify two cases of isolation loss depending on whether or not the change has actually occurred at or before the stage in which we announce the isolation decision:
 - (a) *Loss due to false alarm.* Let us denote by a_{0j} the isolation cost on $\{\tau < \theta, d = j\}$ for every $j \in \mathcal{M}$. Then the expected false alarm cost for δ is $\mathbb{E}[a_{0d} \mathbf{1}_{\{\tau < \theta\}}]$.

(b) *Loss due to false isolation.* Let us denote by a_{ij} the isolation cost on the event $\{\theta \leq \tau < \infty, d = j, \mu = i\}$ for every $i, j \in \mathcal{M}$. Then the expected false isolation cost for δ is $\mathbb{E}[a_{\mu d} \mathbf{1}_{\{\theta \leq \tau < \infty\}}]$.

Here, $a_{ij}, i, j \in \mathcal{M}$ are known nonnegative constants, and $a_{ii} = 0$ for every $i \in \mathcal{M}$; i.e., no cost incurred for making a correct terminal decision.

Accordingly, for every sequential decision strategy $\delta = (\tau, d) \in \Delta$, we define a *Bayes risk function*

$$(2.2) \quad R(\delta) = c \mathbb{E}[(\tau - \theta)^+] + \mathbb{E}[a_{0d} \mathbf{1}_{\{\tau < \theta\}} + a_{\mu d} \mathbf{1}_{\{\theta \leq \tau < \infty\}}]$$

as the expected diagnosis cost: the sum of the expected detection delay cost and the expected terminal decision cost upon alarm. The problem is to find a sequential decision strategy $\delta = (\tau, d) \in \Delta$ (if it exists) with the *minimum Bayes risk*

$$(2.3) \quad R^* \triangleq \inf_{\delta \in \Delta} R(\delta).$$

3. POSTERIOR ANALYSIS AND FORMULATION AS AN OPTIMAL STOPPING PROBLEM

In this section we show that the Bayes risk function in (2.2) can be written as the expected value of the running and terminal costs driven by a certain Markov process. We use this fact to recast the minimum Bayes risk in (2.3) as a Markov optimal stopping problem.

Let us introduce the posterior probability processes

$$\Pi_n^{(0)} \triangleq \mathbb{P}\{\theta > n \mid \mathcal{F}_n\} \quad \text{and} \quad \Pi_n^{(i)} \triangleq \mathbb{P}\{\theta \leq n, \mu = i \mid \mathcal{F}_n\}, \quad i \in \mathcal{M}, n \geq 0.$$

Having observed the first n observations, $\Pi_n^{(0)}$ is the posterior probability that the change *has not* yet occurred at or before stage n , while $\Pi_n^{(i)}$ is the posterior joint probability that the change *has* occurred by stage n and that the hypothesis $\mu = i$ is correct. The connection of these posterior probabilities to the loss structure for our problem is established in the next proposition.

Proposition 3.1. *For every sequential decision strategy $\delta \in \Delta$, the Bayes risk function (2.2) can be expressed in terms of the process $\Pi \triangleq \{\Pi_n = (\Pi_n^{(0)}, \dots, \Pi_n^{(M)})\}_{n \geq 0}$ as*

$$R(\delta) = \mathbb{E} \left[\sum_{n=0}^{\tau-1} c (1 - \Pi_n^{(0)}) + \mathbf{1}_{\{\tau < \infty\}} \sum_{j=1}^M \mathbf{1}_{\{d=j\}} \sum_{i=0}^M a_{ij} \Pi_\tau^{(i)} \right].$$

While our original formulation of the Bayes risk function (2.2) was in terms of the values of the unobservable random variables θ and μ , Proposition 3.1 gives us an equivalent version of the Bayes risk function in terms of the posterior distributions for θ and μ . This is particularly effective in light of Proposition 3.2, which we state with the aid of some additional notation that is referred to throughout the paper. Let

$$S^M \triangleq \left\{ \pi = (\pi_0, \pi_1, \dots, \pi_M) \in [0, 1]^{M+1} \mid \pi_0 + \pi_1 + \dots + \pi_M = 1 \right\}$$

denote the standard M -dimensional probability simplex. Define the mappings $D_i : S^M \times E \mapsto [0, 1]$, $i \in \mathcal{M}$ and $D : S^M \times E \mapsto [0, 1]$ by

$$(3.1) \quad D_i(\pi, x) \triangleq \begin{cases} (1-p)\pi_0 f_0(x), & i = 0 \\ (\pi_i + \pi_0 p \nu_i) f_i(x), & i \in \mathcal{M} \end{cases}, \quad D(\pi, x) \triangleq \sum_{i=0}^M D_i(\pi, x),$$

and the operator \mathbb{T} on the collection of bounded functions $f : S^M \mapsto \mathbb{R}$ by

$$(3.2) \quad (\mathbb{T}f)(\pi) \triangleq \int_E m(dx) D(\pi, x) f\left(\frac{D_0(\pi, x)}{D(\pi, x)}, \dots, \frac{D_M(\pi, x)}{D(\pi, x)}\right) \text{ for every } \pi \in S^M.$$

Proposition 3.2. *The process Π possesses the following properties:*

- (a) *The process $\Pi^{(0)} \triangleq \{\Pi_n^{(0)}, \mathcal{F}_n\}_{n \geq 0}$ is a supermartingale, and $\mathbb{E} \Pi_n^{(0)} \leq (1-p)^n$ for every $n \geq 0$.*
- (b) *The process $\Pi^{(i)} \triangleq \{\Pi_n^{(i)}, \mathcal{F}_n\}_{n \geq 0}$ is a submartingale for every $i \in \mathcal{M}$.*
- (c) *The process $\Pi = \{(\Pi_n^{(0)}, \dots, \Pi_n^{(M)})\}_{n \geq 0}$ is a Markov process, and*

$$(3.3) \quad \Pi_{n+1}^{(i)} = \frac{D_i(\Pi_n, X_{n+1})}{D(\Pi_n, X_{n+1})}, \quad i \in \{0\} \cup \mathcal{M}, \quad n \geq 0,$$

with initial state $\Pi_0^{(0)} = 1 - p_0$ and $\Pi_0^{(i)} = p_0 \nu_i$, $i \in \mathcal{M}$. Moreover, for every bounded function $f : S^M \mapsto \mathbb{R}$ and $n \geq 0$, we have $\mathbb{E}[f(\Pi_{n+1}) | \Pi_n] = (\mathbb{T}f)(\Pi_n)$.

Remark 3.3. *Since $1 = \sum_{i=0}^M \Pi_n^{(i)}$, the vector $(\Pi_n^{(0)}, \dots, \Pi_n^{(M)}) \in S^M$ for every $n \geq 0$. Since Π is uniformly bounded, the limit $\lim_{n \rightarrow \infty} \Pi_n$ exists by the martingale convergence theorem. Moreover, $\lim_{n \rightarrow \infty} \Pi_n^{(0)} = 0$ a.s. by Proposition 3.2(a) since $p \in (0, 1)$.*

Now, let the functions h, h_1, \dots, h_M from S^M into \mathbb{R}_+ be defined by

$$h(\pi) \triangleq \min_{j \in \mathcal{M}} h_j(\pi) \quad \text{and} \quad h_j(\pi) \triangleq \sum_{i=0}^M \pi_i a_{ij}, \quad j \in \mathcal{M},$$

respectively. Then, we note that for every $\delta = (\tau, d) \in \Delta$, we have

$$\begin{aligned} R(\tau, d) &= \mathbb{E} \left[\sum_{n=0}^{\tau-1} c(1 - \Pi_n^{(0)}) + \mathbf{1}_{\{\tau < \infty\}} \sum_{j=1}^M \mathbf{1}_{\{d=j\}} h_j(\Pi_\tau) \right] \\ &\geq \mathbb{E} \left[\sum_{n=0}^{\tau-1} c(1 - \Pi_n^{(0)}) + \mathbf{1}_{\{\tau < \infty\}} h(\Pi_\tau) \right] = R(\tau, \tilde{d}) \end{aligned}$$

where we define on the event $\{\tau < \infty\}$ the terminal decision rule \tilde{d} to be any index satisfying $h_{\tilde{d}}(\Pi_\tau) = h(\Pi_\tau)$. In other words, an optimal terminal decision depends only upon the value of the Π process at the stage in which we stop. Note also that the functions h and h_1, \dots, h_M are bounded on S^M . Therefore, we have the following:

Lemma 3.4. *The minimum Bayes risk (2.3) reduces to the following optimal stopping of the Markov process Π :*

$$R^* = \inf_{(\tau, d) \in \Delta} R(\tau, d) = \inf_{(\tau, \tilde{d}) \in \Delta} R(\tau, \tilde{d}) = \inf_{\tau \in \mathbb{F}} \mathbb{E} \left[\sum_{n=0}^{\tau-1} c(1 - \Pi_n^{(0)}) + \mathbf{1}_{\{\tau < \infty\}} h(\Pi_\tau) \right].$$

We simplify this formulation further by showing that it is enough to take the infimum over

$$(3.4) \quad C \triangleq \{\tau \in \mathbb{F} \mid \tau < \infty \text{ a.s. and } \mathbb{E}Y_\tau^- < \infty\},$$

where

$$(3.5) \quad -Y_n \triangleq \sum_{k=0}^{n-1} c(1 - \Pi_k^{(0)}) + h(\Pi_n), \quad n \geq 0$$

is the minimum cost obtained by making the best terminal decision when alarm is set at time n . Since $h(\cdot)$ is bounded on S^M , the process $\{Y_n, \mathcal{F}_n; n \geq 0\}$ consists of integrable random variables. So the expectation $\mathbb{E}Y_\tau$ exists for every $\tau \in \mathbb{F}$, and our problem becomes

$$(3.6) \quad -R^* = \sup_{\tau \in \mathbb{F}} \mathbb{E}Y_\tau.$$

Observe that $\mathbb{E}\tau < \infty$ for every $\tau \in C$ because $\infty > (1/c)\mathbb{E}Y_\tau^- \geq \mathbb{E}(\tau - \theta)^+ \geq \mathbb{E}(\tau - \theta) \geq \mathbb{E}\tau - \mathbb{E}\theta \geq \mathbb{E}\tau - (1/p)$. In fact, we have $\mathbb{E}Y_\tau > -\infty \Leftrightarrow \mathbb{E}Y_\tau^- < \infty \Leftrightarrow \mathbb{E}\tau < \infty$ for every $\tau \in \mathbb{F}$. Since $\sup_{\tau \in \mathbb{F}} \mathbb{E}Y_\tau \geq \mathbb{E}Y_0 > -h(\Pi_0) > -\infty$, it is enough to consider $\tau \in \mathbb{F}$ such that $\mathbb{E}\tau < \infty$. Namely, (3.6) reduces to

$$(3.7) \quad -R^* = \sup_{\tau \in C} \mathbb{E}Y_\tau.$$

4. SOLUTION VIA OPTIMAL STOPPING THEORY

In this section we derive an optimal solution for the sequential change diagnosis problem in (2.3) by building on the formulation of (3.7) via the tools of optimal stopping theory.

4.1. The optimality equation. We begin by applying the method of truncation with a view of passing to the limit to arrive at the final result. For every $N \geq 0$ and $n = 0, \dots, N$, define the sub-collections

$$C_n \triangleq \{\tau \vee n \mid \tau \in C\} \quad \text{and} \quad C_n^N \triangleq \{\tau \wedge N \mid \tau \in C_n\}$$

of stopping times in C of (3.4). Note that $C = C_0$. Now, consider the families of (truncated) optimal stopping problems corresponding to $(C_n)_{n \geq 0}$ and $(C_n^N)_{0 \leq n \leq N}$, respectively, defined by

$$(4.1) \quad -V_n \triangleq \sup_{\tau \in C_n} \mathbb{E}Y_\tau, \quad n \geq 0 \quad \text{and} \quad -V_n^N \triangleq \sup_{\tau \in C_n^N} \mathbb{E}Y_\tau, \quad 0 \leq n \leq N, \quad N \geq 0.$$

Note that $R^* = V_0$.

To investigate these optimal stopping problems, we introduce versions of the *Snell envelope* of $(Y_n)_{n \geq 0}$ (i.e., the smallest regular supermartingale dominating $(Y_n)_{n \geq 0}$) corresponding to $(C_n)_{n \geq 0}$ and $(C_n^N)_{0 \leq n \leq N}$, respectively, defined by

$$(4.2) \quad \gamma_n \triangleq \operatorname{ess\,sup}_{\tau \in C_n} \mathbb{E}[Y_\tau | \mathcal{F}_n], \quad n \geq 0 \quad \text{and} \quad \gamma_n^N \triangleq \operatorname{ess\,sup}_{\tau \in C_n^N} \mathbb{E}[Y_\tau | \mathcal{F}_n], \quad 0 \leq n \leq N, \quad N \geq 0.$$

Then through the following series of lemmas, whose proofs are deferred to the Appendix, we point out several useful properties of these Snell envelopes. Finally, we extend these results to an arbitrary initial state vector and establish the optimality equation. Note that each of the ensuing (in)equalities between random variables are in the \mathbb{P} -almost sure sense.

First, these Snell envelopes provide the following alternative expressions for the optimal stopping problems introduced in (4.1) above.

Lemma 4.1. *For every $N \geq 0$ and $0 \leq n \leq N$, we have $-V_n = \mathbb{E}\gamma_n$ and $-V_n^N = \mathbb{E}\gamma_n^N$.*

Second, we have the following backward-induction equations.

Lemma 4.2. *We have $\gamma_n = \max\{Y_n, \mathbb{E}[\gamma_{n+1} | \mathcal{F}_n]\}$ for every $n \geq 0$. For every $N \geq 1$ and $0 \leq n \leq N - 1$, we have $\gamma_N^N = Y_N$ and $\gamma_n^N = \max\{Y_n, \mathbb{E}[\gamma_{n+1}^N | \mathcal{F}_n]\}$.*

We also have that these versions of the Snell envelopes coincide in the limit as $N \rightarrow \infty$. That is,

Lemma 4.3. *For every $n \geq 0$, we have $\gamma_n = \lim_{N \rightarrow \infty} \gamma_n^N$.*

Next, recall from (3.2) and Proposition 3.2(c) the operator \mathbb{T} and let us introduce the operator \mathbb{M} on the collection of bounded functions $f : S^M \mapsto \mathbb{R}_+$ defined by

$$(4.3) \quad (\mathbb{M}f)(\pi) \triangleq \min\{h(\pi), c(1 - \pi_0) + (\mathbb{T}f)(\pi)\}, \quad \pi \in S^M.$$

Observe that $0 \leq \mathbb{M}f \leq h$. That is, $\pi \mapsto (\mathbb{M}f)(\pi)$ is a nonnegative bounded function. Therefore, $\mathbb{M}^2 f \equiv \mathbb{M}(\mathbb{M}f)$ is well-defined. If f is nonnegative and bounded, then $\mathbb{M}^n f \equiv \mathbb{M}(\mathbb{M}^{n-1} f)$ is defined for every $n \geq 1$, with $\mathbb{M}^0 f \equiv f$ by definition. Using operator \mathbb{M} , we can express $(\gamma_n^N)_{0 \leq n \leq N}$ in terms of the process Π as stated in the following lemma.

Lemma 4.4. *For every $N \geq 0$, and $0 \leq n \leq N$, we have*

$$(4.4) \quad \gamma_n^N = -c \sum_{k=0}^{n-1} (1 - \Pi_k^{(0)}) - (\mathbb{M}^{N-n} h)(\Pi_n).$$

The next lemma shows how the optimal stopping problems can be rewritten in terms of the operator \mathbb{M} . It also conveys the connection between the truncated optimal stopping problems and the initial state Π_0 of the Π process.

Lemma 4.5. *We have*

$$(a) \quad V_0^N = (\mathbb{M}^N h)(\Pi_0) \text{ for every } N \geq 0, \text{ and}$$

$$(b) V_0 = \lim_{N \rightarrow \infty} (\mathbb{M}^N h)(\Pi_0).$$

Observe that since $\Pi_0 \in \mathcal{F}_0 = \{\emptyset, \Omega\}$, we have $\mathbb{P}\{\Pi_0 = \pi\} = 1$ for some $\pi \in S^M$. On the other hand, for every $\pi \in S^M$ we can construct a probability space $(\Omega, \mathcal{F}, \mathbb{P}_\pi)$ hosting a Markov process Π with the same dynamics as in (3.3) and $\mathbb{P}_\pi\{\Pi_0 = \pi\} = 1$. Moreover, on such a probability space, the preceding results remain valid. So, let us denote by \mathbb{E}_π the expectation with respect to \mathbb{P}_π and rewrite (4.1) as

$$-V_n(\pi) \triangleq \sup_{\tau \in C_n} \mathbb{E}_\pi Y_\tau, \quad n \geq 0, \quad \text{and} \quad -V_n^N(\pi) \triangleq \sup_{\tau \in C_n^N} \mathbb{E}_\pi Y_\tau, \quad 0 \leq n \leq N, \quad N \geq 0$$

for every $\pi \in S^M$. Then Lemma 4.5 implies that

$$(4.5) \quad V_0^N(\pi) = (\mathbb{M}^N h)(\pi) \text{ for every } N \geq 0, \quad \text{and} \quad V_0(\pi) = \lim_{N \rightarrow \infty} (\mathbb{M}^N h)(\pi)$$

for every $\pi \in S^M$. Taking limits as $N \rightarrow \infty$ of both sides in $(\mathbb{M}^{N+1} h)(\pi) = \mathbb{M}(\mathbb{M}^N h)(\pi)$ and applying the monotone convergence theorem on the right-hand side yields $V_0(\pi) = (\mathbb{M}V_0)(\pi)$. Hence, we have shown the following result.

Proposition 4.6 (Optimality equation). *For every $\pi \in S^M$, we have*

$$(4.6) \quad V_0(\pi) = (\mathbb{M}V_0)(\pi) \equiv \min\{h(\pi), c(1 - \pi_0) + (\mathbb{T}V_0)(\pi)\}.$$

Remark 4.7. *By solving $V_0(\pi)$ for any initial state $\pi \in S^M$, we capture the solution to the original problem since property (c) of Proposition 3.2 and (3.7) imply that*

$$R^* = V_0(1 - p_0, p_0\nu_1, \dots, p_0\nu_M).$$

4.2. Some properties of the value function. Now, we reveal some important properties of the value function $V_0(\cdot)$ of (4.5). These results help us to establish an optimal solution for $V_0(\cdot)$, and hence an optimal solution for R^* , in the next subsection.

Lemma 4.8. *If $g : S^M \mapsto \mathbb{R}$ is a bounded concave function, then so is $\mathbb{T}g$.*

Proposition 4.9. *The mappings $\pi \mapsto V_0^N(\pi)$, $N \geq 0$ and $\pi \mapsto V_0(\pi)$ are concave.*

Proposition 4.10. *For every $N \geq 1$ and $\pi \in S^M$, we have*

$$V_0(\pi) \leq V_0^N(\pi) \leq V_0(\pi) + \left(\frac{\|h\|^2}{c} + \frac{\|h\|}{p} \right) \frac{1}{N}.$$

Since $\|h\| \triangleq \sup_{\pi \in S^M} |h(\pi)| < \infty$, $\lim_{N \rightarrow \infty} \downarrow V_0^N(\pi) = V_0(\pi)$ uniformly in $\pi \in S^M$.

Proposition 4.11. *For every $N \geq 0$, the function $V_0^N : S^M \mapsto \mathbb{R}_+$ is continuous.*

Corollary 4.12. *The function $V_0 : S^M \mapsto \mathbb{R}_+$ is continuous.*

Note that S^M is a compact subset of \mathbb{R}^{M+1} , so while continuity of $V_0(\cdot)$ on the interior of S^M follows from the concavity of $V_0(\cdot)$ by Proposition 4.8, Corollary 4.12 establishes continuity on all of S^M , including its boundary.

4.3. An optimal sequential decision strategy. Finally, we describe the optimal stopping region in S^M implied by the value function $V_0(\cdot)$, and we present an optimal sequential decision strategy for our problem. Let us define for every $N \geq 0$,

$$\begin{aligned}\Gamma_N &\triangleq \{\pi \in S^M \mid V_0^N(\pi) = h(\pi)\}, & \Gamma_N^{(j)} &\triangleq \Gamma_N \cap \{\pi \in S^M \mid h(\pi) = h_j(\pi)\}, \quad j \in \mathcal{M}, \\ \Gamma &\triangleq \{\pi \in S^M \mid V_0(\pi) = h(\pi)\}, & \Gamma^{(j)} &\triangleq \Gamma \cap \{\pi \in S^M \mid h(\pi) = h_j(\pi)\}, \quad j \in \mathcal{M}.\end{aligned}$$

Theorem 4.15 below shows that it is always optimal to stop and raise an alarm as soon as the posterior probability process Π enters the region Γ . Intuitively, this follows from the optimality equation (4.6). At any stage, we always have two choices: either we stop immediately and raise an alarm or we wait for at least one more stage and take an additional observation. If the posterior probability of all possibilities is given by the vector π , then the costs of those competing actions equal $h(\pi)$ and $c(1 - \pi_0) + (\mathbb{T}V_0)(\pi)$, respectively, and it is always better to take the action that has the smaller expected cost. Stopping costs less (and therefore it is optimal) if $h(\pi) \leq c(1 - \pi_0) + (\mathbb{T}V_0)(\pi)$, equivalently, if $V_0(\pi) = h(\pi)$. Likewise, if at most N stages are left, then stopping is optimal if $V_0^N(\pi) = h(\pi)$ or $\pi \in \Gamma_N$.

For each $j \in \{0\} \cup \mathcal{M}$, let $e_j \in S^M$ denote the unit vector consisting of zero in every component except for the j th component, which is equal to one. Note that e_0, \dots, e_M are the extreme points of the closed convex set S^M , and any vector $\pi = (\pi_0, \dots, \pi_M) \in S^M$ can be expressed in terms of e_0, \dots, e_M as $\pi = \sum_{j=0}^M \pi_j e_j$.

Theorem 4.13. *For every $j \in \mathcal{M}$, $(\Gamma_N^{(j)})_{N \geq 0}$ is a decreasing sequence of non-empty, closed, convex subsets of S^M . Moreover,*

$$\begin{aligned}\Gamma_0^{(j)} &\supseteq \Gamma_1^{(j)} \supseteq \dots \supseteq \Gamma^{(j)} \supseteq \{\pi \in S^M \mid h_j(\pi) \leq \min\{h(\pi), c(1 - \pi_0)\}\} \ni e_j, \\ \Gamma &= \bigcap_{N=1}^{\infty} \Gamma_N = \bigcup_{j=1}^M \Gamma^{(j)}, \quad \text{and} \quad \Gamma^{(j)} = \bigcap_{N=1}^{\infty} \Gamma_N^{(j)}, \quad j \in \mathcal{M}.\end{aligned}$$

Furthermore, $S^M = \Gamma_0 \supseteq \Gamma_1 \supseteq \dots \supseteq \Gamma \not\supseteq \{e_1, \dots, e_M\}$.

Lemma 4.14. *For every $n \geq 0$, we have $\gamma_n = -c \sum_{k=0}^{n-1} (1 - \Pi_k^{(0)}) - V_0(\Pi_n)$.*

Theorem 4.15. *Let $\sigma \triangleq \inf\{n \geq 0 \mid \Pi_n \in \Gamma\}$.*

- (a) *The stopped process $\{\gamma_{n \wedge \sigma}, \mathcal{F}_n; n \geq 0\}$ is a martingale.*
- (b) *The random variable σ is an optimal stopping time for V_0 , and*
- (c) $\mathbb{E} \sigma < \infty$.

Therefore, the pair (σ, d^*) is an optimal sequential decision strategy for (2.3), where the optimal stopping rule σ is given by Theorem 4.15, and, as in the proof of Lemma 3.4, the optimal terminal decision rule d^* is given by

$$d^* = j \quad \text{on the event} \quad \{\sigma = n, \Pi_n \in \Gamma^{(j)}\} \quad \text{for every } n \geq 0.$$

Accordingly, the set Γ is called the *stopping region* implied by $V_0(\cdot)$, and Theorem 4.13 reveals its basic structure. We demonstrate the use of these results in the numerical examples of Section 5.

Note that we can take a similar approach to prove that the stopping rules $\sigma_N \triangleq \inf\{n \geq 0 \mid \Pi_n \in \Gamma_{N-n}\}$, $N \geq 0$ are optimal for the truncated problems $V_0^N(\cdot)$, $N \geq 0$ in (4.5). Thus, for each $N \geq 0$, the set Γ_N is called the stopping region for $V_0^N(\cdot)$: it is optimal to terminate the experiments in Γ_N if N stages are left before truncation.

5. SPECIAL CASES AND EXAMPLES

In this section we discuss solutions for various special cases of the general formulation given in Section 2. First, we show how the traditional problems of Bayesian sequential change detection and Bayesian sequential multi-hypothesis testing are formulated via the framework of Section 2. Then we present numerical examples for the cases $M = 2$ and $M = 3$. In particular, we develop a geometrical framework for working with the sufficient statistic developed in Section 3 and the optimal sequential decision strategy developed in Section 4. Finally, we solve the special problem of detection and identification of primary component failure(s) in a system with suspended animation.

5.1. A. N. Shiryaev's sequential change detection problem. Set $a_{0j} = 1$ for $j \in \mathcal{M}$ and $a_{ij} = 0$ for $i, j \in \mathcal{M}$, then the Bayes risk function (2.2) becomes

$$\begin{aligned} R(\delta) &= c \mathbb{E}[(\tau - \theta)^+] + \mathbb{E}[a_{0d} \mathbf{1}_{\{\tau < \theta\}} + a_{\mu d} \mathbf{1}_{\{\theta \leq \tau < \infty\}}] = c \mathbb{E}[(\tau - \theta)^+] + \mathbb{E}[\mathbf{1}_{\{\tau < \theta\}}] \\ &= \mathbb{P}\{\tau < \theta\} + c \mathbb{E}[(\tau - \theta)^+]. \end{aligned}$$

This is the Bayes risk studied by Shiryaev [19, 20] to solve the sequential change detection problem.

5.2. Sequential multi-hypothesis testing. Set $p_0 = 1$, then $\theta = 0$ a.s. and thus the Bayes risk function (2.2) becomes

$$R(\delta) = c \mathbb{E}[(\tau - \theta)^+] + \mathbb{E}[a_{0d} \mathbf{1}_{\{\tau < \theta\}} + a_{\mu d} \mathbf{1}_{\{\theta \leq \tau < \infty\}}] = \mathbb{E}[c\tau + a_{\mu d} \mathbf{1}_{\{\tau < \infty\}}].$$

This gives the sequential multi-hypothesis testing problem studied by Wald and Wolfowitz [22], Arrow, Blackwell, and Girshick [1]; see also Blackwell and Girshick [5].

5.3. Two alternatives after the change. In this subsection we consider the special case $M = 2$ in which we have only two possible change distributions, $f_1(\cdot)$ and $f_2(\cdot)$. We describe a graphical representation of the stopping and continuation regions for an arbitrary instance of the special case $M = 2$. Then we use this representation to illustrate geometrical properties of the optimal method (Section 4.3) via model instances for certain choices of the model parameters $p_0, p, \nu_1, \nu_2, f_0(\cdot), f_1(\cdot), f_2(\cdot), a_{01}, a_{02}, a_{12}, a_{21}$, and c .

Let the linear mapping $L : \mathbb{R}^3 \mapsto \mathbb{R}^2$ be defined by $L(\pi_0, \pi_1, \pi_2) \triangleq (\frac{2}{\sqrt{3}}\pi_1 + \frac{1}{\sqrt{3}}\pi_2, \pi_2)$. Since $\pi_0 = 1 - \pi_1 - \pi_2$ for every $\pi = (\pi_0, \pi_1, \pi_2) \in S^2 \subset \mathbb{R}^3$, we can recover the preimage π of any




FIGURE 1. Linear mapping L of the standard two-dimensional probability simplex S^2 from the positive orthant of \mathbb{R}^3 into the positive quadrant of \mathbb{R}^2 .

point $L(\pi) \in L(S^2) \subset \mathbb{R}^2$. For every point $\pi = (\pi_0, \pi_1, \pi_2) \in S^2$, the coordinate π_i is given by the Euclidean distance from the image point $L(\pi)$ to the edge of the image triangle $L(S^2)$ that is *opposite* the image point $L(e_i)$, for each $i = 0, 1, 2$. For example, the distance from the image point $L(\pi)$ to the edge of the image triangle opposite the lower-left-hand corner $L(1, 0, 0) = (0, 0)$ is the value of the preimage coordinate π_0 . See Figure 1.

Therefore, we can work with the mappings $L(\Gamma)$ and $L(S^2 \setminus \Gamma)$ of the stopping region Γ and the continuation region $S^2 \setminus \Gamma$, respectively. Accordingly, we depict the decision region for each instance in this subsection using the two-dimensional representation as in the right-hand-side of Figure 1 and we drop the $L(\cdot)$ notation when labeling various parts of each figure to emphasize their source in S^2 .

Each of the examples in this section have the following model parameters in common:

$$p_0 = \frac{1}{50}, \quad p = \frac{1}{20}, \quad \nu_1 = \nu_2 = \frac{1}{2},$$

$$f_0 = \left(\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}\right), \quad f_1 = \left(\frac{4}{10}, \frac{3}{10}, \frac{2}{10}, \frac{1}{10}\right), \quad f_2 = \left(\frac{1}{10}, \frac{2}{10}, \frac{3}{10}, \frac{4}{10}\right).$$

We vary the delay cost and false alarm/isolation costs to illustrate certain geometrical properties of the continuation and stopping regions. See Figures 2, 3, and 4.

Specifically, these examples show instances in which the $M = 2$ convex subsets comprising the optimal stopping region are connected (Figure 2) and instances in which they are not (Figures 3 and 4(a)). Figure 4(b) shows an instance in which the continuation region is disconnected.

Each of the figures in this section have certain features in common. On each subfigure there is a dashed line representing those states $\pi \in S^2$ at which $h_1(\pi) = h_2(\pi)$. Also, each




FIGURE 2. Illustration of *connected* stopping regions and the effects of variation in the false-alarm costs. (a) and (b): $a_{12} = a_{21} = 3, c = 1$. (a): $a_{01} = a_{02} = 10$. (b): $a_{01} = a_{02} = 50$.




FIGURE 3. Illustration of *disconnected* stopping regions and the effects of asymmetric false-isolation costs. (a) and (b): $a_{01} = a_{02} = 10, c = 1$. (a): $a_{12} = a_{21} = 10$. (b): $a_{12} = 16, a_{21} = 4$.

subfigure shows a sample path of $(\Pi_n)_{n=0}^\sigma$ and the realizations of θ and μ for the sample. The shaded area, including its solid boundary, represents the optimal stopping region, while the unshaded area represents the continuation region.

An implementation of the optimal strategy as described in Section 4.3 is as follows: Initialize the statistic $\Pi = (\Pi_n)_{n \geq 0}$ by setting $\Pi_0 = (1 - p_0, p_0 \nu_1, p_0 \nu_2)$ as in part (c) of Proposition 3.2. Use the dynamics of (3.3) to update the statistic Π_n as each observation X_n is realized. Stop taking observations when the statistic Π_n enters the stopping region $\Gamma = \Gamma^{(1)} \cup \Gamma^{(2)}$ for

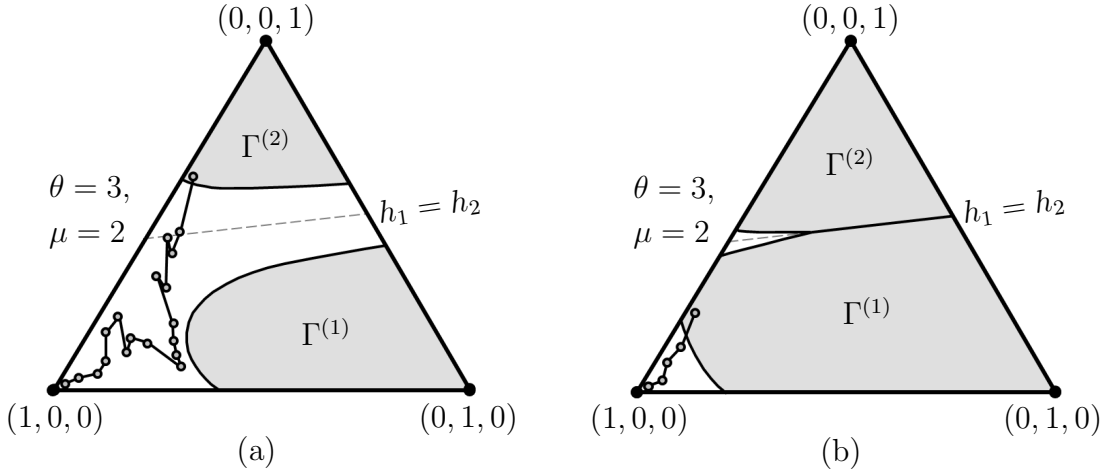


FIGURE 4. Illustration of a *disconnected* continuation region and the effects of variation in the delay cost. (a) and (b): $a_{01} = 14, a_{02} = 20, a_{12} = a_{21} = 8$. (a): $c = 1$. (b): $c = 2$.

the first time, possibly before the first observation is taken (i.e., $n = 0$). The optimal terminal decision is based upon whether the statistic Π_n is in $\Gamma^{(1)}$ or $\Gamma^{(2)}$ upon stopping. Each of the sample paths in Figures 2, 3, and 4 were generated via this algorithm. As Figure 2 shows, the sets $\Gamma^{(1)}$ and $\Gamma^{(2)}$ can intersect on their boundaries and so it is possible to stop in their intersection. In this case, either of the decisions $d = 1$ or $d = 2$ is optimal.

We use value iteration of the optimality equation (4.6) over a fine discretization of S^2 to compute $V_0(\cdot)$ and generate the decision region for each subfigure. Because in the expression $V_0(\pi) = \min\{h(\pi), c(1 - \pi_0) + (\mathbb{T}V_0)(\pi)\}$ the value $V_0(\pi)$ for any fixed initial condition $\Pi_0 = \pi$ on the left depends on the *entire* function $V_0(\cdot)$ on S^M on the right, we have to calculate $V_0(\cdot)$ (or approximate it by $V^N(\cdot)$) on the entire space S^M . The resulting discretized decision region is mapped into the plane via L .

See Bertsekas [4, Chapter 3] for techniques of computing the value function via the optimality equation such as value iteration. Solving the optimality equation by discretizing high-dimensional state-space may not be the best option. Monte Carlo methods based on regression models for the value function seem to scale better as the dimension of the state-space increases; see, for example, Longstaff and Schwartz [15], Tsitsiklis and van Roy [21], Glasserman [10, Chapter 8] for details.

5.4. Three alternatives after the change. In this subsection we consider the special case $M = 3$ in which we have three possible change distributions, $f_1(\cdot)$, $f_2(\cdot)$, and $f_3(\cdot)$. Here, the continuation and stopping regions are subsets of $S^3 \subset \mathbb{R}^4$. Similar to the two-alternatives




FIGURE 5. Illustration of the mapped decision region for an instance of the special case $M = 3$; see also Figure 7 below. A sample path of the process Π is shown in which $\theta = 6$ and $\mu = 3$

case, we introduce the mapping of $S^3 \subset \mathbb{R}^4$ into \mathbb{R}^3 via

$$(\pi_0, \pi_1, \pi_2, \pi_3) \mapsto \left(\sqrt{\frac{3}{2}}\pi_1 + \frac{1}{2}\sqrt{\frac{3}{2}}\pi_2 + \frac{1}{2}\sqrt{\frac{3}{2}}\pi_3, \frac{3}{2}\sqrt{\frac{1}{2}}\pi_2 + \frac{1}{2}\sqrt{\frac{1}{2}}\pi_3, \pi_3 \right).$$

Then we use this representation—actually a rotation of it—to illustrate in Figure 5 an instance with the following model parameters:

$$\begin{aligned} p_0 &= \frac{1}{50}, & p &= \frac{1}{20}, & \nu_1 &= \nu_2 = \nu_3 = \frac{1}{3}, \\ f_0 &= \left(\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}\right), & f_1 &= \left(\frac{4}{10}, \frac{3}{10}, \frac{2}{10}, \frac{1}{10}\right), & f_2 &= \left(\frac{1}{10}, \frac{2}{10}, \frac{3}{10}, \frac{4}{10}\right), & f_3 &= \left(\frac{3}{10}, \frac{2}{10}, \frac{2}{10}, \frac{3}{10}\right) \\ c &= 1, & a_{0j} &= 40, & a_{ij} &= 20, & i, j &= 1, 2, 3. \end{aligned}$$

Note that Figure 5 can be interpreted in a manner similar to the figures of the previous subsection. In this case, for every point $\pi = (\pi_0, \pi_1, \pi_2, \pi_3) \in S^3$, the coordinate π_i is given by the (Euclidean) distance from the image point $L(\pi)$ to the face of the image tetrahedron $L(S^3)$ that is opposite the image corner $L(e_i)$, for each $i = 0, 1, 2, 3$.

5.5. Detection and identification of component failure(s) in a system with suspended animation. Consider a system consisting initially of two working concealed components (labeled 1 and 2) such that upon the failure of either component, the system goes into a state of suspended animation. That is, while both components are still working normally, observations of output of the system have density $f_0(\cdot)$, but upon failure of either component the density of observations changes thereafter (until an alarm is raised) to one of two alternatives: if component 2 fails before component 1, then post-failure observations have density $f_2(\cdot)$, otherwise they have density $f_1(\cdot)$. The problem is to detect quickly when there has been a component failure *and* to identify accurately which component has actually failed based only on sequential observations of output of the system.

Let the random variables

$$\theta := \theta_1 \wedge \theta_2 = \min\{\theta_1, \theta_2\} \quad \text{and} \quad \mu := \begin{cases} 1 & \text{if } \theta_1 \leq \theta_2 \\ 2 & \text{if } \theta_1 > \theta_2 \end{cases}$$

be respectively the time of failure of the first failed component of the system and the corresponding index of this component, where the failure time θ_i of the i th component is a random variable having a geometric distribution with failure probability p_i , $i = 1, 2$. It can be shown easily that when the disorder times θ_1 and θ_2 are independent, the random variable θ has a geometric distribution with failure probability $p := p_1 + p_2 - p_1 p_2$ (or equivalently, θ has a zero-modified geometric distribution with parameters $p_0 = 0$ and p) and that it is independent of the random variable μ , which has distribution $\nu_1 = p_1/p$ and $\nu_2 = 1 - \nu_1$. So although the failure type (i.e., which component has failed) is a function of the failure times of each component, it turns out that this problem fits properly within the Bayesian sequential change diagnosis framework.

This problem can be extended naturally to several components and solved via the technology of Sections 3 and 4. In fact, it can be configured for a variety of scenarios. For example, series-connected components where malfunction of one component suspends immediately the operation of all the remaining components can appear in various electronic relays and multicomponent electronic devices which have fuses to protect the system from the misbehavior of one of its components. Since the system may react differently to diagnostics run by the operators, post-malfunction behavior can differ according to the underlying cause of the malfunction. See Barlow [2, Section 8.4] for background on series systems with suspended animation. Consider also a manufacturing process where we perform a quality test on the final output produced from several processing components. If a component is highly reliable then a geometric distribution with a low failure rate can be a reasonable choice for the lifetime of the component. Moreover, since the typical duration between successive component failures widens over time we can often treat the remaining components as if they enter a state of suspended animation under certain cost structures. That is, we can expect the remaining components to outlive the alarm. For example, suppose that two independent

geometric random variables have expected lifetime of 1000 each. Then the first failure will occur at about time 500 on average, while the second failure will take an additional 1000 periods on average to occur. As illustrated in Figures 2 and 4, respectively, lower false-alarm costs promote raising the alarm earlier, while a higher delay cost discourages waiting for more than relatively few additional periods to raise the alarm.

Specifically, suppose that in a “black box” there are K components whose lifetimes are independent and geometrically distributed. Observations have initially distribution $f_0(\cdot)$ while the system is working, but upon failure of a single component (or simultaneous failure of multiple components), the remaining components enter a state of suspended animation, and the post-failure distribution of observations is determined by the failed component(s). We want to detect the time when at least one of them fails as soon as possible. Moreover, when we raise an alarm we would like to be able to make as accurately as possible diagnoses such as (1) *how many* of the components have actually failed, and (2) *which* ones.

Again, let the failure time θ_k of the k th component be a random variable having a geometric distribution with failure probability p_k , $k \in \mathcal{K} := \{1, 2, \dots, K\}$, and define

$$\theta := \theta_1 \wedge \theta_2 \wedge \dots \wedge \theta_K = \min_{k \in \mathcal{K}} \theta_k$$

as the time when at least one of the K components fails. Let the mapping $\varphi : 2^{\mathcal{K}} \mapsto \{0, 1, \dots\}$ be a nonnegative-integer-valued measure on the discrete σ -algebra $2^{\mathcal{K}}$ of the set $\mathcal{K} = \{1, 2, \dots, K\}$ of component indices, and define the random variable

$$\mu := \varphi(\{k \in \mathcal{K} \mid \theta = \theta_k\})$$

as an index function on the set of indices of the failed components. When the random variables $\theta_1, \dots, \theta_K$ are independent, it can be shown that the random variable θ has a geometric distribution with failure probability

$$p := 1 - \prod_{i \in \mathcal{K}} (1 - p_i)$$

and that it is independent of the random variable μ , which has distribution

$$\nu_k := \frac{1}{p} \sum_{A \in \varphi^{-1}(k)} \prod_{i \in A} p_i \prod_{j \in \mathcal{K} \setminus A} (1 - p_j), \quad k \in \mathcal{M} := \{1, 2, \dots, M := \varphi(\mathcal{K})\}.$$

So, the preceding example of two components corresponds to the special case where $K = 2$ and $\varphi(A) = \min A$ for $A \in \{\{1\}, \{1, 2\}, \{2\}\}$. We can handle the other two aforementioned objectives as follows:

- (1) Let $\varphi(A) = |A|$, $A \in 2^{\mathcal{K}}$. Then the random variable μ represents how many components fail.
- (2) Let $\varphi(A) = \sum_{i \in A} 2^{i-1}$, $A \in 2^{\mathcal{K}}$. Then the mapping φ is one-to-one, the random variable μ takes values in $1, 2, \dots, 2^K - 1$, and the set $\varphi^{-1}(\mu)$ consists of the indices

of the components which fail; i.e., the random variable μ identifies uniquely which components fail.

6. ON THE COMPUTER IMPLEMENTATION OF THE CHANGE-DIAGNOSIS ALGORITHM.

Updating posterior probability process Π online with a computer by using the recursive equations in (3.1) and (3.3) is fast. However, programming a computer to check online whether this process has just entered the optimal stopping region is a challenging task. This is especially so because (i) the critical boundaries of stopping regions do not have known closed-form expressions, and (ii) extensive online computations to determine if one of these boundaries is crossed can take excessive time and defeat the purpose of quickest change detection. Here we outline an implementation strategy that should perform well in general.

The strategy is based on sparse offline representations of critical boundaries between stopping and continuation regions. Suppose that the posterior probability process Π has just been updated to some $\pi = (\pi_0, \pi_1, \dots, \pi_M) \in S^M$. An alarm has to be raised if and only if $\pi \in \Gamma \equiv \Gamma^{(1)} \cup \dots \cup \Gamma^{(M)}$. Checking $\pi \in \Gamma^{(i)}$ for every $i = 1, \dots, M$ (in the worst case) is, however, unnecessary because

$$\pi \in \Gamma \iff \pi \in \Gamma_i \quad \text{if} \quad h_i(\pi) = h(\pi) \equiv \min_{1 \leq j \leq M} h_j(\pi).$$

In other words, one should

- (i) find $i = \arg \min_{1 \leq j \leq M} h_j(\pi)$ first, and
- (ii) raise an alarm and declare that a change of type i has happened if $\pi \in \Gamma^{(i)}$, or
- (iii) wait for at least one more period before raising any alarm otherwise.

Let us suppose that $i = \arg \min_{1 \leq j \leq M} h_j(\pi)$. Checking if $\pi \in \Gamma^{(i)}$ will be fast if both π and $\Gamma^{(i)}$ are represented in terms of polar coordinates, set up locally relative to the corner of S^M confined in the convex set $\Gamma^{(i)}$.

To illustrate the ideas with simple pictures, we will focus on the case that there are $M = 2$ alternatives after the change; see Figure 6. If $\pi = \pi^\circ$ (respectively, $\pi = \pi^\diamond$) as in Figure 6(a), then $h_1(\pi) \leq h_2(\pi)$ and $i = 1$ (respectively, $h_1(\pi) \geq h_2(\pi)$ and $i = 2$). In either case, π can be identified relative to each corner in terms of (i) the Euclidean distance to that corner (denoted by $r_j(\pi)$, $j = 0, 1, 2$) and (ii) one arbitrary but fixed angle (say, by $\beta_j(\pi)$, $j = 0, 1, 2, 3$ indicated on Figure 6(a)) between the line connecting π and the corner and the rays forming the same corner. Every point on the critical boundary of the stopping region $\Gamma^{(j)}$, $j = 1, 2$ admits the same representation. Let us express by $r = g_j(\beta)$ the critical boundary of the stopping region $\Gamma^{(j)}$ in terms of the polar coordinates (β, r) measured locally with respect to the corner of the simplex confined in $\Gamma^{(j)}$ for $j = 1, 2$. Then $\pi^\circ \in \Gamma$ if and only if $r_1(\pi^\circ) \leq g_1(\beta_1(\pi^\circ))$, and $\pi^\diamond \in \Gamma$ if and only if $r_2(\pi^\diamond) \leq g_2(\beta_2(\pi^\diamond))$; see Figures 6(b) and 6(c), respectively.

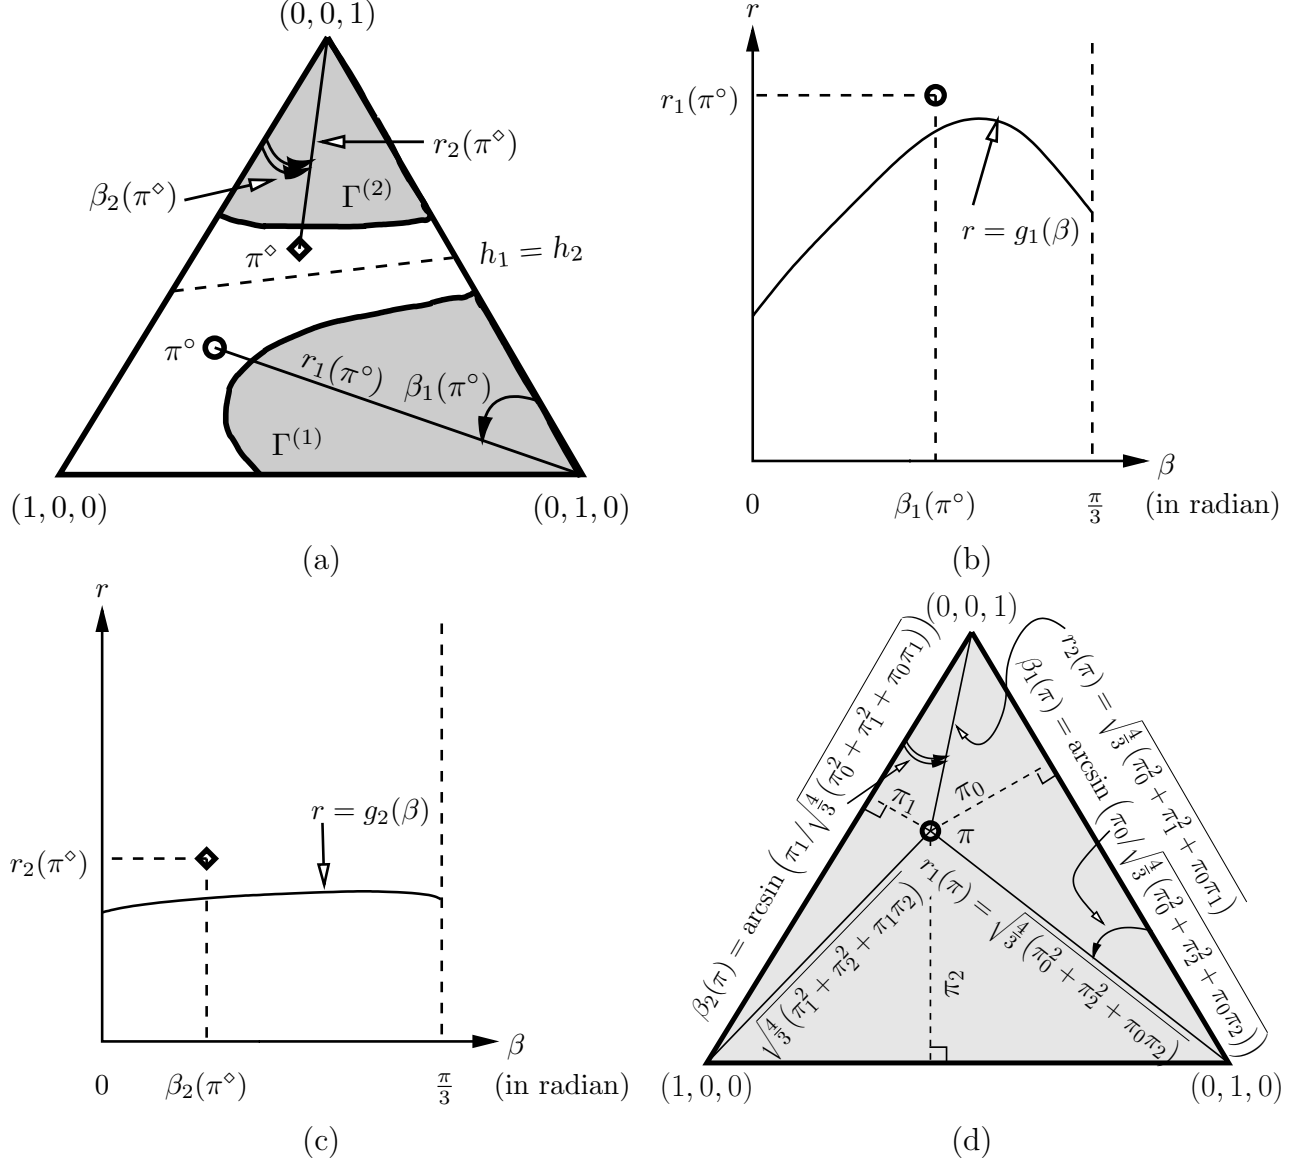


FIGURE 6. For the sample problem displayed in Figure 4(a) ($M = 2$), optimal stopping regions and local polar coordinate systems are shown in (a). The critical boundaries of the stopping regions $\Gamma^{(1)}$ and $\Gamma^{(2)}$ are expressed in terms of local polar coordinates in (b) and (c), respectively. In (d), polar coordinates of π are stated in terms of its Cartesian coordinates. As in Section 5.3, we drop L from $L(\pi^\diamond)$, $L(\Gamma^{(1)})$, $L(1,0,0)$, etc. and simply write π^\diamond , $\Gamma^{(1)}$, $(1,0,0)$. In (a) $h_1(\pi^\diamond) \leq h_2(\pi^\diamond)$ and $h_1(\pi^\diamond) \geq h_2(\pi^\diamond)$.

The plan outlined above works well (i) if the local polar coordinates of π can be identified online quickly, (ii) if the local representations $g_j(\cdot)$, $j = 1, 2$ of the critical boundaries can be stored efficiently to the computer memory, and (iii) if from there they can be retrieved and evaluated fast on demand. Next we will explain how these requirements can be achieved.

Recall from Section 5.3 and Figure 1 that $\pi \in S^2 \subset \mathbb{R}^3$ is embedded into the equilateral triangle $L(S^2) \subset \mathbb{R}^2$ by means of a linear map $\pi \mapsto L(\pi)$. In this natural representation of posterior distributions, $\pi = (\pi_0, \pi_1, \pi_2)$ is mapped to the point $L(\pi) = \left(\frac{2}{\sqrt{3}}\pi_1 + \frac{1}{\sqrt{3}}\pi_2, \pi_2\right)$, whose Euclidean distance to the images $L(1, 0, 0) = (0, 0)$, $L(0, 1, 0) = \left(\frac{2}{\sqrt{3}}, 0\right)$, $L(0, 0, 1) = \left(\frac{1}{\sqrt{3}}, 1\right)$ (corners of the equilateral triangle $L(S^2)$) of $(1, 0, 0)$, $(0, 1, 0)$, $(0, 0, 1)$ are

$$\begin{aligned} r_0(\pi) &\triangleq \|L(1, 0, 0) - L(\pi)\| = \sqrt{\frac{4}{3}(\pi_1^2 + \pi_2^2 + \pi_1\pi_2)}, \\ r_1(\pi) &\triangleq \|L(0, 1, 0) - L(\pi)\| = \sqrt{\frac{4}{3}(\pi_0^2 + \pi_2^2 + \pi_0\pi_2)}, \\ r_2(\pi) &\triangleq \|L(0, 0, 1) - L(\pi)\| = \sqrt{\frac{4}{3}(\pi_0^2 + \pi_1^2 + \pi_0\pi_1)}, \end{aligned}$$

respectively; in a more compact way,

$$(6.1) \quad r_i(\pi) = \sqrt{\frac{4}{3} \left(-\pi_i + \sum_{0 \leq j < k \leq 2} \pi_j \pi_k \right)}, \quad i = 0, 1, 2;$$

see Figure 6(d). Because the Euclidean distance of $L(\pi)$ to the edges opposite to the corners $L(1, 0, 0)$, $L(0, 1, 0)$, $L(0, 0, 1)$ are π_0 , π_1 , and π_2 , respectively, the angles identified in Figure 6(a) can also be calculated easily by

$$\beta_1(\pi) = \arcsin \frac{\pi_0}{\sqrt{\frac{4}{3}(\pi_0^2 + \pi_2^2 + \pi_0\pi_2)}} \quad \text{and} \quad \beta_2(\pi) = \arcsin \frac{\pi_1}{\sqrt{\frac{4}{3}(\pi_0^2 + \pi_1^2 + \pi_0\pi_2)}};$$

or more compactly by

$$(6.2) \quad \beta_i(\pi) = \frac{\pi_{i+2 \bmod 3}}{r_i(\pi)} = \frac{\pi_{i+2 \bmod 3}}{\sqrt{\frac{4}{3} \left(-\pi_i + \sum_{0 \leq j < k \leq 2} \pi_j \pi_k \right)}}, \quad i = 0, 1, 2.$$

Recall that at any $\pi \in S^2$ one has to calculate $\beta_i(\pi)$ and $r_i(\pi)$ only for $i = \arg \min_{1 \leq j \leq 3} h_j(\pi)$ and check if $r_i(\pi) \leq g_i(\beta_i(\pi))$ before raising an alarm.

Unfortunately, an exact/closed-form representation $r = g_j(\beta)$ of the critical boundary of the stopping region $\Gamma^{(j)}$, $j = 1, 2$ in terms of the local polar coordinates (β, r) relative to the corner confined in $\Gamma^{(j)}$ will almost never be available. Instead, only noisy observations (due to the discretization of the state-space S^2 and termination of the value iteration at some finite stage) of that relation can be obtained from the pairs $(\beta_j(\pi), r_j(\pi))$ for every grid-point π on the (approximate) critical boundary of $\Gamma^{(j)}$ for every $j = 1, 2$. Interpolation between those points will certainly give an approximation for $r = g_j(\beta)$ for $j = 1, 2$, but this may waste a lot storage space and computational time during online evaluations, especially when the grid on S^2 is fine. Instead, one can use some statistical smoothing technique to compress the data with minimum loss of information.

Let us suppose that N observations $(\beta^{(k)}, r^{(k)})$, $k = 1, \dots, N$ follow the model $r^{(k)} = g_1(\beta^{(k)}) + \varepsilon^{(k)}$ for every $k = 1, \dots, N$ and that $\varepsilon^{(k)}$, $k = 1, \dots, N$ are i.i.d. random variables with zero mean and some finite common variance. Because $\Gamma^{(1)}$ is convex, the function $\beta \mapsto g_1(\beta)$ is concave, namely, fairly smooth. It may be plausible to approximate it by a cubic spline (twice continuously differentiable piecewise cubic polynomial). The unique curve $\beta \mapsto \widehat{g}_1(\beta)$ that has the minimum penalized sum of squared errors

$$(6.3) \quad S_\lambda(\widehat{g}_1) \triangleq \sum_{k=1}^N [r^{(k)} - \widehat{g}_1(\beta^{(k)})]^2 + \lambda \int_{\mathbb{R}} [\widehat{g}'_1(\beta)]^2 d\beta,$$

for any arbitrary but fixed smoothing parameter $\lambda > 0$, among all twice-differentiable curves is known to exist and belong to the family of cubic splines whose break-points are at β_1, \dots, β_N ; see, for example, de Boor [7], Green and Silverman [11], Ramsay and Silverman [17]. This optimality property and the ability to control the smoothness continuously through λ make cubic splines an attractive candidate for an approximate $g_1(\cdot)$. If the variation of the original curve $\beta \mapsto g_1(\beta)$ is moderate, then the number of break-points $0 \leq K \leq N$ can be taken significantly less than the number of measurements N , and there are $O(K)$ -algorithms that find the cubic spline minimizing (6.3) with the given K break-points; see, for example, Green and Silverman [11, Section 2.3.3] for Reinsch algorithm. Other algorithms represent the solution as a basis-function expansion

$$\widehat{g}_1(\beta) = \sum_{j=1}^{K+3} c_j \Phi_j(\beta)$$

in terms of $K + 3$ spline basis functions $\Phi_1, \dots, \Phi_{K+3}$, and solve the minimization problem in (6.3) by finding the coefficients c_1, \dots, c_{K+3} using multiple-regression; see Green and Silverman [11, Section 3.6], Ramsay and Silverman [17, Section 3.5 and Chapter 5]. Thus, the approximation $\widehat{g}_1(\cdot)$ of $g_1(\cdot)$ can be stored to the computer memory for online use of the change-diagnosis algorithm by means of only $K + 3$ numbers c_1, \dots, c_{K+3} . The basis functions Φ_1, Φ_2, \dots are cubic splines with compact support and can be stored easily and evaluated fast online.

All of the above ideas apply without affecting significantly the online performance of the diagnosis algorithm when the number of alternatives M after change is larger than two. For example, if $M = 3$, then $S^3 \subset \mathbb{R}^4$ is embedded into a tetrahedron $L(S^3) \subseteq \mathbb{R}^3$ by a linear map $\pi \mapsto L(\pi)$ defined in Section 5.4. The Euclidean distance of $L(\pi)$ to the images $L(1, 0, 0, 0)$, $L(0, 1, 0, 0)$, $L(0, 0, 1, 0)$, $L(0, 0, 0, 1)$ of $(1, 0, 0, 0)$, $(0, 1, 0, 0)$, $(0, 0, 1, 0)$, $(0, 0, 0, 1)$ are given

by

$$\begin{aligned} r_0(\pi) &\triangleq \|L(1, 0, 0, 0) - L(\pi)\| = \sqrt{\frac{3}{2} (\pi_1^2 + \pi_2^2 + \pi_3^2 + \pi_1\pi_2 + \pi_1\pi_3 + \pi_2\pi_3)}, \\ r_1(\pi) &\triangleq \|L(0, 1, 0, 0) - L(\pi)\| = \sqrt{\frac{3}{2} (\pi_0^2 + \pi_2^2 + \pi_3^2 + \pi_0\pi_2 + \pi_0\pi_3 + \pi_2\pi_3)}, \\ r_2(\pi) &\triangleq \|L(0, 0, 1, 0) - L(\pi)\| = \sqrt{\frac{3}{2} (\pi_0^2 + \pi_1^2 + \pi_3^2 + \pi_0\pi_1 + \pi_0\pi_3 + \pi_1\pi_3)}, \\ r_3(\pi) &\triangleq \|L(0, 0, 0, 1) - L(\pi)\| = \sqrt{\frac{3}{2} (\pi_0^2 + \pi_1^2 + \pi_2^2 + \pi_0\pi_1 + \pi_0\pi_2 + \pi_1\pi_2)}, \end{aligned}$$

respectively; or more compactly

$$(6.4) \quad r_i(\pi) = \sqrt{\frac{3}{2} \left(-\pi_i + \sum_{0 \leq j \leq k \leq 3} \pi_j \pi_k \right)}, \quad i = 0, 1, 2, 3;$$

see Figure 7. Because the Euclidean distance of $L(\pi)$ to the faces of the tetrahedron opposite to the corners $L(1, 0, 0, 0)$, $L(0, 1, 0, 0)$, $L(0, 0, 1, 0)$, $L(0, 0, 0, 1)$ are π_0 , π_1 , π_2 , and π_3 , respectively, the distance $r_i(\pi)$ and two arbitrary but fixed angles, $\beta_i(\pi) = (\beta_{i1}(\pi), \beta_{i2}(\pi))$, out of three angles defined by

$$(6.5) \quad \arcsin \frac{\pi_j}{r_i(\pi)} = \arcsin \frac{\pi_j}{\sqrt{\frac{3}{2} (-\pi_i + \sum_{0 \leq k \leq \ell \leq 3} \pi_k \pi_\ell)}}, \quad 0 \leq j \leq 3, j \neq i$$

form the local polar coordinates $(\beta_i(\pi), r_i(\pi))$ with respect to the corner of the simplex confined in $\Gamma^{(i)}$, $0 \leq i \leq 3$ and determine $L(\pi)$ uniquely.

The critical boundary between stopping region $\Gamma^{(i)}$, $1 \leq i \leq 3$ and the continuation region can be represented by some concave surface $r = g_i(\beta)$ in terms of the same local polar coordinate system (β, r) just defined above in the vicinity of $\Gamma^{(i)}$, where $\beta = (\beta_1, \beta_2)$ is now a vector. If $(\beta^{(k)}, r^{(k)})$, $k = 1, \dots, N$ are the pairs $(\beta_1(\pi), r_1(\pi))$ evaluated at grid-points π on the approximate boundary of $\Gamma^{(1)}$, then one can fit a *thin plane spline* $\hat{g}_1(\cdot)$, which is twice continuously differentiable and minimizes the penalized sum of squared errors

$$S_\lambda(\hat{g}_1) \triangleq \sum_{k=1}^N [r^{(k)} - \hat{g}_1(\beta^{(k)})]^2 + \lambda \sum_{1 \leq i, j \leq 2} \iint_{\mathbb{R}^2} \left(\frac{\partial^2 \hat{g}_1}{\partial \beta_i \partial \beta_j} \right)^2 (\beta_1, \beta_2) d\beta_1 d\beta_2$$

among all twice-differentiable curves on \mathbb{R}^2 for every arbitrary but fixed smoothing parameter $\lambda > 0$. As before, $\hat{g}_1(\beta) = \sum_{j=1}^{K+3} c_j \Phi_j(\beta)$ admits a basis-function expansion, and the coefficients c_1, \dots, c_{K+3} can be found by using multiple-regression and stored in the computer memory for the online use of change-diagnosis algorithms. See Green and Silverman [11, Chapter 7] for statistical data smoothing in three and higher dimensional Euclidean spaces by using thin plate splines. The similarity of the local polar coordinates (6.1), (6.2)




FIGURE 7. Polar coordinates of π (after transformation by L ; see Section 5.4) in terms of its Cartesian coordinates.

for $M = 2$ and (6.4), (6.5) for $M = 3$ suggest that for general $M \geq 2$ and for a suitable constant $c_M > 0$

$$r_i(\pi) = \sqrt{c_M \left(-\pi_i + \sum_{0 \leq j \leq k \leq M} \pi_j \pi_k \right)}, \quad i = 0, 1, \dots, M$$

and $M - 1$ arbitrary but fixed angles, $\beta_i(\pi) = (\beta_{i,1}, \dots, \beta_{i,M-1})$, out of M angles defined by

$$\arcsin \frac{\pi_j}{r_i(\pi)} = \arcsin \frac{\pi_j}{\sqrt{c_M \left(-\pi_i + \sum_{0 \leq k \leq \ell \leq M} \pi_k \pi_\ell \right)}}, \quad 0 \leq j \leq M, j \neq i$$

form a local polar coordinate system $(\beta_i(\pi), r_i(\pi))$ with respect to the corner of the simplex $S^M \subset \mathbb{R}^M$ confined in stopping region $\Gamma^{(i)}$, $1 \leq i \leq M$ after a suitable linear transformation L into \mathbb{R}^{M-1} .

ACKNOWLEDGMENT

The authors thank two anonymous referees and the associate editor for their careful reading and valuable suggestions that improved the readability of the manuscript. The research of Savas Dayanik was supported partially by the Air Force Office of Scientific Research, under

grant AFOSR-FA9550-06-1-0496 and by the U.S. Department of Homeland Security through the Center for Dynamic Data Analysis for Homeland Security administered through ONR grant number N00014-07-1-0150 to Rutgers University. The research of H. Vincent Poor was supported in part by the U.S. Army Pantheon Project.

APPENDIX A. PROOFS

A.1. Proof of Proposition 3.1. Note that since $\{\tau > n\} \in \mathcal{F}_n$ for every $n \geq 0$, we have

$$\mathbb{E}[(\tau - \theta)^+] = \mathbb{E}\left[\sum_{n=0}^{\infty} \mathbf{1}_{\{\theta \leq n < \tau\}}\right] = \sum_{n=0}^{\infty} \mathbb{E}[\mathbf{1}_{\{\tau > n\}} \mathbb{P}(\theta \leq n | \mathcal{F}_n)] = \mathbb{E}\left[\sum_{n=0}^{\tau-1} (1 - \Pi_n^{(0)})\right].$$

Moreover, for every $j \in \mathcal{M}$, we have $\{\tau = n, d = j\} \in \mathcal{F}_n$, and $\mathbb{E}[\mathbf{1}_{\{d=j\}} \mathbf{1}_{\{\tau < \theta\}}]$ equals

$$\begin{aligned} \sum_{n=0}^{\infty} \mathbb{E}[\mathbf{1}_{\{\tau=n, d=j\}} \mathbf{1}_{\{\theta > n\}}] &= \sum_{n=0}^{\infty} \mathbb{E}[\mathbf{1}_{\{\tau=n, d=j\}} \Pi_n^{(0)}] = \lim_{N \rightarrow \infty} \sum_{n=0}^N \mathbb{E}[\mathbf{1}_{\{\tau=n, d=j\}} \Pi_n^{(0)}] \\ &= \lim_{N \rightarrow \infty} \mathbb{E}\left[\sum_{n=0}^N \mathbf{1}_{\{\tau=n, d=j\}} \Pi_n^{(0)}\right] = \lim_{N \rightarrow \infty} \mathbb{E}[\mathbf{1}_{\{\tau \leq N, d=j\}} \Pi_N^{(0)}] = \mathbb{E}[\mathbf{1}_{\{\tau < \infty, d=j\}} \Pi_{\tau}^{(0)}] \end{aligned}$$

because of the monotone convergence theorem and that $\lim_{N \rightarrow \infty} \{\tau \leq N\} = \cup_{n=1}^{\infty} \{\tau \leq n\} = \{\tau < \infty\}$; see, for example, Ross [18]. Similarly, $\mathbb{E}[\mathbf{1}_{\{d=j, \mu=i\}} \mathbf{1}_{\{\theta \leq \tau < \infty\}}]$ equals

$$\sum_{n=0}^{\infty} \mathbb{E}[\mathbf{1}_{\{\tau=n, d=j\}} \mathbf{1}_{\{\theta \leq n, \mu=i\}}] = \sum_{n=0}^{\infty} \mathbb{E}[\mathbf{1}_{\{\tau=n, d=j\}} \Pi_n^{(i)}] = \mathbb{E}[\mathbf{1}_{\{\tau < \infty, d=j\}} \Pi_{\tau}^{(i)}],$$

for every $i \in \mathcal{M}$. Plugging these expressions into (2.2) completes the proof. \square

A.2. Proof of Proposition 3.2. Parts (a) and (b). Fix any $A = \{(X_1, \dots, X_n) \in B\} \in \mathcal{F}_n$ for some Borel $B \subset E^n$. Then (2.1) implies that

$$(A.1) \quad \mathbb{P}(A) = \int_B m(dx_1) \cdots m(dx_n) \alpha_n(x_1, \dots, x_n)$$

where $\alpha_n(x_1, \dots, x_n) \triangleq \sum_{i=0}^M \alpha_n^{(i)}(x_1, \dots, x_n)$, and

$$\alpha_n^{(i)}(x_1, \dots, x_n) \triangleq \begin{cases} (1-p_0)(1-p)^n \prod_{l=1}^n f_0(x_l), & i=0, \\ p_0 \nu_i \prod_{j=1}^n f_i(x_j) + (1-p_0) p \nu_i \sum_{k=1}^n (1-p)^{k-1} \prod_{l=1}^{k-1} f_0(x_l) \prod_{j=k}^n f_i(x_j), & i \in \mathcal{M}. \end{cases}$$

Hence, $\alpha_n(x_1, \dots, x_n)$ is the joint probability density function of X_1, \dots, X_n with respect to the measure $m(dx_1) \cdots m(dx_n)$. Now for $i \in \mathcal{M}$,

$$\begin{aligned} \int_A \Pi_n^{(i)} d\mathbb{P} &= \mathbb{E} [\mathbf{1}_A \mathbf{1}_{\{\theta \leq n, \mu=i\}}] = \int_B m(dx_1) \cdots m(dx_n) \alpha_n^{(i)}(x_1, \dots, x_n) \\ &= \int_B m(dx_1) \cdots m(dx_n) \alpha_n(x_1, \dots, x_n) \frac{\alpha_n^{(i)}(x_1, \dots, x_n)}{\alpha_n(x_1, \dots, x_n)} = \int_A d\mathbb{P} \frac{\alpha_n^{(i)}(X_1, \dots, X_n)}{\alpha_n(X_1, \dots, X_n)}. \end{aligned}$$

Hence,

$$(A.2) \quad \Pi_n^{(i)} = \frac{\alpha_n^{(i)}(X_1, \dots, X_n)}{\alpha_n(X_1, \dots, X_n)}, \quad i \in \mathcal{M}, \quad \text{and} \quad \Pi_n^{(0)} = \frac{\alpha_n^{(0)}(X_1, \dots, X_n)}{\alpha_n(X_1, \dots, X_n)},$$

since $\sum_{i=0}^M \Pi_n^{(i)} = 1$. Similar considerations also give

$$\mathbb{P}\{\theta = k, \mu = i \mid \mathcal{F}_n\} = \begin{cases} \frac{p_0 \nu_i}{\alpha_n(X_1, \dots, X_n)} \prod_{j=1}^n f_i(X_j), & k = 0, \\ \frac{(1-p_0)(1-p)^{k-1} p \nu_i}{\alpha_n(X_1, \dots, X_n)} \prod_{l=1}^{k-1} f_0(X_l) \prod_{j=k}^n f_i(X_j), & 1 \leq k \leq n, \\ \frac{(1-p_0)(1-p)^{k-1} p \nu_i}{\alpha_n(X_1, \dots, X_n)} \prod_{l=1}^n f_0(X_l), & k \geq n+1. \end{cases}$$

Observe that for $k \geq n+1$,

$$\mathbb{P}\{\theta = k, \mu = i \mid \mathcal{F}_n\} = \frac{\alpha_n^{(0)}(X_1, \dots, X_n)}{\alpha_n(X_1, \dots, X_n)} (1-p)^{k-n-1} p \nu_i = \Pi_n^{(0)} (1-p)^{k-n-1} p \nu_i.$$

In particular, $\mathbb{P}\{\theta = n+1, \mu = i \mid \mathcal{F}_n\} = \Pi_n^{(0)} p \nu_i$, and $\mathbb{P}\{\theta \leq n+1, \mu = i \mid \mathcal{F}_n\}$ equals

$$\mathbb{P}\{\theta \leq n, \mu = i \mid \mathcal{F}_n\} + \mathbb{P}\{\theta = n+1, \mu = i \mid \mathcal{F}_n\} = \Pi_n^{(i)} + \Pi_n^{(0)} p \nu_i.$$

Note also that $\mathbb{P}\{\theta > n+1 \mid \mathcal{F}_n\}$ equals

$$\sum_{k=n+2}^{\infty} \sum_{i=1}^M \Pi_n^{(0)} (1-p)^{k-n-1} p \nu_i = \Pi_n^{(0)} p \sum_{k=n+2}^{\infty} (1-p)^{k-n-1} = \Pi_n^{(0)} (1-p).$$

Thus, $\mathbb{E}[\Pi_{n+1}^{(0)} \mid \mathcal{F}_n] = \mathbb{P}\{\theta > n+1 \mid \mathcal{F}_n\} = \Pi_n^{(0)} (1-p) < \Pi_n^{(0)}$, and

$$\mathbb{E}[\Pi_{n+1}^{(i)} \mid \mathcal{F}_n] = \mathbb{P}\{\theta \leq n+1, \mu = i \mid \mathcal{F}_n\} = \Pi_n^{(i)} + \Pi_n^{(0)} p \nu_i > \Pi_n^{(i)}, \quad i \in \mathcal{M}.$$

Hence, $\{\Pi_n^{(0)}, \mathcal{F}_n\}_{n \geq 0}$ is supermartingale, and $\{\Pi_n^{(i)}, \mathcal{F}_n\}_{n \geq 0}$, $i \in \mathcal{M}$ are submartingales.

For the proof of Part (c), note first that

$$(A.3) \quad \alpha_{n+1}^{(i)}(x_1, \dots, x_{n+1}) = \begin{cases} [\alpha_n^{(i)}(x_1, \dots, x_n) + p \nu_i \alpha_n^{(0)}(x_1, \dots, x_n)] f_i(x_{n+1}), & i \in \mathbb{M}, \\ (1-p) \alpha_n^{(0)}(x_1, \dots, x_n) f_0(x_{n+1}), & i = 0. \end{cases}$$

Substituting these expressions after writing $\Pi_{n+1}^{(i)}$, $i \in \{0\} \cup \mathbb{M}$ by using (A.2), and then dividing both numerator and denominator by $\alpha_n(X_1, \dots, X_n)$ give (3.3).

Next, we find the conditional distribution of X_{n+1} given \mathcal{F}_n for $n \geq 0$. If $g : E \mapsto \mathbb{R}_+$ is a nonnegative function and $A = \{(X_1, \dots, X_n) \in B\} \in \mathcal{F}_n$, then $\int_A \mathbb{E}[g(X_{n+1}) | \mathcal{F}_n] d\mathbb{P}$ equals

$$\begin{aligned} \int_A g(X_{n+1}) d\mathbb{P} &= \int_{B \times E} g(x_{n+1}) \alpha_{n+1}(x_1, \dots, x_{n+1}) m(dx_1) \cdots m(dx_{n+1}) \\ &= \int_B \left[\int_E g(x_{n+1}) \frac{\alpha_{n+1}(x_1, \dots, x_{n+1})}{\alpha_n(x_1, \dots, x_n)} m(dx_{n+1}) \right] \alpha_n(x_1, \dots, x_n) m(dx_1) \cdots m(dx_n) \\ &= \int_A \left[\int_E g(x_{n+1}) \frac{\alpha_{n+1}(X_1, \dots, X_n, x_{n+1})}{\alpha_n(X_1, \dots, X_n)} m(dx_{n+1}) \right] d\mathbb{P}. \end{aligned}$$

Therefore, we have

$$(A.4) \quad \mathbb{E}[g(X_{n+1}) | \mathcal{F}_n] = \int_E g(x) \frac{\alpha_{n+1}(X_1, \dots, X_n, x)}{\alpha_n(X_1, \dots, X_n)} m(dx) = \int_E g(x) D(\Pi_n, x) m(dx),$$

where the second equality follows from (A.2) after substituting (A.3) into previous equality, and the mapping D was defined by (3.1). Then for every nonnegative function $f : S^M \mapsto \mathbb{R}_+$, (3.3) and (A.4) imply that

$$\mathbb{E}[f(\Pi_{n+1}) | \mathcal{F}_n] = \mathbb{E} \left[f \left(\frac{D_0(\Pi_n, X_{n+1})}{D(\Pi_n, X_{n+1})}, \dots, \frac{D_M(\Pi_n, X_{n+1})}{D(\Pi_n, X_{n+1})} \right) \middle| \mathcal{F}_n \right] = (\mathbb{T}f)(\Pi_n)$$

in terms of the operator \mathbb{T} defined by (3.2), and $\mathbb{E}[f(\Pi_{n+1}) | \mathcal{F}_n] = \mathbb{E}[f(\Pi_{n+1}) | \Pi_n]$. Therefore, the process $\{\Pi_n, \mathcal{F}_n; n \geq 0\}$ is Markov, and the proof of part (c) is completed. \square

A.3. Proofs of Lemmas 4.1 and 4.2. Before proving the lemmas, we state Definition A.1, Theorem A.2, and Lemma A.3 from Chow et al. [6, pp. 62-69] for ease of reference.

Definition A.1. A collection $(\xi_t)_{t \in T}$ of random variables is called directed-upwards if, for every $u, v \in T$, there exists $t \in T$ such that $\xi_t \geq \xi_u \vee \xi_v$.

Theorem A.2. If a collection $(\xi_t)_{t \in T}$ of random variables is directed-upwards, then for every $t_0 \in T$, there exists a non-decreasing sequence $(\xi_{t_n})_{n \geq 0}$ in the collection $(\xi_t)_{t \in T}$ such that

$$\operatorname{ess\,sup}_{t \in T} \xi_t = \lim_{n \rightarrow \infty} \uparrow \xi_{t_n} \geq \xi_{t_0} \quad \text{almost surely.}$$

Lemma A.3. For every $n \geq 0$, the collection $\{\mathbb{E}[Y_\tau | \mathcal{F}_n] | \tau \in C_n\}$ is directed-upwards.

Proof of Lemma 4.1. To prove the lemma, we establish two inequalities. Note that $\gamma_n \geq \mathbb{E}[Y_\tau | \mathcal{F}_n]$ for all $\tau \in C_n$ by definition. So, taking expectations, we obtain $\mathbb{E}\gamma_n \geq \sup_{\tau \in C_n} \mathbb{E}Y_\tau = -V_n$. For the reverse direction, by Theorem A.2 and Lemma A.3 there exists a sequence of stopping times $(\tau_k)_{k \geq 1} \subset C_n$ such that $Y_n \leq \mathbb{E}[Y_{\tau_k} | \mathcal{F}_n] \uparrow \gamma_n$ as $k \rightarrow \infty$. So, by the monotone convergence theorem, we have $\mathbb{E}\gamma_n = \mathbb{E}[\lim_{k \rightarrow \infty} \mathbb{E}[Y_{\tau_k} | \mathcal{F}_n]] = \lim_{k \rightarrow \infty} \mathbb{E}Y_{\tau_k} \leq -V_n$. Proof of the equations $-V_n^N = \mathbb{E}\gamma_n^N$, $0 \leq n \leq N$ is similar. \square

Proof of Lemma 4.2. We have $\gamma_n \leq \max\{Y_n, \mathbb{E}[\gamma_{n+1} | \mathcal{F}_n]\}$, because for every fixed $\tau \in C_n$, the expectation $\mathbb{E}[Y_\tau | \mathcal{F}_n]$ equals

$$\begin{aligned} \mathbb{E}[Y_\tau \mathbf{1}_{\{\tau=n\}} + Y_{\tau \vee (n+1)} \mathbf{1}_{\{\tau \geq n+1\}} | \mathcal{F}_n] &= Y_n \mathbf{1}_{\{\tau=n\}} + \mathbf{1}_{\{\tau \geq n+1\}} \mathbb{E}[\mathbb{E}[Y_{\tau \vee (n+1)} | \mathcal{F}_{n+1}] | \mathcal{F}_n] \\ &\leq Y_n \mathbf{1}_{\{\tau=n\}} + \mathbf{1}_{\{\tau \geq n+1\}} \mathbb{E}[\gamma_{n+1} | \mathcal{F}_n] \leq \max\{Y_n, \mathbb{E}[\gamma_{n+1} | \mathcal{F}_n]\}. \end{aligned}$$

For the reverse direction, note that $\gamma_n \geq Y_n = \mathbb{E}[Y_n | \mathcal{F}_n]$ by definition. Since $\gamma_{n+1} = \text{ess sup}_{\tau \in C_{n+1}} \mathbb{E}[Y_\tau | \mathcal{F}_{n+1}]$, by Theorem A.2 and Lemma A.3 there exists a sequence of stopping times $(\tau_k)_{k \geq 1} \subset C_{n+1}$ such that $Y_{n+1} \leq \mathbb{E}[Y_{\tau_k} | \mathcal{F}_{n+1}] \uparrow \gamma_{n+1}$ as $k \rightarrow \infty$. Since $C_{n+1} \subset C_n$ for all $n \geq 0$, we have $\gamma_n \geq \mathbb{E}[Y_{\tau_k} | \mathcal{F}_n] = \mathbb{E}[\mathbb{E}[Y_{\tau_k} | \mathcal{F}_{n+1}] | \mathcal{F}_n]$ for all $k \geq 1$. Taking the limit as $k \rightarrow \infty$ and applying the monotone convergence theorem, we have $\gamma_n \geq \mathbb{E}[\gamma_{n+1} | \mathcal{F}_n]$. Therefore, $\gamma_n \geq \max\{Y_n, \mathbb{E}[\gamma_{n+1} | \mathcal{F}_n]\}$. By a similar argument, we can establish the other equations of Lemma 4.2. \square

A.4. Proof of Lemma 4.3. Because $(C_n^N)_{N \geq n}$ is increasing for every $n \geq 0$, the sequence $(\gamma_n^N)_{N \geq n}$ is increasing for every $n \geq 0$ and has a limit. Set $\gamma'_n = \lim_{N \rightarrow \infty} \gamma_n^N$, $n \geq 0$. Because $\gamma_{n+1}^N \geq Y_{n+1}$ and Y_{n+1} is integrable, taking limits in $\gamma_n^N = \max\{Y_n, \mathbb{E}[\gamma_{n+1}^N | \mathcal{F}_n]\}$, see Lemma 4.2, and monotone convergence give $\gamma'_n = \max\{Y_n, \mathbb{E}[\gamma'_{n+1} | \mathcal{F}_n]\}$ for every $n \geq 0$. Particularly, $(\gamma'_n)_{n \geq 0}$ is an \mathbb{F} -supermartingale.

Obviously, $\gamma'_n \leq \gamma_n$ for every $n \geq 0$. To prove the reverse inequality, it is enough to show that $\gamma'_n \geq \mathbb{E}[Y_\tau | \mathcal{F}_n]$ for every $\tau \in C_n$. Take any $\tau \in C_n$. Then for every $F \in \mathcal{F}_n$ and $m \geq n$

$$\begin{aligned} \int_F \gamma'_n d\mathbb{P} &= \int_{F \cap \{\tau=n\}} \gamma'_\tau d\mathbb{P} + \int_{F \cap \{\tau > n\}} \gamma'_n d\mathbb{P} \geq \int_{F \cap \{\tau=n\}} \gamma'_\tau d\mathbb{P} + \int_{F \cap \{\tau > n\}} \gamma'_{n+1} d\mathbb{P} \\ &= \int_{F \cap \{n \leq \tau \leq n+1\}} \gamma'_\tau d\mathbb{P} + \int_{F \cap \{\tau > n+1\}} \gamma'_{n+1} d\mathbb{P} \geq \dots \geq \int_{F \cap \{n \leq \tau \leq m\}} \gamma'_\tau d\mathbb{P} + \int_{F \cap \{\tau > m\}} \gamma'_m d\mathbb{P}, \end{aligned}$$

where the inequalities follow from \mathbb{F} -supermartingale property of the process $(\gamma'_n)_{n \geq 0}$. Because $\gamma'_k \geq Y_k$ for every $k \geq 0$, we have $\gamma'_\tau \geq Y_\tau$, and for every $m \geq n$

$$\int_F \gamma'_n d\mathbb{P} \geq \int_{F \cap \{n \leq \tau \leq m\}} Y_\tau d\mathbb{P} + \int_{F \cap \{\tau > m\}} \gamma'_m d\mathbb{P} \geq \int_{F \cap \{n \leq \tau \leq m\}} Y_\tau d\mathbb{P} - \int_{F \cap \{\tau > m\}} (\gamma'_m)^- d\mathbb{P}.$$

Since $Y_\tau = -Y_\tau^-$ is integrable and $\tau < \infty$ a.s., we have $\lim_{m \rightarrow \infty} \int_{F \cap \{n \leq \tau \leq m\}} Y_\tau d\mathbb{P} = \int_F Y_\tau d\mathbb{P}$ by dominated convergence, and the proof will be completed if $\lim_{m \rightarrow \infty} \int_{\{\tau > m\}} (\gamma'_m)^- d\mathbb{P} = 0$. However, since $\gamma'_m \geq Y_m$, we have $(\gamma'_m)^- \leq Y_m^-$, and $\int_{\{\tau > m\}} (\gamma'_m)^- d\mathbb{P}$ is less than or equal to

$$\int_{\{\tau > m\}} Y_m^- d\mathbb{P} \leq \int_{\{\tau > m\}} m d\mathbb{P} + \|h\| \mathbb{P}\{\tau > n\} \leq \mathbb{E}\tau \mathbf{1}_{\{\tau > n\}} + \|h\| \mathbb{P}\{\tau > m\},$$

where $\|h\| \triangleq \sup_{\pi \in S_M} |h(\pi)|$. Since $h(\cdot)$ is bounded, $\mathbb{E}\tau < \infty$ and $\mathbb{P}\{\tau < \infty\} = 1$, the right hand side of the last inequality converges to zero as $n \rightarrow \infty$. \square

A.5. Proof of Lemma 4.4. Fix any $N \geq 1$. The equality holds trivially for $n = N$. On the one hand, the definition of the random variable γ_N^N in (4.2) implies that

$$\gamma_N^N = \operatorname{ess\,sup}_{\tau \in C_N^N} \mathbb{E}[Y_\tau \mid \mathcal{F}_N] = \mathbb{E}[Y_N \mid \mathcal{F}_N] = Y_N$$

because $C_N^N \equiv \{N\}$. On the other hand, by the definition of the operator \mathbb{M} in (4.3) we have $\mathbb{M}^0 h \equiv h$, and

$$-c \sum_{k=0}^{N-1} (1 - \Pi_k^{(0)}) - (\mathbb{M}^{N-N} h)(\Pi_n) = -c \sum_{k=0}^{N-1} (1 - \Pi_k^{(0)}) - h(\Pi_n) \equiv Y_N;$$

thanks to (3.5). Therefore, (4.4) holds for $n = N$. Now suppose that (4.4) is true for some $n \geq 1$. Then $\gamma_{n-1}^N = \max\{Y_{n-1}, \mathbb{E}[\gamma_n^N \mid \mathcal{F}_{n-1}]\}$ equals

$$\begin{aligned} & \max \left\{ -c \sum_{k=0}^{n-2} (1 - \Pi_k^{(0)}) - h(\Pi_{n-1}), \mathbb{E} \left[-c \sum_{k=0}^{n-1} (1 - \Pi_k^{(0)}) - (\mathbb{M}^{N-n} h)(\Pi_n) \mid \mathcal{F}_{n-1} \right] \right\} \\ &= -c \sum_{k=0}^{n-2} (1 - \Pi_k^{(0)}) - \min \left\{ h(\Pi_{n-1}), c(1 - \Pi_{n-1}^{(0)}) + (\mathbb{T}(\mathbb{M}^{N-n} h))(\Pi_{n-1}) \right\} \\ &= -c \sum_{k=0}^{n-2} (1 - \Pi_k^{(0)}) - (\mathbb{M}(\mathbb{M}^{N-n} h))(\Pi_{n-1}) = -c \sum_{k=0}^{n-2} (1 - \Pi_k^{(0)}) - (\mathbb{M}^{N-(n-1)} h)(\Pi_{n-1}). \end{aligned}$$

By induction, the equality holds for all $0 \leq n \leq N$. \square

A.6. Proof of Lemma 4.5. Applying Lemma 4.4 for $n = 0$ yields part (a) since

$$(A.5) \quad V_0^N = -\mathbb{E}\gamma_0^N = -\gamma_0^N = (\mathbb{M}^N h)(\Pi_0), \quad N \geq 0.$$

By Lemma 4.3, $\gamma_n = \lim_{N \rightarrow \infty} \gamma_n^N$, and so $V_n = \lim_{N \rightarrow \infty} V_n^N$ by Lemma 4.1 and the dominated convergence. Since the left-hand side of (A.5) converges to V_0 as $N \rightarrow \infty$, the limit of the right-hand side as $N \rightarrow \infty$ exists and $V_0 = \lim_{N \rightarrow \infty} (\mathbb{M}^N h)(\Pi_0)$, which proves part (b). \square

A.7. Proof of Lemma 4.8. Given $\pi, \pi' \in S^M$, $\lambda \in [0, 1]$, and $\lambda' \triangleq 1 - \lambda$, we have

$$\begin{aligned} & \lambda(\mathbb{T}g)(\pi) + \lambda'(\mathbb{T}g)(\pi') = \lambda \int_E m(dx) D(\pi, x) g \left(\frac{D_0(\pi, x)}{D(\pi, x)}, \dots, \frac{D_M(\pi, x)}{D(\pi, x)} \right) \\ & \quad + \lambda' \int_E m(dx) D(\pi', x) g \left(\frac{D_0(\pi', x)}{D(\pi', x)}, \dots, \frac{D_M(\pi', x)}{D(\pi', x)} \right) \\ &= \int_E m(dx) [\lambda D(\pi, x) + \lambda' D(\pi', x)] \left\{ \frac{\lambda D(\pi, x)}{\lambda D(\pi, x) + \lambda' D(\pi', x)} g \left(\frac{D_0(\pi, x)}{D(\pi, x)}, \dots, \frac{D_M(\pi, x)}{D(\pi, x)} \right) \right. \\ & \quad \left. + \frac{\lambda' D(\pi', x)}{\lambda D(\pi, x) + \lambda' D(\pi', x)} g \left(\frac{D_0(\pi', x)}{D(\pi', x)}, \dots, \frac{D_M(\pi', x)}{D(\pi', x)} \right) \right\} \end{aligned}$$

Now, by the concavity of $g(\cdot)$ and the fact that

$$\frac{\lambda D(\pi, x)}{\lambda D(\pi, x) + \lambda' D(\pi', x)} + \frac{\lambda' D(\pi', x)}{\lambda D(\pi, x) + \lambda' D(\pi', x)} = 1$$

is a convex combination, we continue the chain of inequalities to obtain

$$\begin{aligned} \lambda(\mathbb{T}g)(\pi) + \lambda'(\mathbb{T}g)(\pi') &\leq \int_E m(dx) [\lambda D(\pi, x) + \lambda' D(\pi', x)] \\ &\quad \times g\left(\frac{\lambda D_0(\pi, x) + \lambda' D_0(\pi', x)}{\lambda D(\pi, x) + \lambda' D(\pi', x)}, \dots, \frac{\lambda D_M(\pi, x) + \lambda' D_M(\pi', x)}{\lambda D(\pi, x) + \lambda' D(\pi', x)}\right) \\ &= \int_E m(dx) [D(\lambda\pi + \lambda'\pi', x)] g\left(\frac{D_0(\lambda\pi + \lambda'\pi', x)}{D(\lambda\pi + \lambda'\pi', x)}, \dots, \frac{D_M(\lambda\pi + \lambda'\pi', x)}{D(\lambda\pi + \lambda'\pi', x)}\right) \\ &= (\mathbb{T}g)(\lambda\pi + \lambda'\pi'). \end{aligned}$$

Note that the second to last equality follows from the fact that each of D_0, \dots, D_M, D is linear in its first argument. So, we have established that $\mathbb{T}g$ is concave. \square

A.8. Proof of Proposition 4.9. Since $h(\pi) = \min_{j \in \mathcal{M}} \sum_{i=0}^M \pi_i a_{ij}$ is concave, and since the pointwise minimum of two concave functions is concave, by Lemma 4.8 the function $(\mathbb{M}f)(\pi) = \min\{h(\pi), c(1 - \pi_0) + (\mathbb{T}f)(\pi)\}$ is concave for every bounded concave $f : S^M \mapsto \mathbb{R}$. Therefore, $\mathbb{M}h, \mathbb{M}^2h, \dots$ are concave, and V_0^0, V_0^1, \dots are concave by Lemma 4.5(a). This proves part (a). For part (b), note that Lemma 4.5(b) implies that $V_0(\pi) = \lim_{N \rightarrow \infty} (\mathbb{M}^N h)(\pi)$ for every $\pi \in S^M$; thus, $V_0(\cdot)$ is concave on S^M . \square

A.9. Proof of Proposition 4.10. The inequality $-V_0(\pi) \geq -V_0^N(\pi)$ for every $\pi \in S^M$ and $N \geq 1$ is obvious. Let us prove the second. Fix $N \geq 1$, $\pi \in S^M$, and any $\varepsilon > 0$. Since

$$0 \geq -V_0(\pi) = \sup_{\tau \in C_0} \mathbb{E}_\pi Y_\tau \geq \mathbb{E}_\pi Y_0 \geq -\|h\| > -\infty$$

is finite, there exists some stopping time $\tau_\varepsilon \in C_0$ such that

$$(A.6) \quad -V_0(\pi) - \varepsilon < \mathbb{E}_\pi Y_{\tau_\varepsilon} = \mathbb{E}_\pi \left[-c \sum_{k=0}^{\tau_\varepsilon - 1} (1 - \Pi_k^{(0)}) - h(\Pi_{\tau_\varepsilon}) \right].$$

Observe that $\tau_\varepsilon \wedge N \in C_0^N$ and

$$(A.7) \quad \begin{aligned} -V_0^N(\pi) &\geq \mathbb{E}_\pi Y_{\tau_\varepsilon \wedge N} \geq \mathbb{E}_\pi \left[-c \sum_{k=0}^{\tau_\varepsilon - 1} (1 - \Pi_k^{(0)}) - h(\Pi_{\tau_\varepsilon}) \right] - \|h\| \mathbb{P}_\pi\{\tau_\varepsilon \geq N\} \\ &\geq -V_0(\pi) - \varepsilon - \frac{\|h\|}{N} \mathbb{E}_\pi \tau_\varepsilon. \end{aligned}$$

The last inequality follows by the Markov inequality applied to $\mathbb{P}_\pi\{\tau_\varepsilon \geq N\}$ and since τ_ε is ε -optimal for V_0 . Next, we will bound $\mathbb{E}_\pi \tau_\varepsilon$ from above by using (A.6):

$$\begin{aligned} -\varepsilon - V_0(\pi) &< \mathbb{E}_\pi \left[-c \sum_{k=0}^{\tau_\varepsilon-1} (1 - \Pi_k^{(0)}) - h(\Pi_{\tau_\varepsilon}) \right] \leq \mathbb{E}_\pi \left[-c \sum_{k=0}^{\tau_\varepsilon-1} (1 - \Pi_k^{(0)}) \right] \\ &= -c \mathbb{E}_\pi \tau_\varepsilon + c \mathbb{E}_\pi \sum_{k=0}^{\tau_\varepsilon-1} \Pi_k^{(0)} \leq -c \mathbb{E}_\pi \tau_\varepsilon + c \mathbb{E}_\pi \sum_{k=0}^{\infty} \Pi_k^{(0)} = -c \mathbb{E}_\pi \tau_\varepsilon + c \sum_{k=0}^{\infty} \mathbb{E}_\pi \Pi_k^{(0)}. \end{aligned}$$

Rearrangement after using the inequality $\mathbb{E}_\pi \Pi_k^{(0)} \leq (1-p)^k$ of Proposition 3.2(a) gives

$$\mathbb{E}_\pi \tau_\varepsilon \leq \frac{1}{c} [V_0(\pi) + \varepsilon] + \frac{1}{p} \leq \frac{\|h\| + \varepsilon}{c} + \frac{1}{p}.$$

Now using this bound on $\mathbb{E}_\pi \tau_\varepsilon$ in (A.7) we have

$$-V_0^N(\pi) \geq -V_0(\pi) - \varepsilon - \frac{\|h\|}{N} \left(\frac{\|h\| + \varepsilon}{c} + \frac{1}{p} \right).$$

However, ε was arbitrary, so taking the limit as $\varepsilon \downarrow 0$ we obtain the desired bound. \square

A.10. Proof of Proposition 4.11. Recall that $V_0^0(\pi) = (\mathbb{M}^0 h)(\pi) = h(\pi) = \min_{j \in \mathcal{M}} \sum_{i=0}^M \pi_i a_{ij}$, which is continuous in $\pi \in S^M$. Suppose that $V_0^N : S^M \mapsto \mathbb{R}_+$ is continuous for some $N \geq 0$. Then by (A.5)

$$(A.8) \quad V_0^{N+1}(\pi) = (\mathbb{M}^{N+1} h)(\pi) = (\mathbb{M} V_0^N)(\pi) = \min \{h(\pi), c(1 - \pi_0) + (\mathbb{T}V_0^N)(\pi)\},$$

where (see (3.2))

$$(A.9) \quad (\mathbb{T}V_0^N)(\pi) = \int_E m(dx) D(\pi, x) V_0^N \left(\frac{D_0(\pi, x)}{D(\pi, x)}, \dots, \frac{D_M(\pi, x)}{D(\pi, x)} \right).$$

Note that

- the mapping $\pi \mapsto D(\pi, x)$ is continuous for every $x \in E$,
- for every $x \in E$ such that $D(\pi, x) > 0$ (these are the x -values that matter in the defining integral of $(\mathbb{T}V_0^N)(\pi)$ above), the coordinates, $\frac{D_0(\pi, x)}{D(\pi, x)}, \dots, \frac{D_M(\pi, x)}{D(\pi, x)}$, are continuous,
- since $V_0^N(\cdot)$ is continuous on S^M by the induction hypothesis, the integrand in (A.9) is continuous in π for every fixed $x \in E$ such that $D(\pi, x) > 0$,
- since $0 \leq V_0^N(\cdot) \leq \|h\|$, the same nonnegative integrand is bounded from above by the integrable function $2\|h\| \sum_{i=0}^M f_i(x)$ for every $\pi \in S^M$,
- then the mapping $\pi \mapsto (\mathbb{T}V_0^N)(\pi)$ is continuous by dominated convergence,
- and finally, since $h(\pi)$ and $c(1 - \pi_0) + (\mathbb{T}V_0^N)(\pi)$ are continuous, (A.8) implies that the mapping $\pi \mapsto V_0^{N+1}(\pi)$ is continuous.

Hence, continuity holds for every $N \geq 0$ by induction, and this completes the proof. \square

A.11. Proof of Corollary 4.12. The function $V_0(\pi)$ on the compact space S^M is the limit of the sequence $\{V_0^N(\pi)\}_{N \geq 0}$ of continuous functions, *uniformly* in $\pi \in S^M$ by Proposition 4.10. Therefore, it is continuous. \square

A.12. Proof of Theorem 4.13. By Lemmas 4.4 and 4.1 we have that $(V_0^N)_{N \geq 0}$ is a non-increasing sequence of functions, bounded from above by the function h . Since $h(\cdot)$ and $h_j(\cdot), j \in \mathcal{M}$ are continuous and since $V_0^N(\cdot), N \geq 0$ are continuous on S^M by Proposition 4.11, the set $\Gamma_N^{(j)} = \{\pi \in S^M \mid V_0(\pi) = h(\pi) = h_j(\pi)\}$ is a closed subset of S^M for each $N \geq 0$ and $j \in \mathcal{M}$.

Fix $j \in \mathcal{M}$. Then $V_0^{N+1}(\pi) = h(\pi) = h_j(\pi)$ implies $V_0^N(\pi) = h(\pi) = h_j(\pi)$; and therefore, $\Gamma_{N+1}^{(j)} \subset \Gamma_N^{(j)}$ for every $N \geq 0$. Hence, $(\Gamma_N^{(j)})_{N \geq 0}$ is a non-increasing sequence of closed subsets of S^M . Clearly, $\Gamma_N = \bigcup_{j=1}^M \Gamma_N^{(j)}, N \geq 0$ and $(\Gamma_N)_{N \geq 0}$ is also a non-increasing sequence of closed subsets of S^M . Moreover, since $V_0^N \searrow V_0$ by Proposition 4.10, the limit of the non-increasing sequence $(\Gamma_N)_{N \geq 0}$ is Γ ; i.e., $\bigcap_{N=1}^{\infty} \Gamma_N = \Gamma$. Similarly, $\bigcap_{N=1}^{\infty} \Gamma_N^{(j)} = \Gamma^{(j)}, j \in \mathcal{M}$.

Given $\pi \in S^M$, if the inequality $h_j(\pi) \leq \min\{h(\pi), c(1 - \pi_0)\}$ holds, then $h_j(\pi) \leq h(\pi)$, which implies that $h_j(\pi) = h(\pi)$. Also,

$$h_j(\pi) \leq \min\{h(\pi), c(1 - \pi_0) + (\mathbb{T}V_0)(\pi)\} = V_0(\pi).$$

This follows from the fact that $V_0 \geq 0$ implies $\mathbb{T}V_0 \geq 0$ and from the optimality equation of Proposition 4.6. But, since $V_0 \leq h$ on S^M , we have $V_0(\pi) = h_j(\pi) = h(\pi)$ and thus $\pi \in \Gamma^{(j)}$. As a corollary, since $h_j(e_j) = 0 \leq \min\{h(e_j), c\}$, we have $e_j \in \Gamma^{(j)}$.

In order to prove the convexity of $\Gamma_N^{(j)}$, take $\pi, \pi' \in \Gamma_N^{(j)}$ and show that $\lambda\pi + (1 - \lambda)\pi' \in \Gamma_N^{(j)}$ for every $\lambda \in [0, 1]$. Since $V_0^N(\cdot)$ is concave by Proposition 4.9, we have

$$\begin{aligned} \lambda V_0^N(\pi) + (1 - \lambda)V_0^N(\pi') &\leq V_0^N(\lambda\pi + (1 - \lambda)\pi') \leq h(\lambda\pi + (1 - \lambda)\pi') \leq h_j(\lambda\pi + (1 - \lambda)\pi') \\ &= \lambda h_j(\pi) + (1 - \lambda)h_j(\pi') = \lambda V_0^N(\pi) + (1 - \lambda)V_0^N(\pi'). \end{aligned}$$

Therefore, since $V_0^N(\pi) \leq h(\pi), \pi \in S^M$, we have

$$V_0^N(\lambda\pi + (1 - \lambda)\pi') = h(\lambda\pi + (1 - \lambda)\pi') = h_j(\lambda\pi + (1 - \lambda)\pi')$$

and $\lambda\pi + (1 - \lambda)\pi' \in \Gamma_N \cap \{\pi \in S^M \mid h(\pi) = h_j(\pi)\} = \Gamma_N^{(j)}$. Hence, $\Gamma_N^{(j)}$ is convex. Since an intersection of convex sets is again convex, $\Gamma^{(j)} = \bigcap_{N=1}^{\infty} \Gamma_N^{(j)}$ is convex.

Thus, we have shown that $\Gamma = \bigcup_{i=1}^M \Gamma^{(i)}$ is the union of M non-empty closed convex subsets of S^M . Finally, consider $\pi(\lambda) \triangleq \lambda e_0 + (1 - \lambda)e_j$ for $\lambda \in (0, \frac{c}{a_{0j} + c}]$. Note that $c > 0$ and $a_{0j} \geq 0$ imply that the interval $(0, \frac{c}{a_{0j} + c}]$ is non-empty. The inequality $\lambda \leq \frac{c}{a_{0j} + c}$ implies that $c(1 - \lambda) \geq \lambda a_{0j} = h_j(\pi(\lambda))$. Hence, $h(\pi(\lambda)) \leq h_j(\pi(\lambda)) \leq c(1 - \lambda) \leq c(1 - \lambda) + (\mathbb{T}V_0)(\pi(\lambda))$ and so $V_0(\pi(\lambda)) = h(\pi(\lambda))$ by Proposition 4.6. Therefore, $\Gamma \ni \pi(\lambda) \notin \{e_1, \dots, e_M\}$. \square

A.13. Proof of Lemma 4.14. For every $n \geq 0$, the limit $\lim_{N \rightarrow \infty} \gamma_n^N$ exists a.s. by Lemma 4.3. So, fix n and take the limit as $N \rightarrow \infty$ of the expression in Lemma 4.4. Then apply Lemma 4.5(b) to obtain the result. \square

A.14. **Proof of Theorem 4.15.** Let us prove part (a) first. Note that

$$\sigma = \inf\{n \geq 0 \mid \Pi_n \in \Gamma\} = \inf\{n \geq 0 \mid V_0(\Pi_n) = h(\Pi_n)\} = \inf\{n \geq 0 \mid \gamma_n = Y_n\}.$$

The second equality follows from the definition of Γ and the last equality follows from Lemma 4.14 and the definition of Y_n (3.5). Now, fix n and recall from Lemma 4.2 that $\gamma_n = \max\{Y_n, \mathbb{E}[\gamma_{n+1} \mid \mathcal{F}_n]\}$. Then $\gamma_n = \mathbb{E}[\gamma_{n+1} \mid \mathcal{F}_n]$ on $\{\sigma > n\}$. So,

$$\begin{aligned} \mathbb{E}[\gamma_{(n+1) \wedge \sigma} \mid \mathcal{F}_n] &= \mathbb{E}[\gamma_\sigma \mathbf{1}_{\{\sigma \leq n\}} \mid \mathcal{F}_n] + \mathbb{E}[\gamma_{n+1} \mathbf{1}_{\{\sigma > n\}} \mid \mathcal{F}_n] \\ &= \gamma_\sigma \mathbf{1}_{\{\sigma \leq n\}} + \mathbf{1}_{\{\sigma > n\}} \mathbb{E}[\gamma_{n+1} \mid \mathcal{F}_n] = \gamma_\sigma \mathbf{1}_{\{\sigma \leq n\}} + \gamma_n \mathbf{1}_{\{\sigma > n\}} = \gamma_{n \wedge \sigma}. \end{aligned}$$

This establishes the martingale property of the stopped process $\{\gamma_{n \wedge \sigma}, \mathcal{F}_n\}_{n \geq 0}$.

To prove part (b), we use part (a) and Lemma 4.1 to write

$$-V_0 = \sup_{\tau \in C_0} \mathbb{E}Y_\tau = \gamma_0 = \mathbb{E}[\gamma_{n \wedge \sigma}] = \mathbb{E}[Y_\sigma \mathbf{1}_{\{\sigma \leq n\}}] + \mathbb{E}[\gamma_n \mathbf{1}_{\{\sigma > n\}}].$$

Since $Y_n = -\sum_{k=0}^{n-1} c(1 - \Pi_k^{(0)}) - h(\Pi_n) \leq 0$ for every n , we can use Fatou's Lemma after taking $\limsup_{n \rightarrow \infty}$ of both sides to obtain

$$(A.10) \quad -V_0 \leq \mathbb{E}[Y_\sigma \mathbf{1}_{\{\sigma < \infty\}}] + \mathbb{E}\left[\left(\limsup_{n \rightarrow \infty} \gamma_n\right) \mathbf{1}_{\{\sigma = \infty\}}\right].$$

Since $\limsup_{n \rightarrow \infty} \gamma_n \leq \limsup_{n \rightarrow \infty} -\sum_{k=0}^{n-1} c(1 - \Pi_k^{(0)}) = -\infty$ by Remark 3.3, and $-V_0 > -h > -\infty$, the inequality (A.10) implies that $\mathbb{P}\{\sigma = \infty\} = 0$. Therefore, the same inequality becomes $-V_0 \equiv \sup_{\tau} \mathbb{E}Y_\tau \leq \mathbb{E}Y_\sigma$. To show that σ is optimal, we must prove that $\sigma \in C_0$. Since $\sigma < \infty$ a.s., it is enough to show $\mathbb{E}Y_\sigma^- < \infty$, which is equivalent to showing that $\mathbb{E}\sigma < \infty$ by the discussion before equation (3.7).

However, since $\mathbb{E}Y_\sigma \geq -V_0 > -\infty$, we also have $\mathbb{E}\sigma < \infty$. Indeed,

$$\begin{aligned} -\infty < \mathbb{E}Y_\sigma &= \mathbb{E}\left[-\sum_{k=0}^{\sigma-1} c(1 - \Pi_k^{(0)}) - h(\Pi_\sigma)\right] \leq -c\mathbb{E}\sigma + c\mathbb{E}\left[\sum_{k=0}^{\infty} \Pi_k^{(0)}\right] \\ &= -c\mathbb{E}\sigma + c\sum_{k=0}^{\infty} \mathbb{E}\Pi_k^{(0)} \leq -c\mathbb{E}\sigma + c\sum_{k=0}^{\infty} (1-p)^k = -c\mathbb{E}\sigma + \frac{c}{p} \end{aligned}$$

implies $\mathbb{E}\sigma < \infty$. Here, the last inequality follows from Proposition 3.2(a). This completes the proofs of parts (b) and (c). \square

REFERENCES

- [1] K. J. Arrow, D. Blackwell, and M. A. Girshick. Bayes and minimax solutions of sequential decision problems. *Econometrica*, 17:213–244, 1949.
- [2] R. E. Barlow. *Engineering Reliability*. ASA-SIAM Series on Statistics and Applied Probability. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1998.
- [3] M. Basseville and I. V. Nikiforov. *Detection of Abrupt Changes: Theory and Application*. Prentice Hall Information and System Sciences Series. Prentice Hall Inc., Englewood Cliffs, NJ, 1993.

- [4] D. P. Bertsekas. *Dynamic Programming and Optimal Control. Vol. II*. Athena Scientific, Belmont, MA, second edition, 2001.
- [5] D. Blackwell and M. A. Girshick. *Theory of Games and Statistical Decisions*. Dover Publications Inc., New York, 1979. Reprint of the 1954 edition.
- [6] Y. S. Chow, H. Robbins, and D. Siegmund. *Great Expectations: The Theory of Optimal Stopping*. Houghton Mifflin Co., Boston, Mass., 1971.
- [7] C. de Boor. *A Practical Guide to Splines*, volume 27 of *Applied Mathematical Sciences*. Springer-Verlag, New York, revised edition, 2001.
- [8] V. P. Dragalin, A. G. Tartakovsky, and V. V. Veeravalli. Multihypothesis sequential probability ratio tests. I. Asymptotic optimality. *IEEE Trans. Inform. Theory*, 45(7):2448–2461, 1999.
- [9] V. P. Dragalin, A. G. Tartakovsky, and V. V. Veeravalli. Multihypothesis sequential probability ratio tests. II. Accurate asymptotic expansions for the expected sample size. *IEEE Trans. Inform. Theory*, 46(4):1366–1383, 2000.
- [10] P. Glasserman. *Monte Carlo Methods in Financial Engineering*, volume 53 of *Applications of Mathematics (New York)*. Springer-Verlag, New York, 2004. , Stochastic Modelling and Applied Probability.
- [11] P. J. Green and B. W. Silverman. *Nonparametric Regression and Generalized Linear Models*, volume 58 of *Monographs on Statistics and Applied Probability*. Chapman & Hall, London, 1994. A roughness penalty approach.
- [12] S. A. Klugman, H. H. Panjer, and G. E. Willmot. *Loss Models*. Wiley Series in Probability and Statistics: Applied Probability and Statistics. John Wiley & Sons Inc., New York, 1998. From data to decisions, With the assistance of Gary G. Venter, A Wiley-Interscience Publication.
- [13] T. L. Lai. Sequential multiple hypothesis testing and efficient fault detection-isolation in stochastic systems. *IEEE Trans. Inform. Theory*, 46(2):595–608, 2000.
- [14] T. L. Lai. Sequential analysis: some classical problems and new challenges. *Statist. Sinica*, 11(2):303–408, 2001. With comments and a rejoinder by the author.
- [15] F. A. Longstaff and E. S. Schwartz. Valuing American options by simulation: A simple least-squares approach. *Review of Financial Studies*, 14(1):113–147, 2001.
- [16] I. V. Nikiforov. A generalized change detection problem. *IEEE Trans. Inform. Theory*, 41(1):171–187, 1995.
- [17] J. O. Ramsay and B. W. Silverman. *Functional Data Analysis*. Springer Series in Statistics. Springer, New York, second edition, 2005.
- [18] S. M. Ross. *Stochastic Processes*. Wiley Series in Probability and Mathematical Statistics: Probability and Mathematical Statistics. John Wiley & Sons Inc., New York, 1983. Lectures in Mathematics, 14.
- [19] A. N. Shiryaev. Optimal methods in quickest detection problems. *Teor. Veroyatnost. i Primenen.*, 8:26–51, 1963.
- [20] A. N. Shiryaev. *Optimal Stopping Rules*. Springer-Verlag, New York, 1978. Translated from the Russian by A. B. Aries, Applications of Mathematics, Vol. 8.
- [21] J. N. Tsitsiklis and V. B. Roy. Regression methods for pricing complex American-style options. *IEEE Transactions on Neural Networks*, 12:694–703, July 2001.
- [22] A. Wald and J. Wolfowitz. Bayes solutions of sequential decision problems. *Ann. Math. Statistics*, 21:82–99, 1950.

(S. Dayanik and C. Goulding) DEPARTMENT OF OPERATIONS RESEARCH AND FINANCIAL ENGINEERING,
AND THE BENDHEIM CENTER FOR FINANCE, PRINCETON UNIVERSITY, PRINCETON, NJ 08544

E-mail address: `sdayanik@princeton.edu, cgouldin@princeton.edu`

(H. V. Poor) SCHOOL OF ENGINEERING AND APPLIED SCIENCE, PRINCETON UNIVERSITY, PRINCETON,
NJ 08544

E-mail address: `poor@princeton.edu`